

# THÈSE

PRÉSENTÉE A

**L'UNIVERSITÉ PIERRE ET MARIE CURIE**

ÉCOLE DOCTORALE : Cerveau Cognition Comportement (ED n°158)

Par Florent MEYNIEL

POUR OBTENIR LE GRADE DE

DOCTEUR

SPÉCIALITÉ : Neurosciences cognitives

## **How the human brain allocates physical effort over time: evidence from behavior, neuroimaging and pharmacology**

Directeur de recherche : Mathias PESSIGLIONE

Soutenue le : 11 septembre 2013

Devant la commission d'examen formée de :

M. Mathias PESSIGLIONE	CR	Inserm-ICM	Directeur
M. Christopher SUMMERFIELD	PI	University of Oxford	Rapporteur
M. Philippe TOBLER	PR	University of Zurich	Rapporteur
M. Stanislas DEHAENE	DR	Inserm-CEA	Examineur
M. Predrag PETROVIC	MCU-PH	Karolinska Institutet	Examineur
M. Habib BENALI	DR	Inserm-UPMC	Président



1	Forewords.....	7
1.1	Publication list .....	7
1.2	How to read this manuscript .....	7
1.3	Acknowledgements .....	8
1.4	Résumé français .....	10
2	Introduction.....	13
2.1	Effort allocation over time .....	15
2.1.1	A very simple question .....	15
2.1.2	Basic intuitions .....	16
2.2	The economic view: decision aims at maximizing benefit against cost .....	19
2.2.1	Value-based decision making .....	19
2.2.1.1	Utility: cost traded against benefit.....	19
2.2.1.2	Brain Valuation System .....	22
2.2.1.3	From valuation to decision and learning.....	26
2.2.1.4	Different decision systems .....	35
2.2.2	Effort-based decision making.....	37
2.2.2.1	Physical effort is treated as cost in decision .....	37
2.2.2.2	There are different kinds of effort-related cost .....	40
2.2.3	Summary.....	42
2.3	How efforts are produced .....	43
2.3.1	The motor command.....	43
2.3.1.1	The exercising muscle.....	43
2.3.1.2	The motor drive.....	49
2.3.1.3	Regulation of the central command.....	51
2.3.1.4	Motor preparation.....	53
2.3.2	The motivation to drive the effort .....	56
2.3.2.1	What is motivation? .....	56
2.3.2.2	Incentive motivation of motor behavior in normal condition .....	58
2.3.2.3	Pathological impairment of motor incentive motivation.....	62
2.3.3	Summary.....	64
2.4	How time steps in.....	65
2.4.1	Benefit is not constant over time but discounted by delay .....	65
2.4.1.1	Something special about time .....	65
2.4.1.2	Time, the opportunity cost and lessons from the foraging theories .....	68
2.4.2	Cost is not constant over time but increased by fatigue.....	70
2.4.2.1	What is fatigue?.....	70
2.4.2.2	Where does fatigue come from?.....	71
2.4.2.3	Is fatigue related to perception.....	76

2.4.3	Summary.....	78
2.5	Existing models.....	79
2.5.1	When to stop exertion?.....	79
2.5.1.1	The optimal motor control theory and the movement duration.....	79
2.5.1.2	Foraging theories.....	80
2.5.1.3	Sport literature and the athletic performance.....	82
2.5.2	When to start the exertion?.....	85
2.5.2.1	The opportunity cost .....	85
2.5.2.2	The diving problem.....	87
2.5.3	How to set the work-to-rest ratio? .....	87
2.5.3.1	Economics and the allocation of time .....	87
2.5.3.2	Behavioral ecology and the allocation of time .....	88
2.5.4	Summary.....	89
2.6	A proposal.....	91
2.6.1	The scope.....	91
2.6.2	The model.....	92
2.6.3	Outline of the work .....	94
2.6.3.1	Behavior: what the effort allocation is adapted to .....	95
2.6.3.2	Neuroimaging: functional anatomy of the effort allocation .....	95
2.6.3.3	Manipulation of the brain and effort allocation .....	96
3	Experimental work .....	97
3.1	How is effort allocated over time? Evidence from behavioral data.....	99
3.1.1	Decisions to stop and resume effort are based on the on-line level of cost-evidence. 99	
3.1.1.1	Introduction.....	99
3.1.1.2	Results .....	99
3.1.1.3	Discussion .....	104
3.1.1.4	Methods .....	105
3.1.2	Strategic effects in cost-evidence based effort allocation: when the effort difficulty and the incentive level matter .....	108
3.1.2.1	Introduction.....	108
3.1.2.2	Results .....	109
3.1.2.3	Discussion .....	115
3.1.2.4	Methods .....	116
3.1.3	Computational account of effort allocation.....	119
3.1.3.1	Introduction.....	119
3.1.3.2	Results .....	119
3.1.3.3	Discussion .....	123
3.1.3.4	Methods .....	125

3.1.4	Limitation of introspection.....	128
3.1.4.1	Introduction.....	128
3.1.4.2	Results .....	129
3.1.4.3	Discussion .....	131
3.1.4.4	Method.....	133
3.1.5	Behavioral results: general discussion .....	136
3.1.5.1	Summary.....	136
3.1.5.2	Comments .....	137
3.1.5.3	What we still do not know and how the neuroimaging and manipulation studies could help.....	140
3.2	How does the nervous system solve the effort allocation problem? .....	143
3.2.1	Neuroimaging data support the cost-evidence model of effort allocation .....	143
3.2.1.1	Introduction.....	143
3.2.1.2	Results .....	144
3.2.1.3	Discussion .....	150
3.2.1.4	Methods .....	153
3.2.2	Get back to work: A mechanism to translate incentives into speeded effort resumption 157	
3.2.2.1	Introduction.....	157
3.2.2.2	Results .....	159
3.2.2.3	Discussion .....	165
3.2.2.4	Methods .....	168
3.2.3	Adaptation of the body activation system .....	170
3.2.3.1	Introduction.....	170
3.2.3.2	Results .....	172
3.2.3.3	Discussion .....	173
3.2.3.4	Methods .....	175
3.3	Manipulation of the brain during effort allocation .....	177
3.3.1	Manipulation of nociception with paracetamol.....	177
3.3.1.1	Introduction.....	177
3.3.1.2	Results .....	178
3.3.1.3	Discussion .....	179
3.3.1.4	Methods .....	181
3.3.2	Manipulation of expectations with motivational and analgic hypnotic suggestion ..	182
3.3.2.1	Introduction.....	182
3.3.2.2	Results .....	183
3.3.2.3	Discussion .....	185
3.3.2.4	Methods .....	186

3.3.3	Manipulation of central serotonin levels with Escitalopram .....	187
3.3.3.1	Introduction.....	187
3.3.3.2	Results .....	188
3.3.3.3	Discussion .....	190
3.3.3.4	Methods .....	193
4	General discussion.....	195
4.1	Summary.....	197
4.2	Interpretations .....	199
4.2.1	The dynamic of effort allocation is supported by an accumulation model .....	199
4.2.2	What could be the complete model.....	202
4.2.3	What is cost-evidence?.....	204
4.3	Limitations and unaddressed (related) topics.....	205
4.4	Perspectives.....	207
5	References.....	209
6	Appendices .....	235

# 1 Forewords

## 1.1 Publication list

**Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. Proc Natl Acad Sci USA 110:2641–2646.**

The results and methods of this article nearly appear as copy-paste in the first study of the neuroimaging section. The discussion was complemented in the present dissertation.

**Meyniel F, Pessiglione M, Better get back to work: a role for motor beta de-synchronization in incentive motivation (in revision)**

The results and methods of this article appear nearly as copy-paste in the second study of the neuroimaging section. The introduction and the discussion were complemented in the present dissertation.

**Meyniel F, Safra L, Pessiglione M, How the brain decides when to work and when to rest: evidence for implicit cost-evidence monitoring (submission)**

This article corresponds to the behavioral part of the result section; it was at the stage of a working paper when I wrote the dissertation. The submitted version of this article and the corresponding chapter in this dissertation differ substantially not only because of this timing issue, but also because I prefer to include these behavioral data at the beginning of the dissertation although it is my latest article, which required a different way of presenting the data.

These three articles are appended at the end of the dissertation.

## 1.2 How to read this manuscript

The articles prepared during my PhD were reformatted to allow a coherent and smooth reading. Additional non-published results were also included. Some of these results were not published because as such, there are not interesting enough to deserve a publication but they are relevant in this dissertation. Other results are actually negative results: no effect was found. But since the studies were prepared with care and are relevant to the topic, they are nonetheless included. Finally, the partial results of an on-going study are provided, they may deserve publication but I have to wait for the study to be complete.

### 1.3 Acknowledgements

I have many people to thank.

First, my PhD advisor, Mathias Pessiglione for initiating this exciting topic with me. For giving me the opportunity to carry my own projects after your initial impulse and for your close follow up of my work. Your inputs were always much appreciated. I hope I have inherited your methodological rigor, your skill to operationalize concepts and to explain things in simple and attractive terms.

I thank the members of my jury. It is a distinguished honor to defend in front of such a committee: Habib Benali, Stanislas Dehaene, Predrag Petrovic, and particularly Christopher Summerfield and Philippe Tobler who accepted to review this dissertation.

I am also thankful to Mathias for introducing me to the Motivation Brain and Behavior (MBB) team and the co-PIs: Sebastien Bouret and Jean Daunizeau. It was a great pleasure to share your unconditional enthusiasm Sebastien, your valuable knowledge of the brain physiology and your expertise with primates. Jean, I am now a Bayesian devout, I suspect it is your legacy. Thanks for that, the nice talks and the humor.

My MBB fellows, for your intellectual sharpness and your fun. With the privilege of seniority: Stefano Palminteri for your Mediterranean wit, Raphael Lebouc for your precious and repeated help, Mael Lebreton for your surprising humor and insightful talks, Thomas Mauras for adding the subtitles when it comes to clinic. I am honored and happy that the fruitful scientific collaboration turned to friendship. The new MBB wave: Lionel Rigoux for inspirational discussion on motor control and optimization, Aurore San-Galli for your refreshing directness and sharing your experience of monkeys, Bastien Blain for what is unexpected, Vassilissa Skvortsova, for strength, Marie Devaine for sharing your smartness... Actually that's difficult to find the right word for everybody: thanks to Alizée, Chiara, Fabien, Irma, Raphaëlle and the ones I forgot.

The CENIR team was always helpful, in particular Antoine Ducorps, for his amazing kindness and willingness to share his precious experience on MEG, Denis Schwartz for time frequency tools, Christophe Gitton and Laurent Hugueville for the invaluable assistance in building and repairing the electronic grip devices.

The analysis of MEG data benefited from talks with Catherine Tallon Baudry and Claire Sergent.

I thank Véronique Marchand-Pauvert for her expertise on EMG recording. I am still sorry that I burnt out the amplifier you lent me.

It was also nice to have the interesting experience with the pharmaceutical industry. Thanks to this incredible team, the British research assistants, and Raphael Galliard for giving me this opportunity, and the stimulating talks we had.

Thanks to Yann Cojan and Patrik Vuilleumier for this experience with hypnosis. It was a new world to me.

Thanks to Lou Safra, for following me in your master internship despite I was new in supervising. I am lucky that you made my start easy by being so brilliant.

Thanks to Imen El Karoui, for your significant contribution to this work and much more.

Thanks to the Ecole Normale Supérieure and the Ministère de la Recherche for my scholarship and the financial support. Also thanks to the ICM, this new exciting institute.

## 1.4 Résumé français

Faire le bon choix, c'est trouver le bon compromis entre coût et bénéfice. Cette idée a été formalisée par la théorie économique et les neurosciences ont plus récemment étudié les représentations et mécanismes cérébraux associés à ces décisions reposant sur une analyse coût-bénéfice. Ces études ont notamment permis de mieux comprendre comment les coûts sont évalués et comment le processus motivationnel permet de surmonter ces coûts pour obtenir des bénéfices. Ces deux processus sont surtout étudiés dans le cas de choix simples, à un instant donné, entre plusieurs alternatives différant par leur coût et bénéfice. La dimension temporelle que peuvent prendre ces processus est assez méconnue, nous suggérons qu'elle pourrait être clé pour comprendre comment l'effort physique est alloué dans le temps. L'allocation temporelle de l'effort, c'est-à-dire la décision d'arrêter ou reprendre l'effort, est à notre connaissance un processus encore très inexploré, bien que très courant dans la vie quotidienne. Dans ce travail de thèse, nous avons développé un nouveau paradigme expérimental pour recréer en laboratoire le problème de l'allocation temporelle de l'effort chez le sujet humain sain. Nous proposons également un modèle computationnel pour rendre compte de la gestion de l'effort. Ce modèle repose sur le suivi en ligne du coût estimé de l'effort. Ce coût estimé instantané augmente pendant l'effort à mesure que la fatigue gagne la commande motrice et il diminue au repos à mesure que nous récupérons. Dans ce modèle, la décision d'arrêter l'effort est prise quand le coût estimé atteint une limite et, de façon similaire, la décision de reprendre l'effort est prise quand le coût estimé de l'effort est suffisamment bas.

Le comportement des participants reflète les variations de ce coût estimé et du compromis entre cette variable et le bénéfice attendu lorsque la difficulté et l'enjeu monétaire sont manipulés expérimentalement. Plusieurs scénarios basés sur le modèle d'accumulation et de dissipation du coût estimé ont été comparés avec des techniques Bayésiennes pour rendre compte de l'impact des facteurs expérimentaux comme la difficulté et l'enjeu monétaires sur les paramètres du modèle. Cette comparaison a révélé que la difficulté réelle de l'effort augmente la vitesse d'accumulation du coût perçu pendant l'exercice, que l'enjeu permet de repousser les limites de cette variable et que la vitesse de récupération pendant le repos est accrue par l'utilité estimée, qui peut être dissociée de l'utilité réelle grâce à une manipulation indépendante des difficultés réelle et attendue à l'insu du participant. Notre capacité à accéder par l'introspection à cette variable de coût estimé semble limitée, en particulier le ressenti subjectif d'épuisement est distinct de cette variable. L'ensemble de ces études comportementales montrent qu'une part du processus d'allocation de l'effort est implicite et repose sur une adaptation motrice qui limite les coûts physiologiques et qu'une autre part du processus est explicite et permet d'intégrer des facteurs stratégiques dans la régulation du comportement, comme l'enjeu et la difficulté attendue.

Grâce à la complémentarité de l'imagerie fonctionnelle par résonance magnétique et de la magnétoencéphalographie (MEG), nous avons montré sur deux groupes de sujets distincts que la variable computationnelle de notre modèle, le coût estimé, pourrait être encodée dans les régions proprioceptives du cerveau : l'insula postérieure et le thalamus ventromédian. Ce corrélât neuronal ne correspond pas seulement à l'augmentation et la diminution du coût estimé pendant l'effort et le repos, mais également aux modulations de cette dynamique par les facteurs expérimentaux. La MEG a également révélé que la désynchronisation du rythme beta moteur (13-30Hz) permettait de transformer les différents niveaux d'enjeux en un code moteur, avec une reprise plus rapide de l'effort pour les enjeux importants. Enfin, en accord avec le rôle du système autonome sympathique dans l'activation motivationnelle, le rythme cardiaque des participants corrélait avec le niveau d'enjeux, en sus du niveau de performance physique.

Nous avons cherché à démontrer le rapport entre la variable de coût estimé et le signal nociceptif. Cependant, la manipulation antalgique dans deux études séparées avec une suggestion hypnotique ou une prise unique de paracétamol n'a pas affecté la gestion de l'effort. En revanche, l'étude pharmacologique a révélé une modulation de la gestion de l'effort par les niveaux de sérotonine centraux. Le niveau de sérotonine a été manipulé chez le sujet sain avec un traitement prolongé par un inhibiteur sélectif de recapture de la sérotonine : l'Escitalopram. Le traitement a montré une meilleure énergisation de l'effort sous Escitalopram que sous placebo, bien que ces effets attendent encore confirmation sur la cohorte complète de sujet.

Nos résultats montrent que la gestion de l'effort est adaptée en ligne au coût estimé de façon implicite et que ce processus n'est pas rigide, mais modulable explicitement et stratégiquement en fonction des coûts et des bénéfices attendus. Notre contribution se situe à l'interface entre la médecine du sport, la théorie de la décision et les modèles d'accumulation utilisés en neurosciences. Le mécanisme que nous proposons permet d'optimiser la gestion de l'effort physique en maximisant les gains tout en minimisant les dommages corporels.

Mots clés : effort, accumulation, imagerie fonctionnelle par résonance magnétique, magnétoencéphalographie, prise de décision, motivation



## 2 Introduction



## 2.1 Effort allocation over time

### 2.1.1 A very simple question

*What it was that was actually the matter with us, we none of us could be sure of; but the unanimous opinion was that it – whatever it was – had been brought on by overwork.*

*‘What we want is rest’ said Harris.*

*‘Rest and a complete change’ said George. ‘The overstrain upon our brains has produced a general depression throughout the system. Change of scene and absence of the necessity for thought, will restore the mental equilibrium’.*

*George has a cousin who is usually described in the charge-sheet as a medical student, so that he naturally has a somewhat family-physicianary way of putting things.*

Jerome K. Jerome, *Three men in a boat*, 1889.

The reader of Jerome’s novel has serious doubts that these nineteenth-century idly gentlemen truly suffer from exhaustion; however, no matter how intense it really is, exhaustion is a pervasive feeling we all have experienced. According to the quoted novel, work and rest are seemingly opposite processes acting upon a common resource (whatever it may be): fatigue accumulates during effort such that we need a break at some point to restore it to lower levels. George’s diagnosis suggests that the brain is altered somehow by work. Despite his expertise, he does not make it clear however why the brain could be involved in this process, whether it is in general because the brain regulates the behavior or because it is the very organ that was depressed by strenuous work: thinking. To make things clearer to that respect, the investigation of physical effort could disentangle these two alternative explanations. How do we allocate physical effort over time? When, why and how do we decide to take breaks or to resume our effort?

This is in a nutshell the issue addressed in this PhD dissertation; it will be progressively refined and documented in this introduction. The fridge-moving problem will serve, hopefully, as an inspirational example to guide the reading of this introduction. Imagine you are moving to a new place with all your furniture. Unfortunately, your new place has the charm of ancient buildings, but no elevator, so that you have to move everything up through the stairs. Your fridge, if nothing else, is your heaviest piece of furniture: it is such a burden that you strive to move it and, in the course of your effort, you take several breaks. Figure 1 may represent your course of action. When, why and how do you decide to take breaks or to resume your effort?

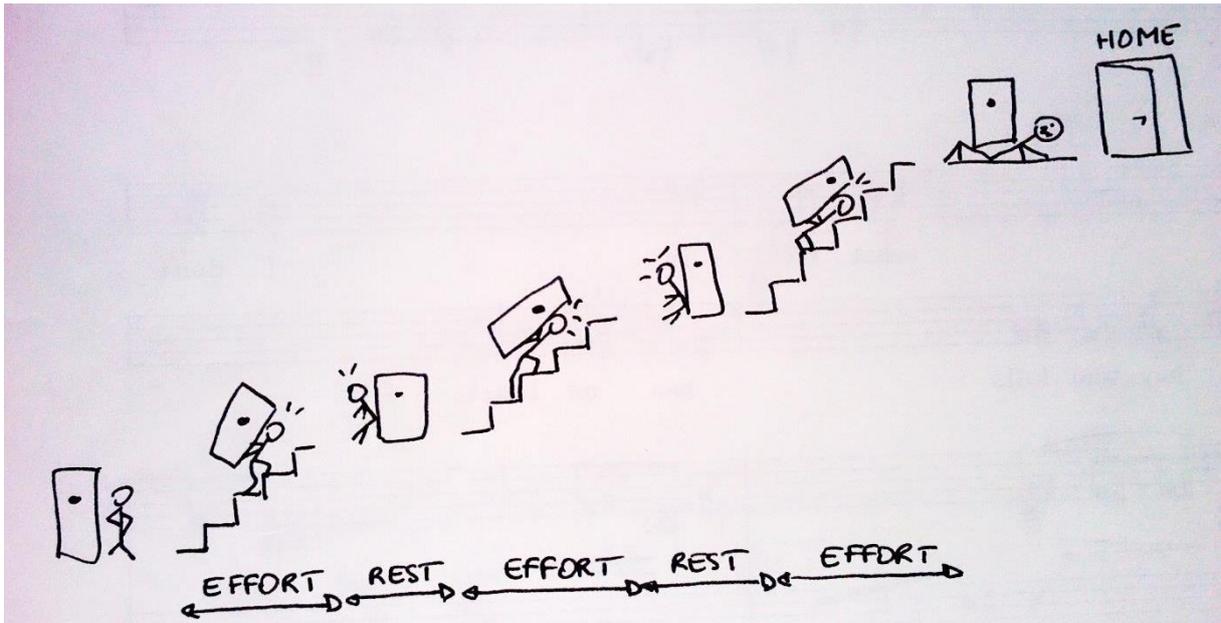


Figure 1 The problem of moving a fridge up through the stairs

The present introduction aims at relating the effort allocation problem to the current scientific context. This review of the literature is an opportunity to bridge several distinct fields of research in the context of this specific problem. The last part of the introduction presents a principled way to answer the question inspired by the state-of-the-art knowledge and then details the outline of my experimental work.

### 2.1.2 Basic intuitions

A quote by the psychologist William James may be useful to share basic intuitions about effort allocation and to serve as a transition to the different literatures reviewed in this introduction. Here are W. James' words (James, 1907):

'Ordinarily, we stop when we meet the first effective layer, so to call it, of fatigue. (...) But if an unusual necessity force us to press onward, a surprising thing can happen. The fatigue gets worse up to a critical point, when gradually or suddenly it passes away (...). We have evidently tapped a level of new energy.'

For the record, the use of this quote here might be a bit far-stretched since W. James intended to dissert on the concept of 'mental energy' from a clinical perspective and at the time scale of days at most, if not years. He pointed out that this concept of mental energy, despite corresponding to a widely shared experience and being potentially a powerful explanatory concept, was both ignored by the scientific community and poorly defined. This concept shares both its fuzziness and its temporal profile with the notion of resource depletion in physical effort allocation: something that increases

during effort and is restored during rest. This resource will be progressively defined in my dissertation, and I will try to uncover the role of the brain in the physical effort allocation process. For the time being, let us just benefit from W. James' intuitions.

Three interesting leads can be found in this quote. First, the notions of 'critical point' and 'level' suggest that some continuous variable accumulates during the effort and that there is a bound to this process to trigger the decision to stop. The decision to stop would therefore be underpinned by a comparison between the level of a variable and a limit. Second, James suggested that this limit can be shifted according to the goal we are pursuing ('unusual necessity'), as if the level of the resource was traded against the value of our goals. How we value the goal of our effort therefore seems to play a key role in the effort allocation problem. Finally, James speaks of 'fatigue' and 'energy': what is the nature of the variable accumulated?

William James clearly proposed a model to account for how we decide when to stop the effort: the effort duration results from a tradeoff between a variable related to effort cost and the value of the goal we pursue through the effort. In simple terms, this decision would be a tradeoff between cost and benefit. This notion of value-based decision-making has been extensively investigated in economics and in neurosciences; it will be reviewed in the first section (2.2). Given that the cost we monitor during effort allocation sounds much physical: 'energy', 'fatigue', the second section will address physiology of the effort production (2.3). Next, the effort allocation problem being essentially dynamical: we have to set the timing of transitions between effort and rest, how does time step into this problem? According to James, the reason is that cost is dynamical. But it could also be that benefits are affected by time. How time affects cost (through fatigue) and benefit (through delay discounting) is addressed in section 2.4. These three first sections are intended to provide ideas on how decisions are made based on value and what the constraints affecting these decisions are in the physical effort allocation problem. Then, I review in section 2.5 some existing explanations and models that account for the timing of effort and rest transitions to show that they provide useful ideas but that none of these models are satisfactory to account for the physical effort allocation problem. The need for a better explanation motivated my experimental and theoretical work, which is introduced in section 2.6.



## 2.2 The economic view: decision aims at maximizing benefit against cost

Following basic intuitions it is likely that effort allocation is impacted by the prospect of a benefit: how motivated we are to move the fridge up through the stairs; and the cost associated to it: how difficult it is to move this fridge. Effort allocation may depend on the comparison between costs and benefits. This tradeoff between cost and benefit is widespread in decision making and it has been extensively investigated. In the first section I discuss value-based decision-making, starting from the view point of the economic literature to formalize useful ideas (section 2.2.1.1), then I have a look into how such decisions may be encoded in the brain (sections 2.2.1.2, 2.2.1.3 and 2.2.1.4). The term effort-based decision making was coined especially when the value involves effort, this is presented in section 2.2.2.

### 2.2.1 Value-based decision making

The concept of value-based decision making is currently used in several research fields (Rangel et al., 2008). For instance in economics; this field contributed to give the concept of value-based decision a strong formalism in order to build a normative view of how decisions should be made. Behavioral psychology also addressed value-based decision-making and operationalized this notion in animal studies, in particular to understand how animals learn to choose the best option among several alternatives. Computing science also tried to formalize these processes with a computational description. These research fields met neuroscience fruitfully in the past two decades in order to bind the different description levels together: behavior, computation and neural implementation. More recently a program emerged from this endeavor to bridge the gap between experimental economics and neurosciences, which was termed 'neuroeconomics' (Glimcher and Rustichini, 2004; Fehr and Rangel, 2011).

In the following, I first introduce the concept of utility (section 2.2.1.1) which is interesting because it was forged to reflect the tradeoff between cost and benefit. In the effort allocation problem, it could be useful to estimate what is the dynamic of utility over time. Second, I discuss how the brain computes utility (section 2.2.1.2). Third, if utility accounts for choice, we need to specify the mechanisms by which valuation is translated into decisions (section 2.2.1.3). In the last section (2.2.1.4) I discuss whether different decision systems should be distinguished.

#### 2.2.1.1 *Utility: cost traded against benefit*

##### 2.2.1.1.1 What utility is for economists

The purpose of the economic theory, right from the beginning, is to explain decisions (Glimcher et al., 2009). In economics, a decision is a choice within a set of alternatives. Two lines of research were

proposed. Both aimed at explaining choices, one by assuming that decisions are underpinned by preferences that satisfy some axioms, the other approach by focusing on the choice level it-self, imposing constraints of consistency that are similar to the first approach. Both formalisms, under certain assumptions, are equivalent. In these formalisms, the preferences are fundamental characteristics of individuals and have relational property for alternatives of a set. This relation is similar to the relational property 'is bigger or equal than' for numbers. It can be applied to any pair of numbers, and it is transitive (if  $x$  is at least bigger than  $y$  and  $y$  at least bigger than  $z$  then  $x$  is at least bigger than  $z$ ). The reason for these constraints is that it helps to keep the problem simple: whatever the alternatives may be, things can be compared; and these preferences are consistent over pairwise comparisons. Choices (something we have a direct access to) reveal preferences (something we do not have access to) in the sense that choices should be consistent with the relational preference properties between the alternatives of a set. Given that the preference relation is similar to the relation 'is bigger or equal than' for numbers, it is convenient to use numbers instead of preference. This version of preference is termed utility; it is quantitative and thus can be conveniently handled. Utility is a function that associates each alternative of a set to a numerical value. Note that there is not a one to one correspondence between a set of preferences and a utility function. For instance two utility functions differing by a constant term theoretically can do the same job. To rephrase in term of utility what was said above for choice and preference, choices are made within the alternatives of a set so as to maximize the utility (Mas-Colell et al., 1995). The utility function is thus a computational hidden variable (we do not have direct access to it) that accounts in a principled way for the choices.

#### 2.2.1.1.2 What are the modulators of utility?

The utility associated to an alternative depends on the outcome entailed by choosing this alternative. Three characteristics of the outcome are fundamental determinants of utility: the uncertainty associated to the outcome delivery, the valence of the outcome and the delay within which the outcome is delivered (Doya, 2008). The matter of delay is crucial in the effort allocation problem; it is reviewed in section 2.4.

##### 2.2.1.1.2.1 *The uncertainty*

There are many situations in which the outcome of a choice is not 'for sure', but probabilistic. For instance, when planning to see a movie, you might have to choose between two theaters: one you like much, with this incredibly large screen and great sound system and the other one that you like less. It seems you would choose the first one. However, you do not have your entrance ticket and you know the first theater suffers from its fame so that it may be full: there is substantial chance that

you cannot see the movie. In the second theater, you know for sure that you can see the movie... So, which option to choose now? John von Neumann and Oskar Mongenster incorporated this probabilistic component into the notion of utility during the 1940s and formalized the so-called expected-utility theory (Glimcher et al., 2009). The idea is that the intrinsic value of seeing the movie should be weighted by the odds for this to happen. Formally, this corresponds to the mathematical expectancy, hence the name of the theory. A current challenge for the research is to understand how the uncertainty of choice is processed in the brain, and to incorporate in the theory other tasks potentially relevant for the clinics besides the tasks classically used in behavioral economics (Schonberg et al., 2011).

#### *2.2.1.1.2.2 The valence of the outcome*

The construct of utility puts on equal footing outcome values with either positive or negative valence, i.e. whether it is something you prefer to get or to avoid through your decision (Rangel et al., 2008). The valence of the outcome is used to distinguish, semantically, positive value as benefit, and negative value as cost. The net value is thus a compound made of cost and benefit (Rangel and Hare, 2010). For instance, there is this picnic organized by friends of yours, with this barbecue food you enjoy so much (something you would like to have), but the party is somehow ruined by the winter-like weather (something you really hate). You cannot get the barbecue without being outside in the cold, so that the utility associated to this event is a mixture of positive and negative features, i.e. benefits and costs. How precisely to weight cost against benefit? Two classes of model can be distinguished (Talmi et al., 2009; Talmi and Pine, 2012): the linear discount (basically 'benefit minus cost') or an interactive discount, e.g. hyperbolic (basically, the ratio between benefit and cost, with cost scaled such that the lowest cost value – no cost – is one, and above one when the cost can no longer be neglected). There are very different kinds of cost: monetary losses, disgusting food, pain, effort to make, etc. (Rangel et al., 2008).

#### *2.2.1.1.2.3 A subjective estimate*

It became rapidly clear while using the concept of utility that subjective estimates differed from the objective data. The reason is that many determinants of utility are subjective.

Probabilities are subjective: people tend to systematically overestimate small probabilities and to underestimate large probabilities (Kahneman and Tversky, 1979). These deviations are rather systematic across individuals so that it was possible to refine the expected utility theory by distorting objective probabilities (see Figure 2). However, these deviations also show significant variations between individuals.

The outcome valence is subjective: it is obvious for food for instance. Between-subject differences may be an issue in the context of experiments: an experimenter may want to manipulate costs by varying the task difficulty; however higher difficulty levels are not necessarily treated as a cost but sometimes as something we seek, e.g. for intellectual and sport exercises (Brehm and Self, 1989; Camerer and Hogarth, 1999). Even extreme cases are not consensual: pain, widely perceived as a cost, may be a benefit for masochist people (Harenski et al., 2012).

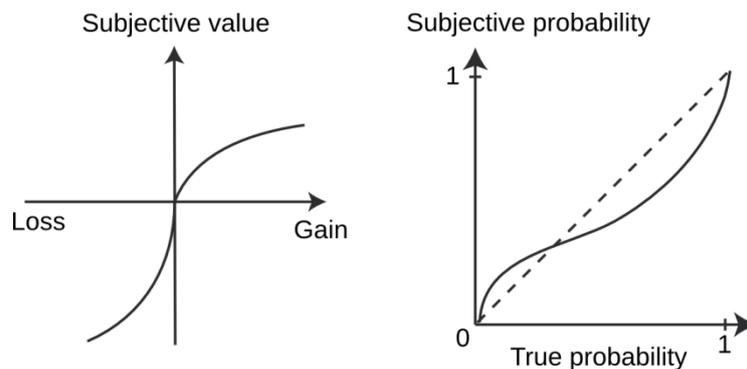


Figure 2 Subjective value and subjective probability

The magnitude of the outcome value is also subjective. In particular, the subjective increase of positive outcome value is smaller than the objective increase; it is the opposite for negative outcome values: subjective decreases are greater than objective decreases. In addition, these deviations from the diagonal are not uniform but have a curvature that differs between the gain and loss domains, see Figure 2. These deviations make us risk averse for negative outcomes and risk seeking for positive outcomes. The difference in the degree of concavity and convexity of the objective-subjective outcome value curve accounts for different risk attitudes across people (Schonberg et al., 2011).

All these differences are observed between subjects, but also at the level of each individual. For instance the utility of food depends on the physiological state (Doya, 2008). Utility estimates also depend on factors as subtle as how the experiment is framed (Tversky and Kahneman, 1981).

### 2.2.1.2 Brain Valuation System

#### 2.2.1.2.1 From economics to neuroscience: Internal currency

The utility has the advantage that it is a quantitative variable that reflects very different determinants of choice. In particular, the utility theory posits that all alternatives can be compared, because a numerical value reflecting relative preference can be assigned to all goods. This numerical value is therefore like a common currency that enables comparison between goods. To bridge the gap between the theory and the neural implementation, people tried to search in the brain for signals that would be similar to this theoretical variable. This search was a major challenge for

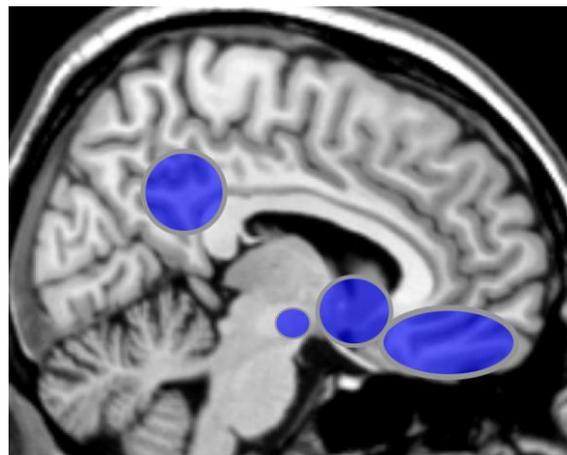
neuroscience and the cumulated evidence tend to support that there are such value signals in the brain, notably in the orbital and ventro-medial prefrontal cortex, based on human and non-human primate data (Montague, 2002; Sugrue et al., 2005; Levy and Glimcher, 2011, 2012).

Although the motivation to search the brain for utility-like signals was to relate the abstract and principled economic view of choices to its implementation in the brain, the metaphor of internal currency should be rounded. On the one hand, brain signals reflecting the utility show some characteristic of a 'currency'. For instance, a crucial property of a currency is that the value is absolute; the value of alternative should therefore be somehow intrinsic and not dependent on the other alternatives in the set. This is a kind of 'menu-invariance' that was evidenced in single unit recordings in the orbito-frontal cortex of monkeys (Padoa-Schioppa and Assad, 2008). On the other hand, the utility theory itself does not impose that the value is absolute, but rather relative to other values. The value of alternatives of a set could therefore be scaled depending on which items are in the current set, while preserving the relative value between all items in the set. The implementation of such a normalization of values within a set would be similar to what happens for sensory systems (the gain is adapted to the range), which has the advantage of optimizing the discrimination between values when signals are noisy. Both theoretical models and experimental data support this notion (Tobler et al., 2005; Seymour and McClure, 2008; Rangel and Clithero, 2012).

#### 2.2.1.2.2 Functional anatomy of valuation

The ventro-medial prefrontal cortex is probably the region that drew most attention with respect to value-based decisions (Rushworth et al., 2011). In fMRI, the activity in this region correlated with the value of items, regardless of how this value was estimated: it could be a rating by the participant (Lebreton et al., 2009) or an estimation of the value derived from choices, and possibly updated through learning (Glascher et al., 2009), or a willingness to pay (Plassmann et al., 2010). The willingness to pay is a method used in behavioral economics by which participants give a monetary value to alternatives they are presented with (Becker et al., 1964). More precisely, the participant offers a price to goods and this price is eventually traded against the good: if the offered price is higher than the true price (or an auction), the participant gets the good, otherwise she just keeps the money that she offered. Note that they are provided with an initial budget that is limited. This rule makes people give the most accurate price based on how worth they think the good is, because an offer that is too low results in a loss: the participant keeps the small amount of money instead on having the higher-value good, and an offer that is too high also results in a loss because the participant spends too much money for a good that does not worth it.

Meta-analyses of published studies reveal that the three studies mentioned above are representative and support that the medial prefrontal cortex is a key brain region involved in valuation (Peters and Büchel, 2010; Kühn and Gallinat, 2012; Bartra et al., 2013; Sescousse et al., 2013). Pattern analysis of the fMRI signal also supported this conclusion (Kahnt et al., 2011). Single unit recordings revealed that neuronal activity in this region correlated with subjective preferences in monkeys (Sugrue et al., 2005; Padoa-Schioppa and Assad, 2006; Bouret and Richmond, 2010). Human surface electrophysiological recordings also showed evidence of value-related processing in the orbital frontal cortex (Doñamayo et al., 2011; Hunt et al., 2012). Lesion studies in this region in humans (Fellows, 2011) and macaque monkeys alter the discrimination between subjective values (Noonan et al., 2010), supporting this region has a causal role in the processing of the utility. However, the medial and lateral orbitofrontal regions are not the only areas involved in valuation. fMRI data also support the role of the ventral striatum and the posterior cingulate cortex (Knutson et al., 2005; Lebreton et al., 2009; Peters and Büchel, 2010). The striatum in particular may have a pivotal role to translate these values into a behavioral response (Schmidt et al., 2008; Liljeholm and O’Doherty, 2012).



*Figure 3 Core regions of the brain valuation system*

*From left to right (i.e. from the posterior to the anterior parts of the sagittal slice), blue blobs are: the posterior cingulate cortex, the ventral tegmental area and the substantia nigra pars compacta, the ventral striatum and the ventromedial prefrontal cortex.*

Two interesting and related questions are the following. First is whether these value signals truly reflect value monotonically, or instead, whether they reflect confounds such as their salience (i.e. the absolute value). A formal difference between value and salience is that extremely negative and positive values elicit the same signal level in the case of salience and distinct signal levels in the case of monotonic value. It seems that several value signals actually correlate with salience (Bartra et al., 2013). Second is whether all kinds of alternatives are processed by the same system. The above-

mentioned areas can be viewed as a core system for valuation, but type-specific areas are also at play and are different, for instance, between food, money or erotic pictures (Sescousse et al., 2013)

#### 2.2.1.2.3 Different valuation systems

Whether there is only one valuation system in the brain, potentially distributed within a network of areas or distinct valuation processes, underpinned by different brain areas is a key issue. At a very fine scale, for instance in the medial prefrontal cortex (which is the best candidate for valuation), single unit recordings reveal that the encoding of value signal is not homogeneous, but involves separate neurons for cost and benefit (Amemori and Graybiel, 2012). Similarly, in a cost-benefit trade-off task, recordings in the substantia nigra pars compacta, a dopaminergic nucleus, revealed that many neurons reflected benefit but few reflected effort, suggesting that cost was not processed in this nucleus (Pasquereau and Turner, 2013). Similar evidence was obtained from voltammetry (Gan et al., 2010). At a coarser level, in the prefrontal cortex, whether the medial and lateral parts compute the same signal is not clear either. The medial part reliably encodes value of actions, outcomes and goals whereas correlates of value in the lateral part are less frequently reported (Peters and Büchel, 2010). It was proposed that the medial to lateral gradient could reflect abstract value processing to more sensory-based value processing, in particular for cost (Kringelbach, 2005). The anatomical pattern of connectivity could support this distinction since the lateral part of the orbitofrontal cortex receives more sensory-related inputs than the medial part (Peters and Büchel, 2010). These different parts of the medial prefrontal cortex, because of their patterns of connectivity, could underpin different types of associations (Euston et al., 2012) and thus potentially different valuations. This issue is still debated, with studies reporting cost and benefit valuation in the same areas (Basten et al., 2010; Plassmann et al., 2010) and other arguably reporting that the processing in the medial frontal cortex reflects only the benefit and that other areas represent cost and benefit at the same time (Croxson et al., 2009; Prévost et al., 2010). Part of the answer might be that there are different kinds of cost, such as the cost to get the outcome (like effort) that would be processed in the anterior cingulate cortex and cost associated to the outcome delivery (like a price paid for some good) would be processed in the orbitofrontal cortex (Rangel and Hare, 2010).

To what extent the valuation process is unique still needs investigations. Some results from genetics support that not all decisions have the same neurobiology. Between-subject differences in decisions involving positive or negative outcomes were related to different polymorphisms (Ramsøy and Skov, 2010): the val158met polymorphism in the catecholamine transporter COMT is associated to variations in decisions involving positive outcomes whereas the 5-HTTLPR polymorphism of the serotonin transporter is associated to variations in decisions involving negative outcomes. This

genetic evidence for distinct valuation systems could however be rounded because the dissociation between behavioral types may not be so sharp and because there is an ambiguity whether the differences are in valuation processes or decision processes. It is not clear either whether there are different value estimations for actions that lead to an outcome, the goal it-self and the stimuli that signal in advance the quality of the outcome (Rushworth et al., 2011). The team led by Antonio Rangel introduced a specific terminology for these values: action-value, goal-value, cue-value etc. (Rangel et al., 2008; Rangel and Hare, 2010). These distinctions may suggest that there are different kinds of values that depend on the structure of the task and the function of these values in the task. An additional category could be the state value, related to the current experience, for which valuation would be automatic, regardless of the participant's actions, and notably even when she is not engaged in a value-based task (Lebreton et al., 2009).

To conclude, it should be stressed that not all signals correlating to the utility are involved in processing this quantity. There are several reports of top-down processing (or at least interpreted as such) in low level sensory areas that reflect the utility, for instance in the visual cortex as demonstrated in humans with both functional MRI and magneto-encephalography (Serences, 2008; Tallon-Baudry et al., 2011). It is also the case for regions like the motor cortex that implements decisions based on valuation and that are thereby impacted by the value despite they are not involved in processing this value.

### *2.2.1.3 From valuation to decision and learning*

#### *2.2.1.3.1 Valuation guides decision*

The notion of utility was introduced above to account for how choices are made. This may be useful because the dynamic of utility could be related to the decision to stop and resume the effort in the effort allocation problem. In other words, the utility has the potential to account for which decisions are made, but it leaves open the issue how they are made. Potential mechanisms to translate utility into a decision are now presented.

##### *2.2.1.3.1.1 Computational views on decision*

The utility theory is framed in such a way that we choose the alternative we prefer. Put this way, this is rather intuitive. What is also intuitive is that the value of alternatives is something we learn. In fact, valuation, decision and learning are tightly related for this reason, and thereby often discussed together in the literature (Rangel et al., 2008). Given that value is not predetermined but learned, it implies that we have to sample and experience values from our environment. Always choosing the best alternative (the one that has been associated so far with our best experience) does not leave

room for this sampling. How to be sure that the option we keep on choosing is still the best one if we do not try something else? This trade-off between how often we should choose our preferred option or try another option is known as the exploitation – exploration tradeoff. Complete exploitation corresponds to the greedy choice of staying strictly on the option we think is the best. This may not be the best strategy however, for instance, when the environment is unstable and the true value (not the estimated value) of alternatives changes over time. Indeed, if the option one keeps on choosing is no longer the best option, one will simply miss better options by sticking on the initial choice. Complete exploration corresponds to choosing randomly, regardless of what we prefer. This strategy ensures that our sampling of the valuation of alternatives is unbiased, but the advantage is not taken upon this good estimation by shifting the choice toward better options. The optimal behavior is a mixed strategy with both exploration and exploitation (Dayan and Abbott, 2005).

A solution to round the maximum rule, i.e. not to always choose the best option, is to follow the so-called soft-max rule. For illustration purpose, if the choice is to accept or reject an offer, the soft-max rule says that the probability to reject, denoted  $p$ , should be set according to the difference of utility between rejecting and accepting the option, denoted  $\delta$ , so that:  $p = \frac{1}{1+e^{-\beta\delta}}$ . The  $\beta$  parameter controls the exploration-exploitation tradeoff of the choice: if  $\beta=0$ , the choice is a pure exploration, completely random, and higher values of  $\beta$  make the choice closer to strict exploitation of the best option. This example is a simplified version of the soft-max rule that can be extended to choice between more than two alternatives.

#### *2.2.1.3.1.2 Neural implementation*

The implementation of value-based choice in the brain leaves open several related issues: Are valuation and selection serial or parallel processes? How does the brain track the opportunity to exploit and explore? How does the brain select an option?

##### *2.2.1.3.1.2.1 Foraging vs. exploiting*

Prior to deciding which of the alternative one prefers within a set, how does the brain decide to restrict choice to this specific set? This is a form of exploration – exploitation tradeoff: should I continue with the current set of alternatives (e.g. looking for job opportunities in my city) or change of set (decide to move to another city for a job)? This comparison also has to take into account the cost to change from the current set. It was proposed by the team of M. Rushworth that tracking the value within the current set is encoded in the medial prefrontal cortex and that the cost and benefit of switching to another set are encoded in the anterior cingulate cortex (Kolling et al., 2012; Rushworth et al., 2012). In this model, several values are tracked in parallel, and the model has the potential to reconcile different views in the debate presented above on why the medial prefrontal

cortex does not show reliable correlation with cost while correlations are robust for benefit (Rushworth et al., 2011).

### 2.2.1.3.1.2.2 Neural models to link estimation and choice

On a different topic, how are valuation and choices processed? Is it serially: first, values are computed, then, they are compared to guide choice; or in parallel: the choice (hence the comparison of values) is intrinsically related to the estimation of values? The current trend is to accept the parallel processing, this view is supported by detailed understanding of the comparison mechanism (Rangel and Hare, 2010; Rushworth et al., 2012). At least two classes of model are discussed to account for such parallel decision process; they may reflect a dual description of the same process, though their neural description differs: the ramping-to-threshold in the temporal domain (drift-diffusion models) and the trajectory of population activity in the state space (recurrent network models) (Wang, 2012). To explain in a nutshell these models, say one has to choose between option 1 and option 2. The drift diffusion model integrates the (noisy) momentary evidence of the value difference between the options. The recurrent network model consists of two populations of neurons, mutually inhibited and excited each in proportion of the value of an option, so that there is a mapping between options and populations. If one value is higher than the other, the competition is biased and the population activity converges to high activity in the population receiving input from the best option (see Figure 4).

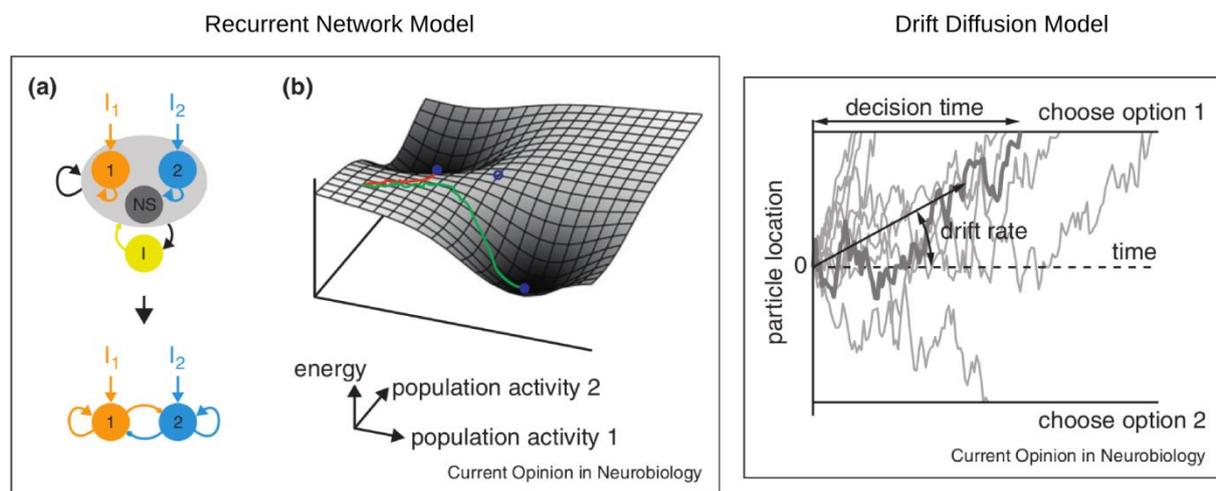


Figure 4 Recurrent network model and Drift diffusion model of choice between two option.

See text for explanation. Picture from (Drugowitsch and Pouget, 2012)

Several important differences should be stressed between these two models:

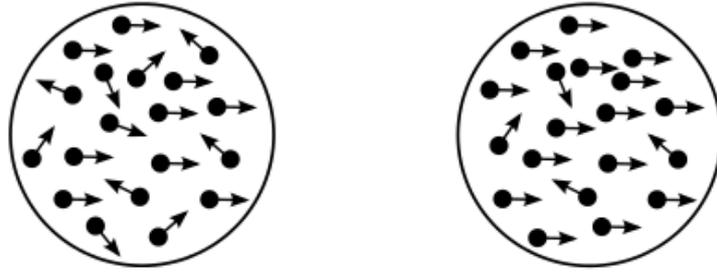
1. Drift-diffusion models can integrate the comparison over an infinite time whereas recurrent network models converge, so that the integration time is finite.

2. The momentary evidence is weighted equally by drift-diffusion models whereas early evidence is weighted more by recurrent-network models.
3. Drift-diffusion models can be encoded by single neurons whereas recurrent network models characterize the population dynamics.
4. Drift-diffusion models lead to 'ramping' dynamics whereas recurrent network models can display both ramping and jumping (abrupt transition) dynamics.

Note that the time constant of the accumulation process in drift diffusion models is several orders of magnitude higher than the neural time constant: the neural implementation of such model is not an easy task a priori, and several models strive to explain why the accumulation is not washed out by fast neural dynamics. Ratchet models have been proposed, in which incremental variations are made possible by local stabilities at each increment (Cain and Shea-Brown, 2012). Regarding the biological plausibility, the potential of recurrent network models may have been overlooked. It is a matter of fact that in neuroscience, drift-diffusion models are by far more widely used; a reason may be that it is easier to predict the corresponding dynamics in electrophysiological signals and functional MRI. Another good reason is that drift-diffusion models are more optimal than the recurrent network models in terms of choice accuracy (Drugowitsch and Pouget, 2012). Recently, another model that is both more optimal than drift diffusion models and more rooted in the neurobiology has been proposed with probabilistic population codes (Ma et al., 2006; Drugowitsch and Pouget, 2012). However, drift-diffusion models are the most commonly used in neuroscience, and are therefore presented in detail below.

#### *2.2.1.3.1.2.3 Drift-diffusion models in perceptual and value-based decisions*

Drift-diffusion models were intensively used to account for perceptual decision making (Gold and Shadlen, 2001, 2007; Heekeren et al., 2008). These models account for the speed-accuracy tradeoff observed behaviorally in perceptual decisions and relate this tradeoff mechanistically to the activity of neurons. This was demonstrated in particular in the random dot motion task, see Figure 5 (Shadlen et al., 2006). In this task, the participant is presented with moving dots; their motion can vary from complete randomness to strong coherence for a certain direction, e.g. left or right. The subjects are asked to estimate this direction.



*Figure 5 Random Dot Motion Task.*

*The actual visual display is made of moving dots, the little arrows indicate the direction of each dots. The right compared to left panel reveals different coherence in the motion: dots move more coherently (to the right) on the right panel than on the left panel. From (Drugowitsch et al., 2012).*

The drift-diffusion model works like an integrator of a noisy input that accumulates the evidence in favor of a certain direction. Another equivalent view of drift-diffusion is that the integration over the noisy input is like a random walk process: the random walk is biased toward more extreme values and this bias in the drift is all the stronger that the motion coherence is high, i.e. that it is clear that the motion is biased toward a certain direction. Placing a bound to this accumulation process enables to translate the estimation of evidence into a decision. Varying the bound level corresponds to varying the conservativeness in the response: low thresholds are reached quickly but are prone to errors since the accumulation may be dominated by noise; high thresholds are reached after longer delays but make the result less likely to be due to noise, i.e. more reliable as the noise is progressively averaged out. This tradeoff accounts for the speed-accuracy tradeoff observed behaviorally.

The firing rate of neurons in the lateral intraparietal cortex reflects this biased random walk: they encode the sensory evidence that the motion is in the direction encoded by their receptive field. Therefore, the firing rate of these neurons reflects the integration of the incoming evidence in their receptive field. Interestingly, neurons in the lateral intraparietal cortex can also combine more complex kinds of evidence, implementing probabilistic reasoning when different pieces of evidence need to be combined and weighted (Yang and Shadlen, 2007) and can also reflect the confidence in the decision (Kiani and Shadlen, 2009). From a formal point of view it is not clear in the random dot motion task whether the evidence accumulated is purely decisional, i.e. the probability that the response is correct, or related to reward (the expected reward magnitude). Attempts to disentangle the two, by manipulating the subjective value with variable delays to get the reward, suggests that the accumulated variable is close to a reward signal (Louie and Glimcher, 2010). No matter whether this variable is purely a choice probability or a value signal, the evidence in the lateral intraparietal cortex is primarily driven by the perceptual evidence. This accumulation process is also tightly

coupled to the motor preparation of the response (Gold and Shadlen, 2000); however although the gradual motor preparation is fed by the evidence accumulation, the two processes are still dissociable (O'Connell et al., 2012). Human surface electrophysiology also suggested that perceptual evidence is accumulated in the parietal cortex (Wyart et al., 2012) and may be cross-modal, in particular independent from whether the stimulus is visual or auditory (O'Connell et al., 2012).

The same drift-diffusion models were also proposed to account for the estimation of values in the economic sense. It was proposed that the value assigned to items is integrated, and that the accumulation could depend on attention: the more we pay attention to items, the more we accumulate evidence in favor of choosing them (Krajbich et al., 2010). In more complex cases, the accumulation would not reflect the evidence in favor of choosing an alternative against another, it could also integrate the evidence in the item utility by combining cost and benefit (Basten et al., 2010). As for perceptual decision and confidence, it was proposed that the evidence accumulation also determine the confidence in the response (De Martino et al., 2013). Note however a striking difference between perceptual decision making and value-based decision making: the momentary evidence is conditioned by the environment in the former case and by internal constraints in the latter. What does it mean to integrate the value of eating an apple or an orange? Why an integration method should be used to read out a variable (the value) that is internal? In other words, why should we sample ourselves to know whether we prefer apple to orange? It might be that the valuation system truly integrates evidence from other brain areas to read out information, for instance from the memory (what was it like the last time I tried apples and oranges?).

As presented in the case of perceptual and value-based decision, the drift-diffusion model directly relates the estimation of values (accumulation of evidence) to a decision (a choice triggered by reaching a threshold). From a theoretical point of view, this choice heuristic is an approximation of the optimal algorithm that should determine choice: the sequential probability ratio test (Cain and Shea-Brown, 2012; Churchland and Ditterich, 2012). This optimal test was developed during the Second World War to automatize the decoding of the German secret coding scheme or choosing the best ballistic method for firing: the idea is to update progressively the posterior probability of alternative hypotheses and to commit to a choice when a given threshold is reached (Drugowitsch and Pouget, 2012). In other words, the algorithm collects evidence as long as the likelihood ratio between the alternatives is not sufficient, and it stops this process when a threshold is reached. This threshold controls that (on average) the best decision accuracy is reached within a minimum duration (Churchland and Ditterich, 2012).

#### *2.2.1.3.1.2.4 Choice with more than two alternatives*

Choice in the decision making literature is often binary: choose or reject; or choose between two alternatives. The drift-diffusion model is a rather intuitive way of implementing this kind of choice: the decision variable accumulated is framed such that the drift can go in two opposite directions; for each of these directions a threshold triggers the decision to accept or reject. Krajbich and Rangel extended their drift-diffusion model of binary value-based decision to trinary choices (Krajbich et al., 2010; Krajbich and Rangel, 2011). Instead of integrating the value difference between two alternatives and committing decision when there is sufficient evidence that this value difference is positive or negative, these authors proposed that there are separate integrators representing each option. Each integrator accumulates the evidence that the value of the option it represents is higher than any other value; in other words, it computes the difference between the value of the alternative it represents and the best value among the remaining alternatives. The choice is made in favor of the integrator that first reaches a given threshold.

A similar drift-diffusion model, but more detailed in the description of the neural implementation, was proposed by another team (Churchland and Ditterich, 2012). Each integrator is a set of neurons with auto-excitation, there are inputs that reflect the values of options such that each input excites its respective pool of neurons and inhibits other pools with feedforward inhibition. Recurrent network models can also implement choice among more than two alternatives with competition between neural pools, but with notable differences in the implementation compared to the drift-diffusion model: excitatory inputs reflecting values target only their respective pool of neurons and the competition between pools is implemented by shared inhibition (Wang, 2012).

### 2.2.1.3.2 Valuation can be updated through learning

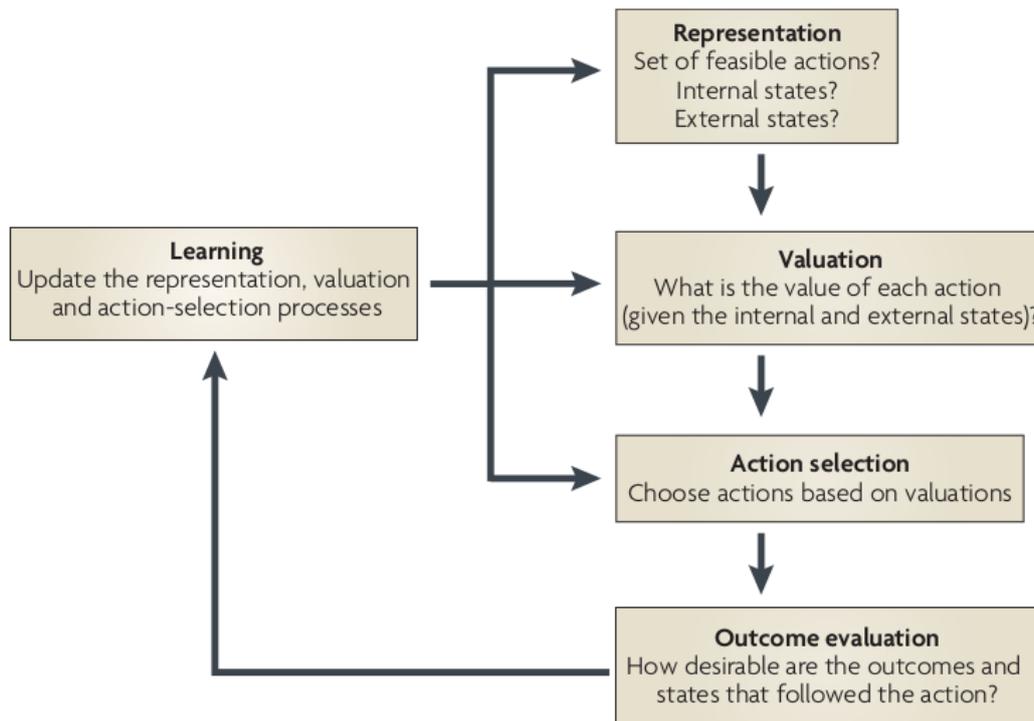


Figure 6 Learning and value-based decisions. From (Rangel et al., 2008)

Learning is essentially related to valuation and decision, see (Rangel et al., 2008). The reason is ontological: many values that we assign to things result from learning; and adaptive: to fit in new environments, choices need to leave room for learning. Two concepts are essential in learning: the expectations and the update of these expectations based on new experience. These concepts are introduced here not only because they are related to decision in general, but also because they may be related to our problem in particular. When moving a fridge, our expectation on how difficult the effort is and the potential update of this expectation could modulate our effort allocation.

#### 2.2.1.3.2.1 Computational views on learning

Learning is a modification as a consequence of an experience. Rescorla and Wagner proposed in the 1960s a simple equation to model how our expectations are updated depending on our experience. The success of this simple equation is considerable, it was and still is often used to model learning.

In simple terms, this model posits that we should revise our expectations relatively to how much our experience differs from our expectations. The prediction error quantifies the discrepancy between experience and expectation. Given a quantitative variable, e.g. the utility,  $U$ , our estimation of  $U$  is updated such that  $U \rightarrow U + \alpha\delta$ , where  $\delta$  is the prediction error, the expected value of  $U$  minus the experienced value. Positive prediction errors are 'good surprises': it is better than expected; negative prediction errors are bad surprises: it is worse than expected. The  $\alpha$  parameter scales how much prediction errors should update our estimation of  $U$ : this parameter is therefore termed the learning

rate. If the value of  $U$  is stationary, the learning rate controls the speed at which the estimate of  $U$  converges to the true value (Dayan and Abbott, 2005). Higher is not always better: intermediate learning rates can provide accurate estimates. When the experienced utility is not constant but noisy, large learning rates result in unstable utility estimates, capturing noise instead of converging to the mean value. The knowledge on how likely things are to change, which is called the volatility, can also be taken into account to optimally scale the learning rate (Behrens et al., 2007).

#### *2.2.1.3.2.2 Neural implementation*

To keep things simple, learning put forward two concepts: expectation and update.

Neuroeconomics is particularly concerned with expected value signals that correspond to anticipations of the outcome delivery. These expectation signals are pervasive across recording and analysis technics and species. For instance in the monetary incentive task with delayed outcomes, the participant is presented with a cue corresponding to an outcome that is probabilistic and delivered after a delay, so that it leaves time to search for signals that represent the expected outcome that may, or may not, be eventually delivered. The BOLD signal measured by functional MRI supported that regions of the brain valuation system encode this expectation (Knutson et al., 2005). This anticipation of reward within this system is also supported by pattern analysis of BOLD signal over voxels (Kahnt et al., 2010). Using a very different signal, the frequency content of surface electrophysiological recordings in human, it was shown that the beta band activity (20-30 Hz) above frontal areas also indexes the size of the expected value (Bunzeck et al., 2011). Finally, the recording of firing rate of neurons in the raphe-nucleus suggests that some of these neurons encode the progression into the task toward the expected reward (Bromberg-Martin et al., 2010). This short list of studies strongly suggests that our brain computes expectations of the outcome value.

The update of expectations is driven by prediction errors. The finding of this error signal in the activity of dopamine neurons of the ventral tegmental areas drew considerable attention (Schultz et al., 1997; Schultz, 2000, 2006). Dopaminergic projection to the striatum and the prefrontal cortex could guide, with this error signal, the update of learned estimates. Genetic polymorphism of dopamine transporters in the human brain accounts for part of the inter-individual difference in learning ability supporting the key role of this neurotransmitter in learning (Frank et al., 2007, 2009). These results accord well with the idea that phasic dopamine supports an error prediction signal that guides learning. Indeed, these polymorphisms change the tonic level of dopamine, thereby modifying the impact that phasic changes may have. Pharmacological manipulation with dopamine replacement therapy in Parkinson disease patients provided the same consistent evidence (Frank et al., 2004) and the opposite effect of enhancement and reduction of dopamine tonic levels has been

further evidenced with Gilles de la Tourette and Parkinson patients in a subliminal instrumental learning task (Palminteri et al., 2009).

The striatum is a key area in which the activity reflects prediction errors (Montague, 2002; Rutledge et al., 2010). The evidence that this striatal prediction error signal is under dopaminergic control was brought by functional MRI and pharmacological manipulation of dopamine levels in normal human subjects during a reinforcement learning task (Pessiglione et al., 2006). And there is now also evidence in humans that the dopaminergic neurons of the ventral tegmental area support the prediction error signal with high resolution functional MRI (D'Ardenne et al., 2008). The update of estimates that follows prediction errors is not observed just in fMRI, but also in surface electrophysiology, with well-known signatures of error processing in the prefrontal cortex (Boksem et al., 2006; van de Vijver et al., 2011).

#### *2.2.1.4 Different decision systems*

In the above sections on value-based decision making, I presented how the utility of options may be computed and how this estimation guides the selection process to make a decision. The purpose was to capture general ideas that could be useful to understand the effort allocation problem, so that there was more emphasize on common properties than specificities. The only exception was the review of different valuation systems depending on the valence of the utility (cost and benefits) or the nature of the thing to value (e.g. the value of the action leading to an outcome or the value of the outcome). It might also be that there are different kinds of decisions. In principle, this issue arises in the effort allocation problem: are decisions to stop or resume the effort different? Is there an overarching decision to monitor performance and sub-decision systems to decide when to stop or resume the effort? To support the idea that there are several decision systems, I present evidence in section 2.2.1.4.1 that learning from benefit and cost may be underpinned by different systems, second, I discuss in section 2.2.1.4.2 the distinction between different kinds of decision and behavior: Pavlovian, habitual and goal-directed. Finally, I suggest in section 2.2.1.4.3 that there is a hierarchy between decisions, so that decisions can be distinguished based on the time scale and the abstractedness of the goal pursued.

##### *2.2.1.4.1 Learning from cost and benefit: the stick and the carrot*

From a mathematical point of view, the difference between positive and negative utilities (benefit and cost) is just a matter of sign. An intriguing result of the behavioral economics is that it is more than a change of sign: costs and benefits are not symmetrical in the sense that they do not have opposed effect on the decision as could be predicted by just a change of sign. There are at least two noticeable differences that deviate from the symmetry. First, costs are weighted more than benefits

(Kahneman and Tversky, 1979). Second, costs and benefits are not treated equally in terms of uncertainty: people are more prone to avoid risk when a loss is at stake than when it is a gain (Kahneman and Tversky, 1979; Schonberg et al., 2011).

A related idea drew much attention in neuroscience: the fact that we do not learn similarly from our mistakes and our successes. The striatum, as an associative hub, is a key player of learning in the brain (Liljeholm and O'Doherty, 2012). The difference between learning from gain and loss has been mapped with functional MRI during a learning task onto a dorso-ventral gradient of error processing, with losses being encoded more dorsally than gains (Seymour et al., 2007a), suggesting that reward and punishment learning are processed by distinct pathways (Seymour et al., 2007b; Brooks and Berns, 2013). Using similar learning tasks in patients that presented lesions in the brain induced by glioma and Huntington disease, neuropsychological evidence suggested that the dorsal striatum (which is the first target of degeneration in the Huntington disease) and anterior insula were specifically related to impaired learning in the loss domain (Palminteri et al., 2012).

#### 2.2.1.4.2 Pavlovian, habitual and goal-directed behaviors

Different kinds of behavior are often distinguished, which correspond to different historical fields of research. Besides the historical reason, it is a shared consensus that these behaviors are underpinned by different systems (Seymour et al., 2007b; Rangel et al., 2008; Dayan, 2012), and reflect very different kinds of associations between our actions and the outcome of these actions. Pavlovian behaviors are much automated; it is likely that they were implemented in our brain through evolutionary selection. Examples of Pavlovian behaviors are the tendency to approach food or to flee on seeing a snake. Habitual behaviors are learned by individuals, they are also fairly automated, like heading toward the coffee machine when getting in the office in the morning. Goal-directed behaviors are not learned like habitual behaviors; rather, they required the different options to be evaluated and selected at the time of the decision. There is evidence that these different kinds of behaviors are processed by different brain regions, in particular in the striatum, respectively in the ventral part, dorso-lateral part, and dorso-medial part (Liljeholm and O'Doherty, 2012).

#### 2.2.1.4.3 Monitoring different goals

The essence of value-based decision is to commit choice, to select the best response. This implies the ability to accord our actions with respect to our goals. This general ability is termed cognitive control in psychology. It has been shown that more sophisticated, abstract constraints on this control require more anterior parts of the prefrontal cortex (Koechlin et al., 2003) and that this control integrates our motivations (Kouneiher et al., 2009). Cognitive control may be a key ability to plan actions. In particular, how to arbitrate between immediate goals and more distal ones? The anterior posterior

gradient of cognitive control has been proposed to underpin the monitoring of goals at different time scales (Koechlin and Hyafil, 2007). Beside the temporal dimension, it is also common that different goals are pursued in parallel, at the same time. It was also proposed that there is an anatomical segregation in the prefrontal cortex of the processing of concurrent goals (Charron and Koechlin, 2010). We encountered previously this notion that the prefrontal cortex could track simultaneously distinct types of actions, such as exploiting vs. exploring (Kolling et al., 2012), though this mapping in the prefrontal cortex is not fully consistent with the one proposed by the team of Etienne Koechlin.

## 2.2.2 Effort-based decision making

In the effort allocation problem, the decisions to stop and resume the effort may be underpinned by monitoring the dynamic of effort utility. Given the effortful nature of the task, costs are likely to play a key role in this process. Effort is often distinguished from other kinds of cost in the decision-making literature, so that a dedicated sub-field emerged: effort-based decision-making. In the following, I present evidence that physical effort is treated as a cost in the economic sense, based on rodent, human and non-human primate data (in section 2.2.2.1) and then I discuss whether there are different kinds of effort cost (in section 2.2.2.2).

### 2.2.2.1 Physical effort is treated as cost in decision

There are several experimental paradigms that present animals with choices that trade effort against benefit. These paradigms are adapted to the species investigated and brought valuable behavioral data to characterize effort as a cost, i.e. as something that is avoided. Studies with rodents enabled the extensive use of pharmacological or physical lesions as well as unit recording; some data were also obtained from monkeys; so that together with neuroimaging data from humans, the key neural processes involved in effort-based decision making are getting progressively uncovered.

#### 2.2.2.1.1 Behavioral evidence

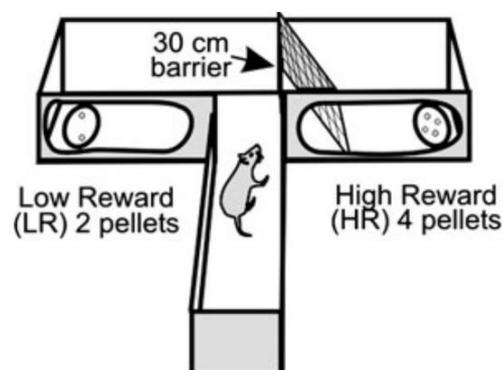


Figure 7 T-maze: an effort-benefit choice task.

From (Floresco and Ghods-Sharifi, 2007)

The T-maze task is often used in rodents: the animal is introduced in the maze and has the possibility to go in two other arms, hence the T-shape of the maze. The two arms offer a reward, e.g. food pellets, and reaching this reward requires a certain effort, e.g. because there is a barrier more or less difficult to climb to get the reward. Different combinations of reward and effort magnitudes are used, like small reward and low effort vs. high reward and high effort, so as to investigate how the animals weight reward and effort. Note that the animal knows which combination is presented in each arms because there is prior training; this paradigm therefore confounds value-based decision and learning processes. The choice made in this paradigm revealed that effort and reward have opposed effect on choices (Cousins et al., 1996; Walton et al., 2002, 2003, 2006; Schweimer et al., 2005; Floresco and Ghods-Sharifi, 2007; Hauber and Sommer, 2009; Hillman and Bilkey, 2010).

Other paradigms involved the rat to press a lever to get a reward, and both the magnitude of the reward delivered and the difficulty associated to the lever press were manipulated experimentally to investigate choices (Walton et al., 2006). Similar level press paradigms are also used in monkeys (Walton et al., 2006); again in this case the effort difficulty (e.g. more lever press needed) has an effect opposed to reward on several motivational markers, such as the error rate and appetitive response, (Bouret et al., 2012). When the error rate was compared in a very simple contrast between performing effort vs. waiting, the comparison of effort and wait matched in durations revealed that monkeys were at most indifferent, or work averse (Minamimoto et al., 2012).

In humans, the effort was manipulated for instance in grip tasks (Kurniawan et al., 2010; Prévost et al., 2010; Burke et al., 2013), in which subjects had to squeeze a handgrip and the difficulty required was varied, or clicking on targets on a screen, more targets to cancel serving as a manipulation of the difficulty (Croxson et al., 2009). These tasks also revealed that effort and reward had opposite valence on behavioral outputs, such as choice or liking rating.

Note that in principle, it is not obvious that effort should be treated as a cost. Utility estimates are highly subjective, as emphasized previously (see page 21). In particular, higher difficulty levels are not always perceived as something to avoid, for instance when we look for achievement or self-esteem (Brehm and Self, 1989). However, in experimental paradigms in which both reward and effort are manipulated, they tend to be treated as benefit and cost (Kurniawan et al., 2011).

#### 2.2.2.1.2 A distinct functional anatomy for physical cost and benefit

Pharmacological experiments in the T-maze task with rodents supported the role of catecholamine compounds (Schweimer et al., 2005), and in particular dopamine (Cousins et al., 1996; Salamone et al., 2007; Salamone and Correa, 2012) to trade effort against reward. In the absence of dopamine, in

particular in the striatum and the anterior cingulate cortex, rodents are biased toward low reward (and low effort) options. In other words, dopamine seems needed to put more effort in order to get higher reward. Note that these results were controlled for any motor impairment that could have been a plausible alternative explanation. The role of dopamine is critical in the cost-benefit integration in several regions (Boksem and Tops, 2008; Salamone and Correa, 2012), in particular in the anterior cingulate cortex (Schweimer et al., 2005). Unit recordings in this latter region revealed that this integration is updated dynamically in the course of the task in rodent (Hillman and Bilkey, 2010). Functional MRI data also supported the integration of cost and benefit in this area (Croxson et al., 2009; Prévost et al., 2010). The anterior cingulate cortex projects to the nucleus accumbens (Hauber and Sommer, 2009), which is another critical brain area for the cost-benefit integration (Walton et al., 2006; Brooks and Berns, 2013; Kurniawan et al., 2013). The anterior cingulate cortex receives input of cost signals from the basolateral amygdala (Floresco and Ghods-Sharifi, 2007) and the insula (Boksem and Tops, 2008). This latter region is a cortical output to the spino-thalamic track that signals the body state, and supports the sensation of this state, a capability called interoception (Craig, 2002, 2003, 2009a).

#### 2.2.2.1.3 Distinction between anticipatory and instrumental costs

The neural representation of the effort can emerge in different contexts, such as the anticipation of a cost and the experience of it (Kurniawan et al., 2011, 2013). The neural representation of these two kinds of costs may be different. Kurniawan et al. noticed that the anticipation of effort involved the anterior cingulate cortex and the dorsal putamen, whereas the integration of cost and benefit are more related to the instrumental role of effort in the reward delivery, involving more ventral parts of the striatum. These distinctions could resolve contradictory results about the striatum that was reported to be more activated (Schmidt et al., 2012; Burke et al., 2013), less activated (Kurniawan et al., 2010) or not affected, at least for dopamine levels (Gan et al., 2010), for higher difficulty levels. Following this distinction, Salamone suggested that the role of the striatum and in particular the striatal dopamine is specific to the case when the effort is instrumental in the response (Salamone and Correa, 2012). Rangel and Hare, from a more general point of view on value-based decision-making, suggested that instrumental costs (i.e. specific to the action that leads to the outcome) and aversive outcomes should be fundamentally distinguished (Rangel and Hare, 2010). This stance makes efforts somehow special among other costs.

### *2.2.2.2 There are different kinds of effort-related cost*

#### *2.2.2.2.1 Physical vs. mental effort*

What is special about effort cost? There may not be a unique kind of effort cost, but rather plenty of them. A first distinction which may come to mind is between physical and mental efforts: are they different for the brain? Obviously, the nature of these efforts is different, so that different brain areas are involved in underpinning these different kinds of effort. Different cost levels within mental and physical efforts have thereby distinct signatures, e.g. the weight load may correlate with the activity of the motor cortex whereas the attentional load may correlate with the activity of the dorso-lateral prefrontal. These distinct correlates still leave open the question whether the representation of these cost levels to guide decision and choice are different across effort types. An intriguing result is that many of the regions involved in coding the cost level of physical effort are also involved in their mental counterpart. For instance, the anterior cingulate cortex is critical to signal physical effort, this region is also more activated for higher cognitive demands (Bush et al., 2000). The insular cortex that signals the body state and the physical cost is also involved in processing monetary costs (Palminteri et al., 2012). It was outlined earlier that the lateral prefrontal cortex is more likely than the medial prefrontal cortex to encode costs because it receives specifically more sensory inputs. It turns out that the lateral prefrontal cortex is also involved in signaling the cognitive cost (McGuire and Botvinick, 2010). The nucleus accumbens is another example of common substrate: besides its implication in the integration of physical costs against benefits, this region may also support the cognitive effort discounting of reward (Botvinick et al., 2009).

Overall, these parallels suggest that, though different in nature, physical and mental effort cost level could share similar processing. Boksem and Tops suggested that this could be a case of evolutionary recycling: the circuits dedicated to physical effort-based decisions could be reused for mental effort (Boksem and Tops, 2008).

#### *2.2.2.2.2 Physical effort and covariates*

The similar network for mental and physical effort-related cost could be the common apex of pathways processing these costs; however, the roots of these pathways may be numerous and correspond to several parameters that co-vary with effort exertion. These parameters differ between effort types. For instance, physical exertion goes along with changes in proprioception and these changes inform the brain about the cost associated to the effort: the weight of the load to lift, the velocity of the effort, etc. This cost may even be painful. Pain is a common aversive stimulus that is often used in experimental designs to study cost-benefit decisions (Brooks and Berns, 2013). Pain is not necessarily entailed by physical effort, but may co-vary with effort parameters, e.g. the excessive

weight to lift. The relationship between pain and effort cost is stressed by the existence of common brain networks to track pain and cost in the general sense, comprising in particular the insula and the anterior cingulate cortex (Peyron et al., 2000; Craig, 2003; Wager et al., 2013) and common neuromodulators such as serotonin (Seymour et al., 2007a; Smith, 2009; Lindstedt et al., 2011; Dayan, 2012).

Time is another covariate of exertion: efforts have duration. This statement may seem trivial but it should be kept in mind because duration makes effort costly for several reasons: longer efforts are *per se* more costly, but also indirectly more costly because they potentially delay the reward. Time is thereby a critical parameter, out of many, that controls effort cost levels.

The list of effort covariates could also include physiological quantities such as the body temperature, the heart beat rate, the carbohydrate resource levels, etc. Overall, it seems that many factors co-vary with effort. They might be epiphenomena or contribute directly to effort-related costs. To better understand what makes the effort costly, I review how efforts are produced in section 2.3 and how time steps in the problem of estimating cost, both through the discount of reward (decrease of benefit) and the build-up of fatigue (increase of cost) in section 2.4.

### 2.2.3 Summary

- When allocating effort over time, like in the fridge problem, it is crucial to make the right decision in choosing when to take a break.
  - Decisions are characterized by the cost and benefit entailed, that are weighted against each other.
  - Computationally, the comparison of cost and benefit is captured by a single quantity: the utility, which is the net benefit (i.e., benefit discounted by the cost).
  - The brain has a core valuation system to encode utility; however there are also specific encoding processes depending on the nature of the option to value (e.g. positive vs. negative, money vs. food, goal vs. action leading to the goal, etc.)
- There is something obvious we may dislike in moving a fridge: it is effortful.
  - Physical efforts are treated as costs by decisional processes, especially when they are manipulated along with benefits.
  - Effort costs can be distinguished from other cost types although the neural encoding of cost levels may be in part common across kinds of cost.
- How effortful it is to move a fridge may be a cost we do not know beforehand.
  - Expected cost and experienced cost may be different. In general they have distinct and dedicated neural representations.
  - Expectations are confronted to experience so that values and decisional policies can be updated through learning.
- Allocating effort over time requires decisions about the timing of actions.
  - Drift diffusion models account for both the decision outcome and the dynamics of the decision process .
  - These models are used to account for perceptual and value-based decisions. In these models, the incoming evidence is integrated and the decision is made by comparing the accumulated evidence to a pre-determined threshold.

## 2.3 How efforts are produced

Given that costs entailed by effort may be crucial determinants of the effort allocation problem, and that how efforts are produced may impose constraints onto how this problem is solved, it is worth to better understand effort production. Two notions characterize effort. On the one hand, physical effort is a quantifiable physical production of the body. It requires a bodily activation, in particular from the muscle, to run, lift, push, etc. On the other hand, the notion of effort also points to the activation process itself and the feeling it entails: how effortful is an action corresponds to a subjective estimation of how difficult this action is. This section is split into these two components of effort: the physiology of the motor production (section 2.3.1) and the physiology of the motivational process that drives the behavioral activation required by the effort (section 2.3.2).

### 2.3.1 The motor command

This review does not aim at digging in the details of the muscle and the motor command physiology without purpose, but to specifically understand the two following processes. First, how can biological signals recorded non-invasively be related to the motor command execution and its preparation? A particular attention will be given to non-invasive whole-brain technics such as functional MRI and electrophysiology, which are used routinely in neuroscience with healthy human participants. Second, where does the exhaustion come from during physical exertion? Basic pieces of knowledge to understand fatigue and exhaustion are presented here with a particular emphasize on proprioception and nociception. These two signals are intimately related to effort exertion but are fundamentally different. The regulation of these afferent signals is tightly related to the efferent signals of the motor command.

#### 2.3.1.1 *The exercising muscle*

The ultimate effector in physical effort is the muscle. How do skeletal muscles produce force?

##### 2.3.1.1.1 Force and speed in the muscle physiology

###### 2.3.1.1.1.1 *Different types of skeletal muscles*

The skeletal muscles are the muscles attached to bones and used to produce movements. Different muscles have different output: some produce force steadily whereas some others are activated abruptly and intensely. These muscles look different: the former are darker (the red fibers) and the latter whiter (the white fibers). A good example of this distinction can be found easily in any cooked chicken: the muscles of the leg are darker, used constantly for posture and walk whereas the muscles of the breast are white, used only occasionally (chickens do not fly) to produce short and vigorous movements in order to flee or frighten another animal. Actually, the red and white fibers are mixed

in all muscles, but in different proportions. This distinctive appearance is due to the metabolism that determines the muscle functional properties (Loeb and Ghez, 2000).

Red fibers (also labeled type I fibers) are fibers that produce low force levels with slow rate of change but that are also quite insensitive to fatigue, i.e. they can sustain exertion during long periods. The reason is that they are specialized to produce a very efficient oxidative catabolism. In other words, they produce energy by breaking down energy-rich molecules using the oxygen supplied by the blood. They are supported by an extensive network of capillaries and are also specialized to extract blood oxygen efficiently (hence their dark color) so that virtually, their energy supply depends only on the blood, which explains why they do not run out of energy and are able to sustain exertion.

White fibers (also labeled type II fibers) are faster and produce more force, but at far shorter time scales. The classification is further refined into type II A&B fibers. Type IIB fibers are highly fatigable and powerful. Their metabolism is mostly anaerobic and produces energy by breaking down large internal stores of glycogen, so that they do not depend on blood supply, but also stop the effort production when their internal resources are depleted. These resources then take hours to recover. Type IIA fibers are also powerful, but less fatigable because their metabolism is more oriented to aerobic catabolism than type IIB fiber, so that they have extra resources from blood supply. However, their production is still limited to several minutes.

#### *2.3.1.1.1.2 Recruitment of muscular units*

The muscle fibers are recruited following a strict order, so that fast, highly fatigable and powerful fibers (type II) are recruited only after the slow fibers when extra force is needed. This progressive recruitment is due to the properties of the motoneurons that innervate these fibers: they differ in their size. The size determines the excitability: motoneurons of type II fibers require higher input levels to be excited, and are thereby recruited after the motoneurons of type I fiber because they have a lower excitation threshold (Loeb and Ghez, 2000). Functionally, this recruitment goes from weaker (and fatigue insensitive) to stronger (and fatigable) fibers, so that the muscular resource is preserved until it is really needed.

The force produced by a muscle depends on two critical parameters (on top of the muscle type recruited): the amount of fibers recruited and the firing rate of individual fibers. To increase force production, both the amount and the firing rate of activated fibers increase. This increase is driven by a parallel increase in the recruitment of motoneurons and their firing rate. The muscle excitation is driven by motoneurons, this signal is the muscle efferent signal and it is actually supported by a specific subclass of motoneurons: the alpha class.

#### 2.3.1.1.2 The afferent signal: Feedback from the muscle

The muscle is driven by efferent signals, but it also sends back signals to the spinal cord that represent its physical state. These afferent signals can be functionally split into proprioception and nociception. These two sensory signals are first described; then I discuss why such a distinction is relevant.

##### 2.3.1.1.2.1 *Proprioception*

The British physiologist Charles Sherrington made valuable contributions to the understanding of the muscular reflex. He was co-awarded a Nobel Prize in physiology, with Edgar Adrian, for this contribution. He showed that muscular reflex stems on a loop between the muscle and the spinal cord, in which changes in the muscle state are signaled by the muscle to the spinal cord and the spinal cord sends corrective signals in return to the muscle (Pearson and Gordon, 2000). These signals sent by the muscle are not specific to the reflexes but are pervasive. They are sensory signals generated by the body's own movement, something Sherrington termed proprioception. Proprioception arises from two sensory structures: the muscle spindle that signals the muscle length and the Golgi tendon organ that signals the muscle tension (Pearson and Gordon, 2000).

The muscle spindle is an encapsulated structure made of intrafusal muscle fibers. These fibers are specialized in the sense that they are barely contractile and do not contribute to the muscle contraction. Instead, the center of the non-contractile fiber is surrounded by afferent nerve endings that fire when the fiber is stretched: they signal the muscle length. The muscle spindle is involved in the stretch reflex that we probably have all experienced at the physician's: this is when the knee is gently hit by a hammer, which triggers a jerk of the leg. In this case, the spindles are stretched by the hit and the spinal cord sends a corrective signal to contract the muscle, which produces the jerk. Interestingly, the spindle also receives an efferent signal from the gamma motoneurons. This input, as opposed to the alpha motoneurons, does not contribute to the muscle contraction. They induce small contractions of the spindle polar region, so that the length (hence the response) of the spindle depends on the global contraction of the muscle, and on this additional intrinsic contraction. The gamma motoneurons input thereby participates to the regulation of the excitation of the spindle.

The Golgi tendon organs are not located in the fleshy part of the muscle (like the spindles) but at the junction between the muscle and the tendon. The adherence between the muscle and the tendon is ensured by numerous collagen fibers. Afferent axons are intermixed within these collagen fibers, so that when there is tension on the muscle, the collagen fibers are stiffened and compress these axons, which translates the muscle tension into firing rate. Note that as opposed to the spindle, it is not per se the length of the muscle that impacts the Golgi tendon organ but the muscular tension.

### 2.3.1.1.2.2 Nociception

There are other sensory receptors in the muscle that are not associated to a mechanical structure like the spindle or the Golgi tendon organ. The nerve endings are somehow free, and sensitive to different kinds of stimulation, such as chemical stimulation (e.g. acidification) or temperature (Gardner et al., 2000; Pearson and Gordon, 2000).

### 2.3.1.1.2.3 Separate pathways for proprioception and nociception

Besides their sensory receptors, the muscle afferents can be sorted by their axon type. Mechanical signals are conveyed by large myelinated fibers: type Ia for the primary spindle ending, Ib for Golgi tendon organs and type II for secondary spindle ending and non-spindle endings, and smaller but still myelinated fibers (type III, for some nociceptive fibers) and also smaller, non-myelinated fibers. The size and myelination of the fiber determine the signal conduction speed, so that an equivalent nomenclature has been proposed based on the conduction speed: A $\alpha$ , C, ... the following chart lists the afferents and the nomenclature.

	myelinated	Size nomenclature	Size ( $\mu$ m)	Velocity nomenclature	Velocity (m/s)	Sensory class
Proprioception	Yes	Ia	12-20	A $\alpha$	72-100	Primary spindle ending (length)
	Yes	Ib	12-20	A $\alpha$	72-100	Golgi tendon organ (tension)
	Yes	II	6-12	A $\beta$	36-72	Secondary spindle ending (length)
	Yes	II	6-12	A $\beta$	36-72	Non spindle ending (deep pressure)
Nociception	Yes	III	1-6	A $\delta$	4-36	Free ending, stretch, chemical stimuli, temperature
	No	IV	0.2-1.5	C	0.4-2	chemical stimuli, temperature

Table 1 Afferent fiber classes. Based on (Gardner et al., 2000; Pearson and Gordon, 2000)

Etymologically, nociception means the feeling that there is something wrong and harmful. The nociception afferents differ from the proprioceptive afferents in their sensory receptors, their diameter, their conduction velocity and their myelination, which supports a functional distinction between both types of afferents.

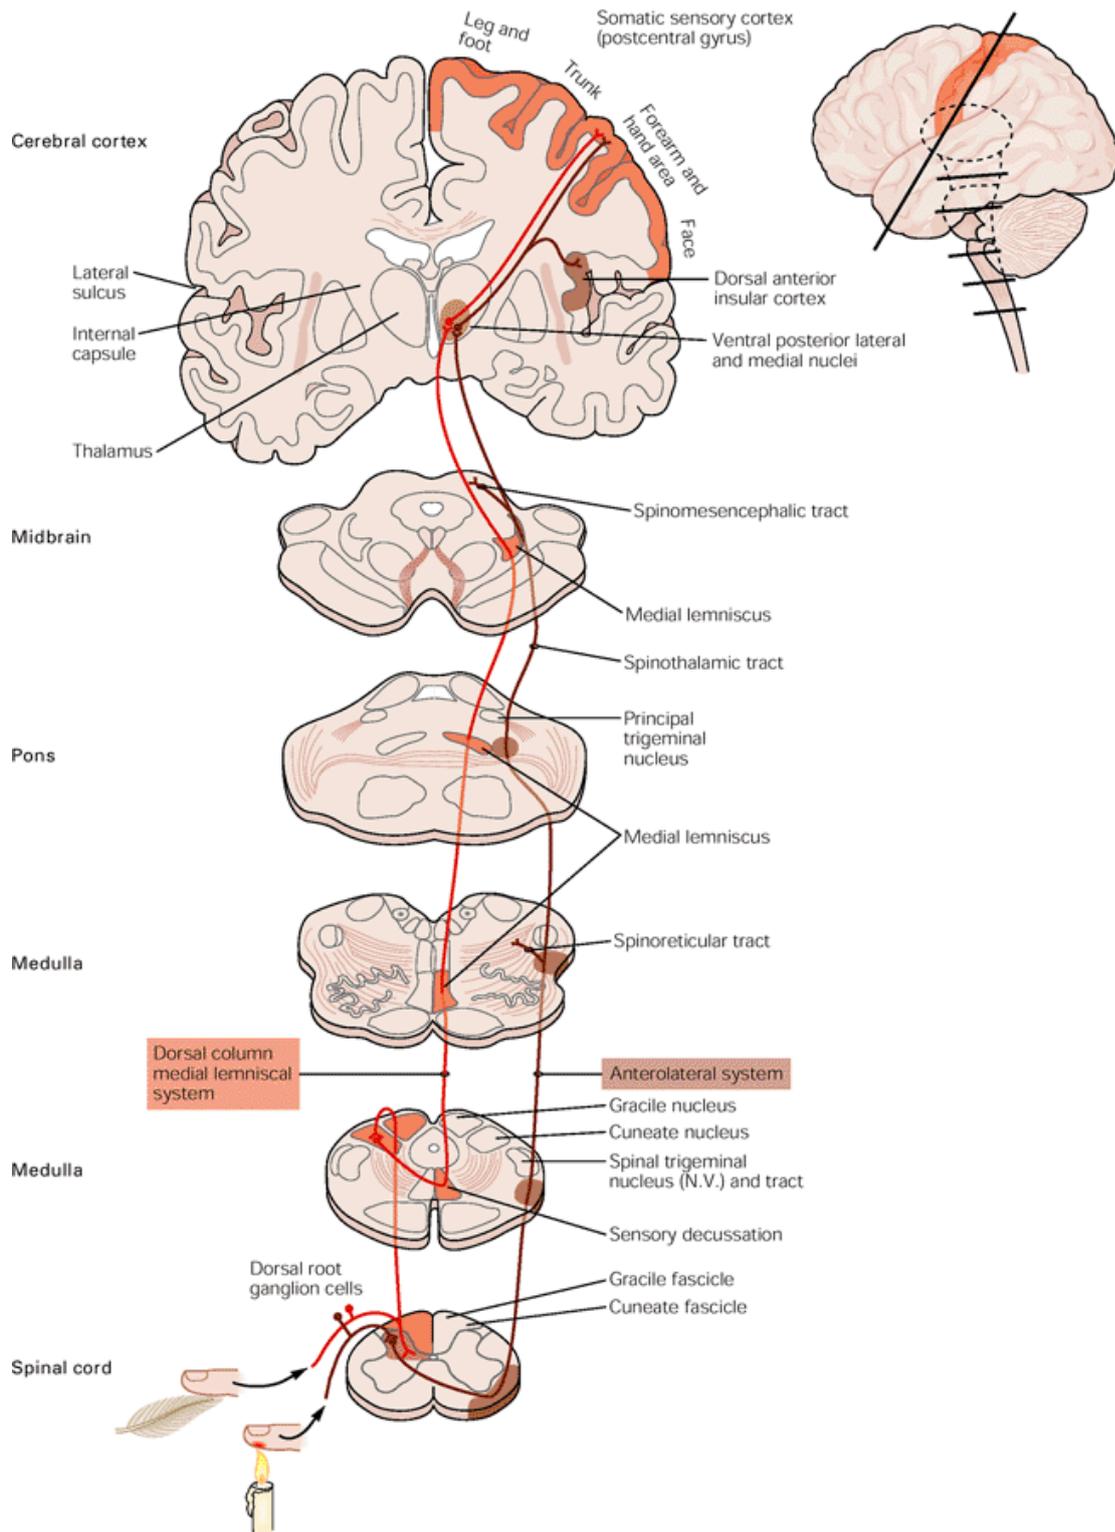


Figure 8 Distinct pathways for nociception and proprioception. From (Gardner et al., 2000)

Interestingly, the functional distinction is also supported by different pathways to the central nervous system, see Figure 8 (Gardner et al., 2000):

- Proprioception: the dorsal column medial lemniscal system

- Afferent enters the spinal cord through the ipsilateral dorsal column
- Decussation in the medulla (inversion of the body map: left fibers go to the right & the other way around)
- Medial lemniscus of the pons
- Medial lemniscus of the midbrain
- Ventro-posterior medial nucleus of the thalamus
- Somato sensory cortex
- Nociception: the anterolateral system
  - Afferent fibers enter the ipsilateral dorsal horn of the spinal cord, they then project to the contralateral part through a synaptic relay to terminate in the anterolateral track (inversion of the body map)
  - Follows the spinothalamic tract
  - Ventro-posterior lateral nucleus of the thalamus
  - Somato-sensory cortex.

These distinct pathways, with a contrast between ipsi and contralateral representations and also dorso-ventral separation explain why nociception and proprioception can be differentially altered by partial spinal injuries.

More recently, a third pathway was put forward to convey interoceptive signals, i.e. signals that inform the brain about the state of the body, how 'well' are body physical conditions. These afferent signals are functionally and anatomically distinct from the nociceptive and proprioceptive signals: they are conveyed by the laminal I spinothalamocortical pathway (Craig, 2002).

#### 2.3.1.1.3 Cross talk between efferent and afferent signals

Charles Sherrington made a significant contribution in putting forward the role of the spinal cord. A major advance since was the discovery that the spinal reflex is also modulated at supraspinal levels, i.e. by the brain. In fact the spinal cord is a major site for integrating reflexes with the central commands and goal-directed movements (Pearson and Gordon, 2000). In particular, efferent (top-down command) and afferent signals interact at the spinal cord level. Two examples are provided here to support this idea of cross-talk between these signals. During the muscle contraction driven by alpha motoneurons, there is a parallel activation of gamma motoneurons to regulate the sensitivity of muscle spindle: thereby, the efferent signal modulates the afferent signal. Another example is the case of type Ia inhibitory interneurons. These neurons receive inputs from the descending central command and participate in the regulation of the balance between the agonist antagonist muscle activation. These neurons also receive inputs from muscle afferents so that the level of activation of

motoneurons is regulated directly and indirectly by the descending pathway and the feedback signal (Ghez and Krakauer, 2000a; Pearson and Gordon, 2000). For instance the spindle return can regulate the type Ia inhibitory activity and facilitate or silence alpha motoneurons activity. The silencing occurs for instance in the stretch reflex to inhibit the antagonist muscle so that only the agonist (and stretched) muscle is contracted (Pearson and Gordon, 2000).

### *2.3.1.2 The motor drive*

#### *2.3.1.2.1 The primary motor cortex drives motor activation*

How does the central nervous system drive muscle contraction? Direct electrical stimulation of the cortical surface in experimental animal preparation or during human surgery (e.g. the work by Penfield & Jasper) revealed that motor responses can be induced by stimulation only in specific cortical areas, mainly on the so-called primary motor cortex (M1, Brodman Area 4) and to a lesser extent on the premotor cortex (Brodman Area 6) (Ghez and Krakauer, 2000b); the proportion of corticospinal neurons from the supplementary motor area is estimated to be only 10% of the corticospinal tract (Nachev et al., 2008). Stimulation studies revealed that there was a body map on the rolandic sulcus: stimulation at precise location along this sulcus induces the activation of similar muscle and movement across individuals. This tight coupling between the motor cortex activity and muscle activation is due to the presence of neurons projecting directly to the spinal cord motoneurons: the cortico-spinal neurons. There are two pathways (Ghez and Krakauer, 2000b):

- The ventral corticospinal tract: the innervation of the muscle of the neck and the trunk, to control the posture.
- The lateral corticospinal tract: including cortico-spinal neurons from motor, premotor (BA 4 and 6) and sensorimotor cortices (BA 1, 2, 3), for the control of limb. In this pyramidal tract, there is a decussation in the brain stem, which makes that the body map is inverted, with cortical stimulation activating contralateral muscles.

The motor cortex can also be stimulated non-invasively with transcranial electric or magnetic stimulations. This method is used in healthy humans to assess the cortico-spinal excitability (McNeil et al., 2013). The comparison between transcutaneous stimulation at the level of the pyramidal decussation and transcranial stimulation of the motor cortex makes it possible to estimate the relative excitability of the motor cortex and the cortico-spinal tract, however, this method is much painful for participants.

It is believed that M1 codes primitives of the movement and that details (for instance the muscle agonist - antagonist activation balance) are controlled at the spinal level (Scott, 2004). In this view,

M1 is involved in representing motor execution at a high level. However, another model posits that there are too many and complex correlations between the motor parameters (speed, direction, load, etc.) for details to be resolved at the spinal level so that low-level properties of muscle activation must be encoded in M1 (Todorov, 2000). The same author also suggested that the sensory feedback is crucial to control movement (Todorov and Jordan, 2002; Todorov, 2004). The issue of whether fine-grain motor description of muscle activation is resolved at the cortical or spinal level might be reconciled with the very idea that sensorimotor integration is crucial to determine the motor control because this integration occurs both at the cortical and the spinal level. Supporting the idea of a fine-grain control in M1, modeling approaches of the motor control suggested that the activity of neurons in the motor cortex encodes the motor drive, i.e. the input to the motoneurons, rather than particular motor parameters such as movement, force, etc. (Guigon et al., 2007).

Finally, the view that the descending motor command simply drives muscle activation should be rounded. The detailed physiology of the spinal cord revealed some complex connectivity patterns, such as the recurrent inhibition of motoneurons by Renshaw cells (Pearson and Gordon, 2000). Renshaw cells are inhibitory neurons, and motoneurons make polysynaptic connections to the muscle they activate and to the Renshaw cells that, in return, inhibit these motoneurons. This negative feedback implements a limitation of the command. This feedback is involved in the stabilization of the motor command, but also in the muscular fatigue, a topic discussed later.

#### 2.3.1.2.2 Functional recording of the primary motor cortex during force production

It was previously mentioned that the muscle produces higher force levels through an increase of the amount of fibers recruited and an increase of their firing rate. Early reports, for instance by Edward Evart in the 1960s showed with unit recordings in the monkey primary motor cortex that the firing rate increases along with the force level produced (Ghez and Krakauer, 2000b). Ever since, the relationship between the activity of the primary motor cortex and the force level produced showed several discrepancies in the functional MRI literature. More precisely, the debate is whether higher force levels should entail more intense or more widespread activations. Some studies found both increased extent and amplitude of BOLD signal for higher force produced (Dai et al., 2001; Spraker et al., 2007), whereas other found increase only in the amplitude (Liu et al., 2003) or the extent (Thickbroom et al., 1998), or none (Ludman et al., 1996). The contrast of power grip and precision grip tended to show higher M1 activations (Ehrsson et al., 2000; Kuhtz-Buschbeck et al., 2008). Some authors have emphasized that the relationship between the force level produced and the motor cortex activation reported by fMRI is highly non-linear (Liu et al., 2002). Some authors also noticed that specific motor parameters can impact the BOLD signal in M1, such as the rate of change of force

(Vaillancourt et al., 2004; Keisker et al., 2010) or the visuo-motor control needed to perform the task (Vaillancourt et al., 2003). Besides some discrepancies in these results, it is clear that the motor cortex is not the only area involved during motor execution: somatosensory, premotor and parietal cortices as well as subcortical structures are also activated (Ehrsson et al., 2000; Dai et al., 2001; Cramer et al., 2002; Liu et al., 2002; Vaillancourt et al., 2003; Spraker et al., 2007; Kuhtz-Buschbeck et al., 2008).

#### 2.3.1.2.3 The issue of sustained vs. transient activations

The discrepancies in the fMRI literature on the activation of the motor cortex for different levels of force produced could be partially accounted for by variations in the motor parameters used in the task, in particular the duration of each effort, and hence, the frequency of the effort. Indeed, there may be distinct responses depending on the duration of the activation, both at the neural level and in the hemodynamic coupling, which could account for these discrepancies.

Transient and sustained BOLD responses were analyzed in a motor task and showed a bias for transient effects (Duff et al., 2007). Similar results were found in other modalities. For auditory stimulations, neural adaptation could explain why there are higher responses for transient than sustained stimulation (Fox et al., 2005; Harms et al., 2005). There might also be an hemodynamic bias: when the same subjects were scanned under both fMRI and magnetoencephalography (MEG) in an auditory stimulation, the BOLD signal better correlated to the transient responses than the sustained responses identified in MEG (Gutschalk et al., 2010a). It is also likely that the hemodynamic in the motor system itself is more prone to transient responses. Hemodynamics models such as the balloon model reveal that there are many reasons for BOLD transients to occur (Ogawa et al., 1993). The parametrization of the balloon model (Nakai et al., 2000) and concurrent recording of blood flow and BOLD signals (Bandettini et al., 1997; Obata et al., 2004) showed evidence for uncoupling between blood supply and consumption in the motor regions (among several others). Such uncoupling produces transient responses.

Overall, these studies showed that sustained neural activity can be translated into sustained BOLD activity but that due to several biases, transient neural events elicit higher BOLD responses. Given that the duration of effort varied between the studies reported above on effort production, these biases could account for some discrepancies.

#### 2.3.1.3 Regulation of the central command

There is a limited set of regions that directly drive the motor command (the motor cortex around the central sulcus) but there are many regions involved in regulating the motor cortex. It seems that the

brain first represents the outcome of the motor action at an abstract level, then translates this representation into the motor code and eventually regulates this motor code online depending on the proprioceptive return and the unexpected perturbations that occur (Ghez and Krakauer, 2000a). To take an example, when writing, we can represent ourselves a letter string, implement a motor program with our usual writing hand but possibly with the other hand, and if somebody moves the sheet we are writing on, we are able to correct our command.

The motor cortex is therefore a hub that converts intentions into actions, regulates the on-going action based on the body state and corrects the command when perturbations are introduced. Distinct brain regions serve these functions. The premotor cortex is connected to posterior parietal regions and the lateral prefrontal cortex. The premotor cortex participates in the formation of motor programs to serve specific goals (Ghez and Krakauer, 2000a). The execution of this command by the motor cortex is performed in interaction with several other areas that can be functionally interpreted as follows (Shadmehr and Krakauer, 2008):

- The computation of an internal model of the motor system and the action to perform involves the cerebellum. Deviations during the course of the action from the predictions of this model and the error signals (visual, proprioceptive) entailed by these deviations lead to corrections that are accurate and fast, by a loop between the motor cortex and the cerebellum through a ventro-lateral thalamic relay. We are unaware of these corrections, provided that the deviation is limited (Frith et al., 2000).
- The monitoring of the proprioceptive return, and hence the motor state of the body, is achieved by the parietal cortex; in particular the posterior parietal cortex integrates information across modalities and conveys it to the motor and premotor cortex.
- The optimization of the motor execution involves the basal ganglia. It is a major site for motor learning (Liljeholm and O'Doherty, 2012), they encode benefits and costs associated to the motor execution and participate to the motor energizing, so that this structure is functionally well suited to optimize motor production (Turner and Desmurget, 2010).

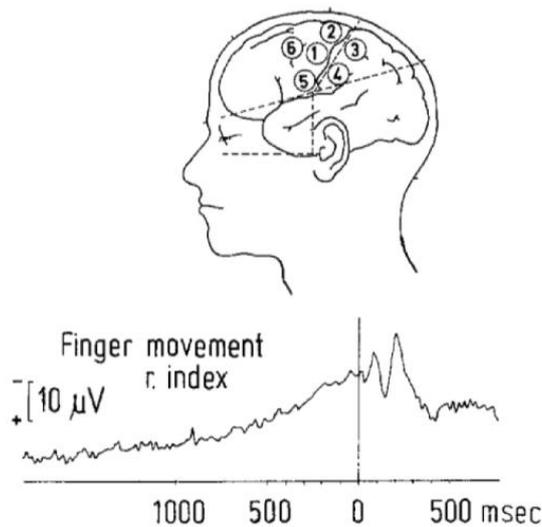
Note that this brief description tries to map the concepts of the optimal control theory: an internal model, the current state and the optimal program to the areas involved in regulating the activity of the motor cortex (Wolpert and Ghahramani, 2000; Todorov, 2006). Note that this integrated view of the motor function is in line with the pragmatic view of cognition that seeks to ground cognition into an action-oriented framework, thereby putting forward the sensory-motor integration (Engel et al., 2013).

### **2.3.1.4 Motor preparation**

An interesting finding with the motor system is that the process associated with motor production is active well before the production itself. There is an anticipatory phase corresponding to the motor preparation, occurring during rest. In the effort allocation problem, this is potentially interesting because the decisions to explain are both when to stop the effort, but also when to start it. Therefore, the timing of the effort resumption could depend in particular on motor preparation. The preparation can be tracked in several neural signals. Here, the so-called readiness potential and the oscillatory desynchronization are presented, together with the increase of cortico-spinal excitability.

#### **2.3.1.4.1 Evidence from the readiness potential**

The readiness potential was first reported in the 1960s in surface electrophysiological recording, as a deflexion in the electric signal recorded above the head when the signal time series are averaged back on the motor output onset, see Figure 9 (Kornhuber and Deecke, 1965; Deecke et al., 1969). This signal was divided into an early bilateral component and a late contralateral component (the late component being just before the movement is produced). This readiness signal is observed only during voluntary execution, e.g. not in tics and is more pronounced when the movement is natural and complex, as opposed to finger tapping for instance (Shibasaki and Hallett, 2006). This signal is modulated by many motor parameters, such as the motor sequence to produce, the effector to use, the force level required (Lang, 2003). Source localization of this signal is centered on the motor cortex, but is not restricted to this area (Leuthold and Jentsch, 2002; Cunnington et al., 2005). Recording in EEG under fMRI and fMRI alone confirmed that the SMA, the parietal lobule and the primary motor cortex produce this signal (Ball et al., 1999; Cunnington et al., 2003). Direct recording in humans with electrocorticogram confirmed the involvement of the rolandic sulcus (Satow et al., 2003).



*Figure 9 The motor readiness potential*

*Participants performed voluntarily finger movement. Electroencephalographic data were averaged with respect to the detection of movement on the mechanogram ( $t=0$ ). The position of the scalp electrode (#1) is reported above, superimposed on the position of the Rolandic sulcus based on anatomical landmarks. From (Deecke et al., 1969).*

The counterpart of this signal in magnetoencephalography (MEG) was identified in 1982 (Shibasaki and Hallett, 2006), approximately ten years after the magnetoencephalography technique was operationalized for brain signals (Cohen, 1972). Source reconstruction of the readiness field (since MEG records fields, not potentials) was consistent with the EEG results (Praagstra et al., 1999; Erdler et al., 2000; Takahashi et al., 2004). Positron emission tomography yielded similar results (Pedersen et al., 1998). The readiness field (RF), like the readiness potential (RP), is impacted by movement parameters, such as the frequency of tapping (Mayville et al., 2005). It was however noticed that the readiness potential and readiness field may not be exactly the same signals (Nagamine et al., 1996). This is probably due to the fact that the electric and magnetic signals are orthogonal one to another, so that the two techniques have complementary sensitivity to currents that are tangential or orthogonal to the skull (Hämäläinen et al., 1993), which is critical for instance in the case of the SMA.

This signal is associated with the motor preparation. It gradually builds up as the motor response is about to be delivered (Gluth et al., 2013). Since the seminal work of Benjamin Libet on how the readiness potential relates to the awareness of the voluntary movement preparation (Libet et al., 1983), much work and controversies have been produced. It now seems that the intention to move is encoded in the parietal lobe and that the SMA is related to the urge to move (Desmurget and Sirigu, 2009, 2012). However, it was also recently suggested that the readiness potential could signal the random fluctuations of a 'go' signal when the temporal pressure is minimal (Schurger et al., 2012).

#### 2.3.1.4.2 Evidence from motor beta synchrony

The readiness potential and the readiness field (RP, RF) are not the only electrophysiological signals that precede motor production. Desynchronization of electrophysiological signals above the motor cortex in the range of mu (circa 10 Hz) and beta bands (13-30Hz) is also a typical signature that precedes movement initiation (van Wijk et al., 2012). Early reports that the beta band activity is suppressed during, and just before, motor exertion came from direct recording at the cortical surface during human surgery (Jasper and Penfield, 1949).

The desynchronization does not seem to be related to the RP or RF, since they can be observed independently across subjects (Feige et al., 1996). A neural account for this dissociation could be found in different substrates: the desynchronization would originate from small neurons connected to the thalamus and the motor cortex – basal ganglia loop whereas the RP and RF originate from pyramidal neurons (Shibasaki and Hallett, 2006). The source of this signal at the cortical level is focused on the contralateral rolandic area (Jurkiewicz et al., 2006; Tzagarakis et al., 2010), but with significant extension to the ipsilateral hemisphere (Hummel et al., 2003). It is observed both in EEG and MEG. Contrary to the RP / RF, it is difficult to find movement parameters that impact consistently the desynchronization level (Stancák et al., 1997; Cassim et al., 2000; Ehrsson et al., 2000; Tzagarakis et al., 2010).

Interestingly, this preparatory signal corresponds to the gradual commitment toward response, even in complex decisions. For instance, when evidence for a perceptual decision was delivered progressively, the perceptual evidence was sampled gradually and the beta desynchronization followed incremental levels of evidence (Wyart et al., 2012). In another perceptual task where the perceptual evidence was tracked directly in the brain with the gamma activity in the posterior cortex, the beta desynchronization also built up gradually together with the evidence (Donner et al., 2009). This result was also observed in another task, in which the motor beta desynchronization also reflected the decisional evidence, and in this case, it even reflected the bias that were induced experimentally on the decision (de Lange et al., 2013). A study however showed that this relationship between decisional evidence and motor beta desynchronization was observed only when a response was needed, suggesting that the integration of evidence and the gradual motor preparation are two distinct processes (O'Connell et al., 2012).

#### 2.3.1.4.3 Evidence from the electromyogram

The readiness potential and the beta desynchronization favor the corticospinal excitability (Schoffelen et al., 2005; Engel and Fries, 2010). Given the previous reports that these signals are also associated with increases in decisional evidence, the cortico-spinal excitability should also increase

during this process. Consistent evidence was revealed by at least two studies using different methodologies. One of these studies used transcranial magnetic stimulation of the primary motor cortex to estimate the excitability of the corticospinal tract. The authors found that in a choice task, faster responses correlated with speeded build-up of the cortico-spinal excitability (Klein-Flügge et al., 2013). In the other study, the level of accumulated evidence was manipulated in the random dot motion task presented earlier (see p. 29) and the excitability at the spinal level was estimated on the electromyogram after perturbation that were introduced randomly to induce reflex response (Selen et al., 2012). Again, the response reflected a motor preparation that was commensurate to the accumulated level of evidence.

## **2.3.2 The motivation to drive the effort**

Being able to do the effort is one thing; doing it is another thing. Why do I have to move this fridge up through the stairs? If I really want to do it, it is likely that I can put more effort in this action, thereby impacting my effort allocation. But how is more motivation translated into more effort? The previous section was dedicated to how the motor command activates muscles to produce force, but how and why this motor command is activated is the issue of this new section. Given that the concept of motivation has a strong historical background in psychology, some preliminary words are aimed at defining motivation as intended in this work (section 2.3.2.1), then, I describe the neural underpinning of motor motivation (section 2.3.2.2) and last I mention some deficits in this process with a few clinical examples (section 2.3.2.3).

### **2.3.2.1 What is motivation?**

#### **2.3.2.1.1 Intrinsic and extrinsic motivations**

Motivation is a word widely used in the lay language, in clinics, e.g. in the Starkstein scale of apathy (Starkstein et al., 1992), in psychology and more recently in neuroscience. A brief historical presentation of the concept is provided, inspired by (Berridge, 2004). Motivation is about why we do what we do. To answer this question, the notion of drive was put forward at the time of behaviorism. From the crudest perspective, the behaviorist approach is based on observable variables and tries to list the associations between stimuli (that can be controlled experimentally) and behavioral responses (that are measurable). Given this approach, one can notice some associations like: the room temperature is high and somebody takes a bottle from the fridge, or a very salty solution is given to a rat and after that she will drink more water, or after being deprived from water while visiting a city, one can pay at the bar to get a drink. Rather than considering that these stimulus-response associations are independent, we can use the concept of thirst to simplify the story: in all cases, the participant is thirsty and seeks to drink, even if it is costly. In that sense, thirst is a drive, an

intervening variable between stimulus and response. Referring to the notion of drive makes things simpler, instead of considering the full combinations between all these stimuli and responses as unrelated. Referring to this concept also gives some explanatory power to the description of the behavior. It is not circular as it could seem at first sight: we drink because we are driven to drink. Actually, the concept of thirst enables to make predictions between new stimuli and new responses that are likely to be met by observation. This notion of drive, a hidden variable from the point of view of the behavior, was quite inspirational for the concept of motivation.

The notion of drive was often related to physiological needs, e.g. thirst is a signal to recover a good balance in our body water after potential perturbations such as salt intake, temperature etc.: drives are homeostatic. The reason why this notion is so related to homeostasis could be that regulated physiological variables are easy to measure and to manipulate in animal experiments. However, it was acknowledged that not all drives are necessarily homeostatic. A simple example is the binge eating behavior of peanuts at cocktail parties: there is objectively no homeostatic drive, yet, it is difficult to refrain from it! To improve the concept of motivation, the concept of hedonic reward was introduced: we do things because we get pleasure from them. Two concepts should be distinguished regarding hedonia: the wanting aspect (the appetitive phase) and the liking aspect (the consummatory phase). The wanting aspect emphasizes that a motivated behavior is an operant response. In this view, a motivated behavior should be characterized by its flexibility (we can change the response if needed), it implies the expectation of a goal and it is associated to affects: it feels like something special.

These outcomes of the behavior can be rewarding and pleasurable *per se* (e.g. drinking, in the sense of homeostatic drive), or secondarily (e.g. earning money, because money can be exchanged for goods and services that are pleasurable). Following this distinction between primary and secondary rewards, the latter are also called incentives. Note however that the sharp distinction between incentives and intrinsic rewards sometimes need to be rounded; for instance, money can be seen as a tool to get pleasurable experience, but it can also be seen as an intrinsic reward, potentially leading to pathological behaviors (Lea and Webley, 2006). In this work, we use the notion of motivation in the sense of incentive motivation, the process that translates an expected goal into a behavioral activation, and we will use money as an incentive.

#### 2.3.2.1.2 Potential motivation and actual motivation arousal

A sort of puzzle with the notion of motivation was the following: an experimenter asks a monkey in the lab to press a lever at a given difficulty level, that the monkey knows, to get one drop of juice; (A) the monkey does it when the difficulty is 2 Newtons, (B) she also does it for 4 Newtons, but (C) she

declines for 10 Newtons. The comparison of (A) and (C) fits well with the concept of motivation: one scales behavioral activation in proportion to the reward at stake, but then, how to explain that the behavioral activation is higher in (B) and (A), while the reward is kept constant? The distinction between the potential motivation (the maximal activation we may allot) and the actual motivation arousal (the actual activation allotted) settles this issue (Brehm and Self, 1989). The potential motivation is controlled by the reward size, thus it is the same across A, B, C; however the actual motivation arousal needed to perform these actions is different, because the difficulty is different. When we are free to decide whether to perform the action or not, the actual arousal is deployed only if it does not exceed the potential motivation and in case it is deployed, the minimal level necessary to do the task is used, so that costs are minimized.

This optimization perspective on the behavior should nevertheless be rounded, as there are cases when the behavioral energizing is clearly not relevant, or operant to follow the wording of the conditioning literature. An example is the vigor of the response, e.g. how quickly we respond. In many cases, the response is all the quicker that the incentive is high, even though this increased vigor does not impact at all the reward delivery (nor its timing) (Niv et al., 2005). We will see below how this can be explained as a side effect of the neurobiology of motivation.

### ***2.3.2.2 Incentive motivation of motor behavior in normal condition***

#### **2.3.2.2.1 The pivotal role of the basal ganglia**

In the first section of this introduction, I stressed that the striatum plays a major role in encoding values, prediction errors, learning and optimization of effort production. There are therefore in the striatum, a value-based processing and a motor processing, in particular in the ventral part, see Figure 10 (Liljeholm and O'Doherty, 2012). It is thus a natural candidate to translate the prospect of a benefit into the exertion of a motor cost, and has a major role in the activation system (Kalivas and Nakamura, 1999).

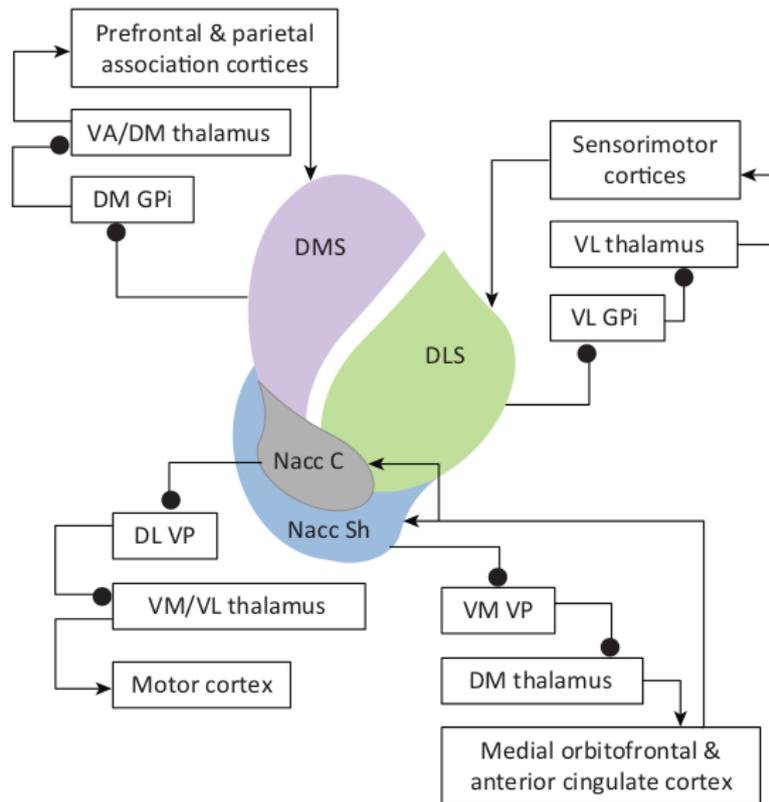


Figure 10 **Connectivity of the striatum**

Four domains can be identified based on the connectivity. Arrows correspond to excitatory connections, circles to inhibitory connections. DMS, dorsomedial striatum; DLS, dorsolateral striatum; GPI, internal segment of globus pallidus; VP, ventral pallidum; VA, ventral anterior; DM, dorsomedial; VL, ventrolateral; VM, ventromedial; Nacc C, nucleus accumbens core; Nacc Sh, nucleus accumbens shell. From (Liljeholm and O'Doherty, 2012)

The incentive motivation process was related to the ventral striatum with functional MRI in a task that motivated the effort production on a hand grip with subliminal incentives: coin images were presented unbeknownst to the subject with a masking procedure (Pessiglione et al., 2007). In this task, the production of higher force levels for higher incentives was operant since the payoff was proportional to the force produced and the incentive level. The results were that higher force levels correlated to the activity in M1, higher incentive levels to the activity in the ventral pallidum (both in the conscious and unconscious perception condition) and participants produced higher force levels for higher incentives, both in the conscious and unconscious condition. In another study, the motivation was also manipulated by monetary incentives and subjects had to produce effort, knowing that they would be rewarded following the same rule as in the previous study. Again, the incentive level was encoded in the ventral striatum, the force level produced was encoded in the motor cortex, but in addition, a significant functional connectivity was found between the two, as if the value signal in the ventral striatum energized the motor cortex to produce higher force levels (Schmidt et al., 2009). In that particular study, the emotional state of the subject was also

manipulated and impacted the effort produced when the emotional stimulus, a picture, was salient (either negatively or positively) as opposed to neutral images. However, this emotional salience was encoded in another part of the brain, the ventro-lateral prefrontal cortex and also showed a significant increase in connectivity with the motor cortex. Overall this study showed that emotional and incentive motivation facilitation of effort production are processed by different routes.

The role of the ventral striatum in energizing behavior was also found in another study in which the ventral striatum was more activated for higher incentives and when subjects better performed in either a physical or mental task. Furthermore, the functional connectivity was switched on and off with the putamen and caudate when the effort was physical or mental respectively, as if the ventral striatum was energizing specifically the cognitive and motor parts of the dorsal striatum depending on the task demand (Schmidt et al., 2012). The combined effect of higher incentive value and higher effort production (when it is operant in the task) was also found in the ventral striatum in another research group (Kurniawan et al., 2010). Similar evidence was found in primate: neurons in the ventral pallidum encoded the expected reward in a task in which saccades toward specific targets were rewarded. Monkeys improved the latency and velocity of their saccades for higher expected rewards; this energizing however was abolished by muscimol inactivation of the ventral pallidum (Tachibana and Hikosaka, 2012). The pivotal role of the ventral striatum to translate the expected value into a motor code that energizes the effort production in other motor cortical and subcortical regions is comforted by the evidence, notably from rodent studies, that the ventral pallidum receives input for other regions, such as the anterior cingulate cortex, to represent the net value of the action to perform (Hauber and Sommer, 2009). The nucleus accumbens in particular has a critical role to accept high effort for high reward; a study using precise targeted reversible inactivation in this structure showed that the shell, not the core, support this function (Ghods-Sharifi and Floresco, 2010).

#### 2.3.2.2.2 Dopamine as a key neuromodulator

The role of striatal dopamine, in particular in the ventral striatum, for learning was stressed earlier in this introduction. A possible misunderstanding of its functional role is that it signals reward, whereas it may actually subserve operant behavior in the general sense. In the case of effort-based decision-making, dopamine may promote the behavioral activation that is often entailed by the prospect of higher reward, but not the reward per se (Salamone et al., 2007). To follow the distinction presented earlier between wanting and liking (appetitive phase vs. consummatory phase), the dopamine would be more associated to the wanting (Dayan, 2012). The role of dopamine in operant behavior would be to boost motivated behavior by overcoming the cost associated to the response (Kurniawan et al.,

2011). Manipulation of dopamine levels with pharmacological tools, such as L-DOPA that enhances dopamine level, has been shown to increase appetitive and approach behaviors specifically, as opposed to when no action is needed (e.g. go vs. no go task) across incentive levels, suggesting that the dopamine is involved in representing the reward associated to actions, but not reward as such (Guitart-Masip et al., 2012). In this study, the behavioral pattern of effort was mirrored by increased activation in the ventral tegmental area and the substantia nigra in fMRI, a dopaminergic nuclei that innervate the striatum. Similarly in rodents, depletion of dopamine in the nucleus accumbens biases rats to choose lower effort and lower payoff (Cousins et al., 1996). Rodent studies however showed that this effect is not specific to the nucleus accumbens, similar results are observed for the anterior cingulate cortex (Schweimer et al., 2005).

Striatal dopamine therefore seems implicated in energizing the behavior when more reward is at stake. This view makes an interesting prediction of the potential side effect of this mechanism: the general increased activation subserved by higher tonic striatal levels could propagate to parameters of motor production that are actually not operant for the task. A critical example of this is the response vigor, i.e. how fast the subject responds. This response vigor is often increased by the reward at stake, even though it does not change the reward size or reward delivery because the schedule is controlled in the task (Niv et al., 2005, 2007). This vigor effect of higher motivation is dependent on dopamine, as revealed by dopaminergic manipulations in healthy humans (Beierholm et al., 2013). Besides pharmacological manipulations, if the tonic dopamine level is increased locally in an operant task, for instance because the reward rate is increased over several consecutive trials, the response vigor should increase. This was found behaviorally in rodents, and critically, this effect was abolished by lesions of the dorso-medial striatum, while the response to trial-to-trial variations of reward was still preserved (Wang et al., 2013).

Dopamine, despite a major role in incentive motivation, may not be the only neuromodulator that affects this process. It was for instance suggested that dopamine is part of an opponent system, the symmetric of which is serotonin (Daw et al., 2002). It was also suggested that among catecholamines, dopamine and noradrenaline could have complementary roles, with noradrenaline contributing to adding extra behavioral activation that is not accounted for by dopamine activation elicited by the reward prospect, so that the extra activation induced by noradrenaline would cope with higher cost demands (Bouret et al., 2012). In that sense, noradrenaline would be crucial to regulate the actual motivation arousal, to follow Brehm's distinction. Histamine levels in the brain could also be an important determinant to the incentive motivation, as reduced histamine levels are associated to decrease motivation, in particular in the appetitive phase (Torrealba et al., 2012).

### *2.3.2.3 Pathological impairment of motor incentive motivation*

The previous section presented the pivotal role of the striatum and the dopamine for incentive motivation, with evidence from normal condition. Some pathological states provide further evidence to support this view.

#### *2.3.2.3.1 Parkinson as a motor motivation deficit*

Parkinson's disease is characterized by a reduction of voluntary movement (hypokinesia) that, in the most extreme form, is close to the absence of voluntary movements (akinesia). This disease is due to abnormally low levels of striatal dopamine, which is the consequence of the degeneration of the substantia nigra pars compacta, a dopaminergic nucleus (DeLong, 2000). The functional interpretation of this pathology can be reframed in a motor motivation deficit: Parkinsonian patients are not able to activate their motor system (Mazzoni et al., 2007). The role of dopamine in these slow movements mirrors the previous discussion on dopamine and the response vigor (Niv et al., 2005). In line with this interpretation in terms of a motor motivation deficit, the motor production of Parkinsonian patients was analyzed following the formal quantitative framework of the optimal motor control theory. The result of this analysis was that Parkinsonian patients have a motor control that is as optimal as it is in normal subjects, the only difference being a lower motor drive. In other words, the motor impairment in this pathology is not a problem of motor control but a problem of motor activation (Baraduc et al., 2013).

#### *2.3.2.3.2 More canonical motivation disorders*

Other pathologies are more canonical motivation deficit from the clinical point of view than Parkinson's disease. A first example is depression. The emotional incentive motivation task mentioned earlier (Schmidt et al., 2009) dissociated two ways of activating the motor system: the emotional arousal and the incentive motivation. Depressed patients were tested in this paradigm, the results showed that they are able to modulate their effort production depending on the emotional arousal, but not depending on the incentive, arguing in favor of a specific deficit in the incentive motivation process, not in the motor command (Cléry-Melin et al., 2011).

Another interesting motivation deficit is the auto-activation deficit disorder. This pathology is a form of apathy in which patients are unable to produce behavioral activation following internal goals: they can spend hours in a quiet room doing nothing. This is not a motor deficit: they can activate their behavior upon external request, e.g. if they are asked to. It seems that the difficulty to follow internal goals is not due to the inability to attribute values to things. In an incentive effort production task, their skin conductance response followed the level of incentive, as it does for control subjects, suggesting that their internal valuation system and their arousal is not completely dysfunctional.

However, as opposed to control subjects, they did not modulate their effort production according to the incentive level (Schmidt et al., 2008). Crucially, these patients are characterized by bilateral lesions of basal ganglia. This finding accords well with the idea that the ventral striatum has a pivotal role to translate an expected reward into a behavioral activation.

### 2.3.3 Summary

- There is something unavoidable when moving a fridge: we have to make efforts.
  - The force is produced by muscles and the central nervous system controls their recruitment levels, hence, the exerted force.
  - The motor command is produced by the motor cortex.
  - Production of higher force levels requires more activation of the motor cortex.
- We can produce efforts... provided that we are motivated to do so.
  - The motor cortex is embedded in a network that monitors motor production and adapts it, not only to unexpected motor perturbations, but also to the goal pursued.
  - Incentive motivation translates the prospect of a benefit into an operant response to obtain this benefit.
  - The basal ganglia, in particular the ventral striatum, are key structures, and dopamine a key neuromodulator to energize behavior in incentive motivation.
- The physiological cost of effort should be limited to avoid damages.
  - Afferent signals inform the central nervous system about cost-related parameters.
  - Proprioceptive afferents signal the muscle mechanical state, like length and tension.
  - Nociceptive afferents signal muscular damages: chemical imbalance, excess of temperature, etc.
  - Nociception and proprioception are separate pathways involving distinct receptors, fibers and projections.
  - The afferent signals serve as a feedback regulation of the efferent command, but the efferent command also modulates afferents: there is cross-talk between signals.
- The temporal effort allocation problem is by essence a problem of dynamic. Effort production entails specific dynamics.
  - There is a hidden dynamic in effort production that is anticipatory: voluntary effort is preceded by preparation, as reflected in the readiness potential, the motor beta de-synchrony and electromyogram recordings.
  - The readiness potential and the motor beta de-synchronization have distinct origins. Furthermore, the former, as opposed to the latter, is impacted by motor parameters.
  - A decreased effort production can be due to 1) a decreased incentive to produce the effort, 2) a decrease efficacy of the process translating incentives into effort or 3) a decreased ability to produce the effort.

## 2.4 How time steps in

In order to understand the effort allocation problem, I introduced above the economic view on cost and benefit, how they are represented in the brain, how they are traded-off in decision-making, how higher benefits may lead us to produce higher physical efforts and how physical efforts are produced. The effort allocation problem addresses the question of why we alternate periods of rest and exertion over time, and how we set the switch between the two. For the ‘why’ question, it is obvious that there is something that changes overtime: at least behaviorally, there are switched between effort and rest. Time therefore seems to be a crucial parameter underlying effort allocation. What hidden determinants of the behavior could explain such a dynamic? Given the theoretical concepts introduced above, a rather simple candidate explanation to start with is that there are changes over time of benefit (section 2.4.1) and cost (section 2.4.2). The last two sections of this introduction (section 2.5 and 2.6) will address the ‘how’ question: how the dynamics of cost and benefit may guide effort allocation.

### 2.4.1 Benefit is not constant over time but discounted by delay

#### 2.4.1.1 *Something special about time*

##### 2.4.1.1.1 Time is different from other costs

The delay within which the outcome might be reached is at play in decisions, not only because it requires prospective thought, but also because delays discount the utility value (Carter et al., 2010). An example might be the following: would you prefer to go on holiday now or six month later, knowing that in the meantime, you can manage to save money so that delaying your trip would improve it? The outcome delay is often taken into account and its effect on decision can be as strong as to reverse our initial preference: if the delay is sufficiently long between the later-but-higher reward and the sooner-but-smaller reward, we may choose the small reward. This is the issue of inter-temporal choice (Peters and Büchel, 2011): how time within which the outcome is reached discounts its value.

This tendency to devaluate the future is pervasive in humans. This bias for present is a form of impulsivity when we discount future benefits (Kalenscher and Pennartz, 2008; Peters and Büchel, 2011) and could also account for procrastination when we discount future costs (Steel, 2007). This bias varies between individual depending on differences in reward valuation processes, cognitive control abilities and abilities in prospective thoughts (Peters and Büchel, 2011). These differences reflect a graded ability to wait for longer delays in order to get larger rewards. Delay discounting and the ability to refrain from always choosing immediate rewards is shared with other primates, other

mammals such as rodents (Cardinal et al., 2001, 2002; Rudebeck et al., 2006) and non-mammalians, such as birds (Hayden and Platt, 2007). To stress the pervasiveness of delay discounting, it should be mentioned that translation from economic-like choices to other domains like motor control (Wu et al., 2009) revealed that motor behavior is also underpinned by delay discounting of reward (Shadmehr et al., 2010).

Data from fMRI provided evidence that delay discounting of reward involved specific brain areas (McClure et al., 2004; Pine et al., 2009). This specificity leads to a functional distinction between delay discounting and effort discounting (Walton et al., 2006; Talmi and Pine, 2012). A few experimental pieces of evidence are listed here to support this distinction. A study contrasted effort and delay discounting under fMRI and found different anatomico-functional correlates: ventromedial prefrontal cortex and ventral striatum for delays; anterior cingulate cortex and anterior insula for efforts (Prévost et al., 2010). Lesion studies in rodents also support a dissociation between the two: lesions of the orbito-frontal cortex made rats more impulsive whereas the effort side of the decision was affected by lesions of the anterior cingulate cortex (Rudebeck et al., 2006). There is extensive neuropsychological evidence from the human clinics supporting a role for the prefrontal cortex in delay discounting but surprisingly limited evidence in rodent (Kalenscher and Pennartz, 2008). The issue of the homology between rodents and humans prefrontal cortex could resolve this discrepancy. Two other rodent studies may be representative of this asymmetry between humans and rodents. Lesions of the nucleus accumbens, but not the anterior cingulate cortex and medial prefrontal cortex, was associated to increased impulsivity in an influential work (Cardinal et al., 2001) and another study also reported a role for the nucleus accumbens, but with a differential role of dopamine and serotonin: dopamine in this nucleus impacted both effort and delay discounting, whereas serotonin impacted only delay discounting (Denk et al., 2005). Despite some discrepancies between studies of the anatomico-functional correlates of effort and delay discounting, there are several reports of within-study dissociation between effort and delay processing.

#### 2.4.1.1.2 A critical methodological issue

The dissociation between effort and delay processing is often impaired in experimental studies when higher difficulty levels co-vary with longer exertion levels. For instance when the cost is manipulated through the amount of lever presses (Bouret et al., 2012), or a number of target to cancel on a screen (Croxson et al., 2009), more difficult efforts are also longer efforts. In this case, it is impossible to unravel the impact of the effort difficulty *per se* from that of the duration of the exertion that delays the reward delivery. Whether the cost variations are purely variations of the delay (with the same effort level) or variations of the effort (with the same delay) is often overlooked. This distinction may

not be crucial when all that matters is that cost is manipulated; however it turns out that when one wants to disentangle the effect of delay and effort, a proper experimental setting is needed for the distinction.

#### 2.4.1.1.3 Why delays violate the standard utility theory

It is not a problem a priori for the standard utility theory that delays affect the utility. In fact, like the expected utility theory was designed to serve as a normative framework of choice under uncertainty (i.e. when outcome delivery is stochastic and not deterministic), the discounted utility theory was designed to serve as a normative framework of choice under delay. Typical applications for this theory are pension funds, saving plans, loans etc. This theory quantifies mathematically the delay discounting. In the simplest form, this discount is linear: a subtraction by a quantity proportional to the delay; however, non-linear temporal discounting, e.g. with exponential or hyperbolic functions better fit the data (Peters and Büchel, 2011). In these two latter models, the discounting is captured by a single variable: the temporal decay parameter.

The discounted utility theory, like the expected utility theory, is designed to fulfill some consistency constraints. One of these constraints is that a choice between  $R_1 \& D_1$  vs.  $R_2 \& D_2$ , two options with a reward level ( $R_1$  or  $R_2$ ) and a delay duration ( $D_1$  or  $D_2$ ), should not be affected by shifting both delays, i.e. by introducing a common lag  $L$ :  $R_1 \& (D_1+L)$  vs.  $R_2 \& (D_2+L)$ . Exponential discounting functions are used in the discounted utility theory because they ensure that the relative order of discounted values is preserved whatever the lag introduced. This constraint however is relaxed when hyperbolic discount functions are used. And unfortunately for the consistency requirement of the theory, hyperbolic functions often provide better fits of the behavioral data than exponential discount functions (Kalenscher and Pennartz, 2008; Peters and Büchel, 2011).

How to account for violations of the standard utility theory? Results from neuroscience could provide potential explanations (Kalenscher and Pennartz, 2008). For instance, evidence from neuroscience, in particular with fMRI results, suggests that there is not a single brain system to discount value when there is a delay, but several systems. In particular, a distinction was found between the limbic system that attributes a penalty to options that are not immediate, whereas the prefrontal cortex attributes gradual discount to future rewards (McClure et al., 2004). Conflict between these systems, or different involvement of these systems depending on other parameters of the choice could account for some inconsistencies between choices with respect to the discount utility theory (Kalenscher and Pennartz, 2008). Another and actually related idea is that the brain uses routines to solve the inter-temporal choice decisions. These routines would have been optimized by evolutionary selection to be nearly optimal in ordinary and frequent situations, which is not necessarily the case of

experiments in behavioral economics, or even more simply, in our modern lives (Kalenscher and Pennartz, 2008).

#### *2.4.1.2 Time, the opportunity cost and lessons from the foraging theories*

This section follows the above intuition that the brain might use routines shaped by the evolution and that more evolutionary rooted or at least more 'natural' ways of addressing the inter-temporal choice could shed interesting light on this process. Foraging may be such a more 'natural' way and potentially of interest for the matter (Stephens, 2008). Why would the foraging problem be similar to inter-temporal choices and how the understanding of the former could enlighten the latter? An experimental example will serve to address this point. In a recent publication (Wikenheiser et al., 2013), rats underwent an inter-temporal choice task framed as a stay or leave decision at spots that led to reward after a delay of variable duration: the rat moved in a maze, entered spots and then could either wait to get a reward (the reward is not delivered until the rat has waited the required duration) or leave and go to another spot. What is optimal to do in such a task can be assessed according to a metric that is the rate of food per unit time (e.g. how many food pellets the rat receives divided by the duration she had to wait to get it). The idea is that the rat should maximize the amount of food and minimize the time spent. The notion of opportunity cost is closely related to the rate of food intake in this task. From the microeconomic point of view (Cahuc and Zylberberg, 2004), the opportunity cost is the value of the alternative forgone, i.e. the forgone utility with respect to the best thing you could have done. In the task by Wikenheiser and colleague, rats were introduced in different environments, in which the mean delay across options was varied from an environment to another. In that sense, the richness of the environment varied: it could be more or less easy (i.e. quick) to get food. Therefore, the foregone cost to wait a unit time, instead of trying to forage at another spot to get food more quickly, increased when the richness of the environment increased: the forgone cost of waiting was higher when the environment led to reward quickly compared to when it took longer. In other words, the opportunity cost of waiting was higher when the environment richness was higher. The rate of food intake and the opportunity cost serve as a reference to benchmark behavioral policies. In the study mentioned rats foraged from spot to spot but tended to accept abnormally long delay, i.e. options for which the delay should be declined and the rat should have left to another spot.

Optimization of the opportunity cost has actually already been presented in this introduction. It was mentioned that motivated behavior tended to optimize this opportunity cost, in particular by shortening delays when it is possible (see section p. 60). This aspect of motivated behavior is under the control of striatal dopamine, as are motivated behaviors as a whole, which could be the reason

why mammals shorten delays when more motivated, even if it is actually useless (Niv et al., 2007). In the experiment by Wikenheiser and colleagues, the results are at odd with this opportunity cost optimization since the rats seemed reluctant to leave a spot at which they had already waited a long time. They tended to persevere in waiting, although it was sub-optimal in this case, as if there were a 'cost to leave', that the author estimated quantitatively. This cost to leave is similar to the 'sunk cost' effect, the aggravation of cost ensued by a reluctance to abandon actions (although it would be better) for which a significant cost has already been expended. A possible explanation is that the 'cost to leave' is part of a routine behavior that is on average optimal in usual foraging conditions. Why should the amount of travels from spot to spot be limited?

To illustrate this point, I introduce a classical foraging problem: the predator in a patchy environment (Charnov, 1976). In this context, the preys are available in patches of the environment. When the predator gets into a new patch and starts feeding, the rate of catch progressively declines as there are gradually fewer preys in the patch (because of the predation-induced mortality, or because preys migrate to another patch to flee the predator, etc.). To face this decline of the patch productivity, the predator has two options: either stay in the low and declining yield patch or migrate to another patch. But crucially: the travel is costly since the predator starves in the meantime she reaches the new patch. Optimal foraging policies can be defined based on the opportunity cost. Such a foraging problem is a 'natural' form of inter-temporal choice: immediate but little food (exploitation of the patch) or delayed but plenty of food (travel to a new patch). The optimal policy is to stay certain duration per patch; the precise duration depends on the characteristics of the environment but is always a tradeoff between not staying too much and not leaving too early. The case of preys and predators is a canonical example, but it can be generalized to many foraging problems, such as rodent foraging.

The optimality of foraging policies can be investigated in the field and in experimental settings. For example, starlings were used for an experiment in which the environment was calibrated so that the opportunity cost was controlled (Bautista et al., 2001). The starlings could walk (low expense, low rate of food intake) or fly (high expense, high food intake). The walk or fly decision, based on the amount and distribution of food, revealed that starlings modulate their decisions according to opportunity cost (and other factors, such as the predation risk, etc.).

The routine to achieve good foraging policies may be hard wired in our brain. The understanding of such 'natural' inter-temporal choice could be an evolutionary account of why we penalize delay (discount of the reward magnitude) but still are able to avoid too impulsive behaviors and preserve an ability to wait (Stephens, 2008).

## 2.4.2 Cost is not constant over time but increased by fatigue

In the preceding section I showed that time could alter benefit. It is also the case for physical cost, although the reasons and mechanisms are completely different. Physical cost is a highly dynamical variable that grows up along with the exertion. This dynamic is due to fatigue. In the first section (2.4.2.1) I define the notion of fatigue; then I review the mechanisms by which fatigue increases during exertion because of alterations at every level of the motor command (section 2.4.2.2). In the last section (2.4.2.3) I discuss the relation between fatigue, pain, and more generally sensations.

### 2.4.2.1 What is fatigue?

There is no need to be an experienced physiologist to know that muscular exercise cannot be sustained indefinitely and that the maximal amount of force produced decreases over time (see Figure 11).

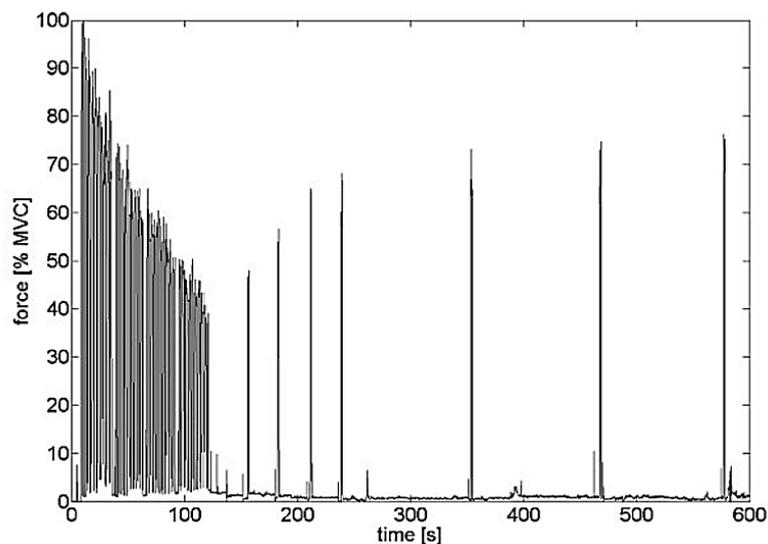


Figure 11 Reduction of maximal voluntary contraction (MVC) over time.

The exerted force level is normalized with respect to the maximal value. In the first phase (from 0 to 120s), there are repeated and frequent efforts, then more rest is allowed. From (Scheidtger et al., 2010).

There are several definitions of fatigue, two of which are particularly interesting for this work. A definition could be ‘any exercise-induced reduction in the ability to exert muscle force or power regardless of whether or not the task can be sustained’ (Gandevia, 2001), a view shared by other authors, see (Tanaka and Watanabe, 2012). Another definition has been proposed and complement the first one with a sensation aspect: ‘an acute impairment of performance that includes both an increase in the perceived effort necessary to exert a desired force and an eventual inability to produce this force’ (Enoka and Stuart, 1992). This latter review was updated 15 years later, and the authors noticed that the research in the field has been much more concerned with the behavioral and physiological aspect of fatigue than with the subjective aspect (Barry and Enoka, 2007).

One of the hot topics addressed in the field is the origin of fatigue, in particular the parts of it that is accounted for by peripheral cause and central cause (spinal and supraspinal). This issue can be dated back a century ago when the Italian physiologist Angelo Mosso and the British physiologist (and Nobel Prize) Archibald Vivian Hill argued about whether physical fatigue was more central or muscular (Gandevia, 2001). It is now agreed that they are both causes of physical fatigue and that supraspinal fatigue can account for up to 25% of the physical fatigue during exertion (Taylor et al., 2006). However, such a clear cut distinction between peripheral and central may not be fully relevant since central fatigue partly emerges from feedback signals arising from within muscles (Barry and Enoka, 2007).

#### *2.4.2.2 Where does fatigue come from?*

The relationship between fatigue and effort production is tight because fatigue is a consequence of regulatory mechanisms that are intrinsic to the motor command at different levels and to the homeostatic controls of the body and brain variables that are perturbed by strenuous exercise. Fatigue thus reflects direct and indirect feedback controls so that exertion cannot go without fatigue.

##### *2.4.2.2.1 Regulation of the motor command*

###### *2.4.2.2.1.1 How peripheral and supraspinal cause can be disentangled*

The reduction of force production shown in Figure 11 has muscular and central causes, how to distinguish these causes experimentally? A set of stimulation techniques makes it possible to unravel the origin of fatigue at different levels and their relative weights. The trick is that the motor command can be shunt by direct electric (or magnetic) stimulations at different levels: the motor cortex, the pyramidal decussation, the spinal cord and the efferent nerve between spinal cord and the muscle (McNeil et al., 2013). An example will serve as an explanation (see Figure 12). An athlete was tested before and after running a marathon (42 km). A motor response was evoked on the electromyogram by ankle flexion after transcranial magnetic stimulation of the motor cortex and transcutaneous stimulation of the efferent nerve. After the marathon, the evoked response was reduced compared to the pre-marathon testing, which is, by definition, a measure of fatigue. Note that a reduction observed after the stimulation can be accounted for by fatigue at all the downstream stages of the motor command: just the efferent and muscle response in the case of the transcutaneous stimulation, and all the stages in the case of the transcranial motor cortex stimulation. By contrasting the two results, the spinal and supraspinal fatigue can be told apart from the peripheral fatigue.

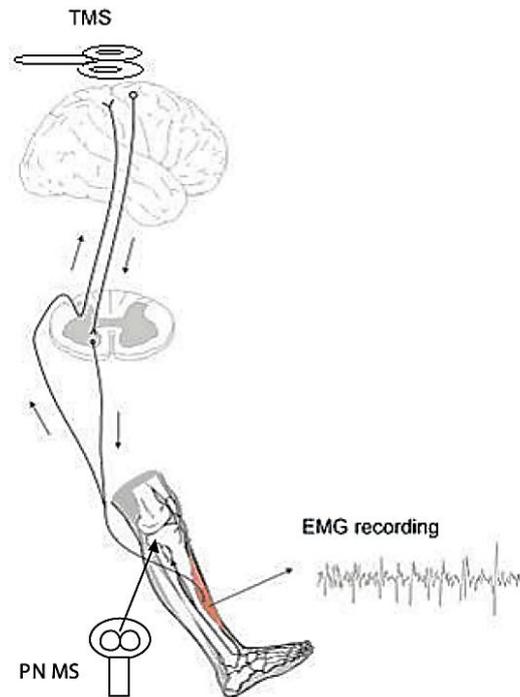


Figure 12 **Central and peripheral origin of fatigue.**

The force produced by ankle flexion and the EMG of the according muscle is estimated after stimulation of the motor cortex and the efferent nerve. From (Ross et al., 2007).

This method was developed in the 1950s and is called the twitched interpolation technique, since what is informative is the additional force produced (the twitch) between two conditions, like between stimulations at two different stages of the motor command, or between a stimulation and the voluntary command. This additional effort fraction is produced by shunting the command at a given level and hence, shunting the potential fatigue from upper levels (Merton, 1954). This technic provides evidence of spinal and supraspinal origin of fatigue, in particular in the motor cortex (Amann and Dempsey, 2008).

Note that a simpler way of providing evidence for supraspinal source of fatigue is the case when fatigue is shared between a muscle that was exercising, and a muscle that was kept idle during the exercise. This can be tested for instance on distinct limbs (Rasmussen et al., 2010).

The two following sections review some causes of muscular and central fatigue.

#### 2.4.2.2.1.2 *Loss of capacity in the muscle*

The causes of muscle fatigue were investigated by muscle physiologists: many candidate causes were found. For the sake of brevity, they are listed below, after (Boyas and Guével, 2011; Tanaka and Watanabe, 2012):

- Loss of the muscular junction excitation capability

- Lack of intracellular metabolic substrate (glycogen depletion, phosphocreatine, ATP)
- Impairment of the muscle contractile apparatus (e.g. when calcium ions are not released sufficiently)
- Increase of damaging metabolic by-products, such as the reactive oxygen species.

An example of substrate depletion is illustrated in Figure 13. The data show a clear depletion of the creatine content during exertion. Creatine is a support of energy in the muscle.

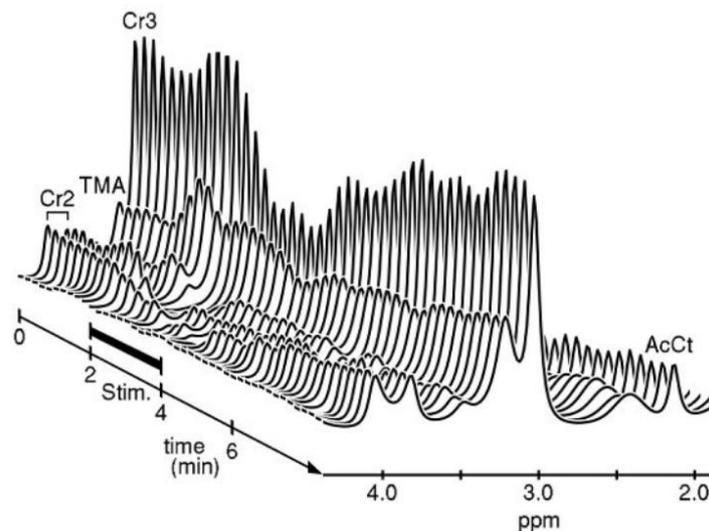


Figure 13 *Repeated mass-spectroscopy data on an exercising muscle.*

The x-axis correspond to compound mass, the compound name is provided for each peak (total creatine: Cr2 and Cr3, AcCt: acetylcarnitine, TMA: trimethylamine). The vertical axis (not plotted) corresponds to the compound quantity. On the time axis, the black line indicates the duration of muscle exertion. From (Scheidegger et al., 2010).

The energy supply to the muscle can also be limited during exercise, which is also a cause of muscle fatigue.

#### 2.4.2.2.1.3 Feedback inhibition of the central command

There is direct evidence with electrophysiological correlates of fatigue, for instance on the readiness potential, that the motor command is altered in the brain during muscular fatigue (Berchicci et al., 2013). Where do these changes come from? As discussed in a previous section (see p. 43), the muscular drive results from a balance between the feedforward motor command and feedback inhibition. The motor command can therefore be viewed as under the control of two opponent systems that regulate fatigue (Boyas and Guével, 2011; Tanaka and Watanabe, 2012), see Figure 14.

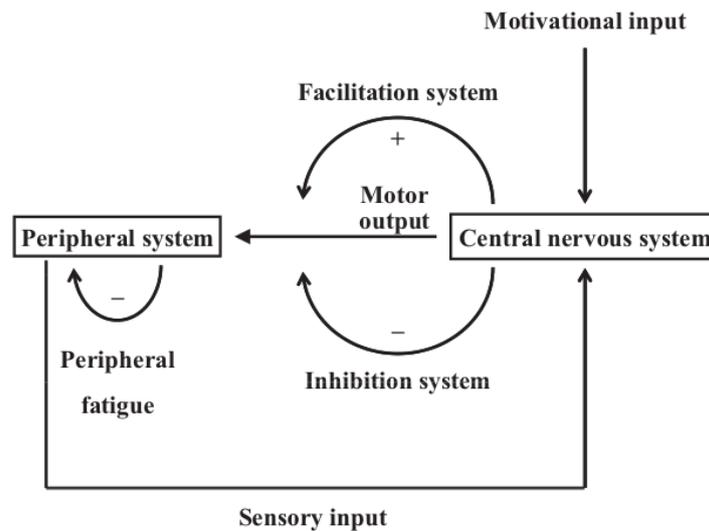


Figure 14 *Opponent control on the motor output.*

From (Tanaka and Watanabe, 2012)

The reduction of the drive results from the muscle afferent signals (fibers Ia, Ib, II, III, IV) that inhibit the drive to ensure muscle protection. This feedback inhibition on the drive was referred to as the 'muscle wisdom' (Enoka and Stuart, 1992; Gandevia, 2001; St Clair Gibson et al., 2001).

Among the afferent fibers, the type III and IV that convey alarm signals (see p. 45) drew a particular attention. The role of the nociceptive fibers on muscle fatigue was evidenced by manipulating the pain level in the muscle with injection of intramuscular saline solution in healthy participants. The response evoked on the electromyograms under pain was increased by cortico-spinal stimulation and reduced by transcranial stimulation, suggesting that increased activity of type III & IV fibers facilitates motoneurons and depresses the cortical excitability (Martin et al., 2008). There are several other reports of similar results revealing a dis-facilitation of the motor command by these fibers (for a review: (Taylor and Gandevia, 2008)). A more recent study used intrathecal injection of  $\mu$ -opioid agonist (fentanyl) in healthy participants with the twitched interpolation technic. Under opioids, the EMG response was increased and the silent cortical period that follows transcranial stimulation was reduced, which again points toward a role of the nociceptive feedback in fatigue at supraspinal levels (Hilty et al., 2011b). Note that during physical exertion, the level of beta-endorphin release (an endogenous opioid) is increased in the brain, which could relieve effort-induced pain and fatigue (Gandevia, 2001).

The role of serotonin is often mentioned as a key neuromodulator of supraspinal fatigue, in particular from animal studies, though there is little evidence in humans (Gandevia, 2001; Millet and Lepers, 2004; Nybo and Secher, 2004). For instance, muscle fatigue was investigated in turtles, and

stimulation of the raphe nucleus significantly induced more fatigue; these authors concluded that serotonin is involved in supraspinal regulation of fatigue (Cotel et al., 2013).

#### 2.4.2.2.2 Homeostatic perturbations

Besides changes in the muscle physiology and the motor command which characterize muscular fatigue, there are also other perturbations in the body that are caused indirectly by physical exertion and that may also participate in the fatigue process itself.

##### 2.4.2.2.2.1 *At the body scale*

There are many homeostatic perturbations at the scale of the whole body during strenuous physical exercise. Several of them may participate in the down regulation of the effort production, i.e. in fatigue. The cardiovascular and aerobic capability could impose a higher bound on exertion (Abbiss and Laursen, 2005), even though whether this factor is truly the limitation of effort production is debated; for a controversy, see (Noakes, 2000) against (Shephard, 2009). Other potential limiting factors include thermoregulation. The effort production is a major cause for the increase of the body temperature, which is a variable that is highly regulated in the body (Abbiss and Laursen, 2005). Energy depletion ensued by a shortage in the available substrate is also a cause of body-scale homeostatic limitation on effort production (Noakes, 2000).

##### 2.4.2.2.2.2 *In the brain*

The energy supply of the brain is a critical issue on top of that of the body in general. Neurons have a specific metabolism based on glucose and lactate (Peters et al., 2004; Barros, 2013). During exercise the part of lactate in this metabolism increases (Ide et al., 1999), and the aerobic balance is shifted: the glucose and lactate intake overrides the oxygen intake during physical effort (Dalsgaard et al., 2004; Dalsgaard, 2005). This non-oxidative glucose consumption is not specific to physical effort, it is also observed during cognitive tasks (Fox et al., 1988).

The fuelling of cerebral blood is also affected by exertion (Rasmussen et al., 2006). This global reduction of blood flow could be a response against hyperthermia (Nybo et al., 2002).

The accumulation of serotonin levels in the brain during exercise was mentioned above (Millet and Lepers, 2004; Cotel et al., 2013). This effect could result from an imbalance in the amino-acid content of the blood ensued by muscular exertion (Gandevia, 2001). The muscle exertion could also have a feedback on brain processing mediated by the blood, i.e. through a humoral response. Indeed the levels of several cytokines, in particular the interleukins, are elevated in the brain and this immunity stress response is likely to be triggered by the muscle activity (Gandevia, 2001; Nybo and Secher, 2004)

### 2.4.2.3 *Is fatigue related to perception*

I stressed above that the muscle afferent fibers play a major role in the facilitation / dis-facilitation of the motor command. These afferent fibers also support our proprioceptive and nociceptive abilities (see the dedicated section of the introduction p. 45). Thereby it seems that there might be a relation between fatigue and perception. This link is discussed in this section.

#### 2.4.2.3.1 *Is fatigue related to pain?*

It was mentioned previously that beta-endorphins are released in the brain during physical exertion and that intrathecal injection of  $\mu$ -opioid agonist alleviates fatigue. The opioid system has a major role in regulating our perception of pain (Petrovic et al., 2002, 2008; Benedetti, 2008; Staahl et al., 2009). Blockers of opioid receptors can mitigate our analgic abilities (Taylor et al., 2013).

To what extent the perception of fatigue is distinct from pain is certainly an open question. Interestingly, the brain mechanism at play when pain is regulated, either externally with opioid injection or endogenously like in the placebo effect involved a system comprising in particular the insula (Petrovic et al., 2002; Wager et al., 2004, 2013). This region receives input from the lamina I spinothalamicocortical (Craig, 2002), which is not the same as, but neither independent from, the nociceptive pathway presented earlier (see p. 46). This pathway is believed to support our interoception ability, which is the sense of our body state (Craig, 2002, 2003). Taken together, it seems that fatigue is related to proprioception, nociception, interoception altogether.

#### 2.4.2.3.2 *Can fatigue be introspected?*

That fatigue should co-vary with proprioception, nociception and interoception in normal conditions still leaves open the issue of whether and how fatigue may be introspected. As mentioned earlier in the introduction, some authors do not reduce fatigue to a motor impairment, but include the sensation entailed by this impairment (Enoka and Stuart, 1992; Barry and Enoka, 2007). In line with this idea, several scales of perceived exertion have been proposed (Borg, 1982). The so-called Borg-scale of perceived exertion is now widely used. Borg constructed a categorical scale in which the levels are not just relative from an individual to another, but have an intrinsic meaning. To achieve the requirement of an absolute scale, Borg deformed his scale so that categories could match the linear increase of heart beat rate and the oxygen consumption during exertion. This scale is rather robust and can be used for diagnostic of heart disease (Borg, 1990). The deformation of the scale produces on purpose a violation of the true growth of the perceived intensity of the effort exertion in order to map the sensation onto physiological variables. As such, it is interesting that this perception actually matched a physiological set of variables, even though the relation is not linear.

Yet, the Borg scale may not answer directly the issue of fatigue since it corresponds to an exertion level. Both are related, but not equivalent and can be dissociated. Fatigue ratings were asked to patients with multiple sclerosis, a disease characterized by chronic sensation of fatigue. These patients showed uncorrelated fatigue ratings and actual behavioral performance (DeLuca et al., 2008).

There is another view on the sense on effort. Instead of considering it as emerging from feedback signals, it was proposed that it could emerge from the forward signal. In other words, the sense of effort would be commensurate to the drive sent to produce the effort (Marcora, 2009; Smirmaul, 2012). An experimental argument is that curare-induced weakness of muscle does not affect afferent signals but increases the perceived exertion (Marcora, 2009).

### 2.4.3 Summary

- When we allocate effort over time, like when moving a fridge, we may want to minimize durations because time, in general, is costly.
  - Time to get a reward is treated as a cost in decision: delay discounts rewards. This discount may not be consistent over time, in particular when it is hyperbolic.
  - There are several systems in the brain that contribute to discount reward.
  - The delay-discounting processes may be phylogenetically ancient and related to problems such as foraging.
  - Time is also a cost when it is associated to an opportunity cost (i.e. when doing something else than waiting would pay more).
- An additional reason why time is costly in the context of effort is that exertion (hence its duration) is costly.
  - Effort duration and effort difficulty are in principle dissociable kinds of cost, although they are often confounded.
  - The effort cost is related to time through its dynamic: physical fatigue can be taken as a proxy for physical cost induced by exertion.
  - Physical fatigue is pervasive: it increases rapidly from any exertion.
- Fatigue is intrinsically related to exertion.
  - Fatigue is caused by regulatory mechanisms in the central nervous system, both in the spinal cord and in the brain.
  - Fatigue is also induced by homeostatic regulation, including the monitoring of the cardiovascular and aerobic system, body temperature and energy stores. It can also be caused by substrate depletion.
  - Proprioceptive and nociceptive afferents contribute to fatigue by down-regulating the motor drive at the spinal and supraspinal levels.
  - Opioids and serotonin (among others neuromodulators) regulate fatigue at central levels.
- Fatigue, besides affecting our performance, is something that we feel.
  - The perceived exertion could co-vary with proprioception, nociception, interoception and variables such as the cardiac activity.
  - An alternative view is that perceived exertion co-varies with the motor drive.
  - Perceived fatigue and actual fatigue may be dissociable.

## 2.5 Existing models

When moving a fridge up through the stairs, it is crucial to allocate the effort efficiently over time. This allocation may involve breaks to recover from fatigue, so that part of the effort allocation problem is to set the durations of effort and rest. It would be useful to look at other problems and models that address similar issues, though in different contexts, like the duration of rest, the duration of work or the ratio between the two. This new section reviews some of these models. In order to enlighten the effort allocation problem, these models are presented following distinctions that are relevant to effort allocation: first, models and theories that bring something to the understanding of when to stop the effort; second, those useful to account for when to start it; and last those that suggest how the proportion of effort and rest should be balanced. This outline is tailored for the effort allocation problem so that the scope of some of the theories presented actually overlaps these subsections.

### 2.5.1 When to stop exertion?

#### 2.5.1.1 *The optimal motor control theory and the movement duration*

One, out of many, interesting behavioral findings about our motor control is that the speed of our movement and their accuracy are inversely related. If we have to point with our finger to a target location, the velocity profile has a bell shape, with maximal speed around the middle of the movement and gradual changes when leaving from the start point and when arriving at the target. The more accurate we need to be and the slower is our movement. This relationship is the Fitt's law (Todorov and Jordan, 2002), named after the American psychologist Paul Fitt who worked on the human movement in the 1950s. This relationship is accounted for by the properties of our motor system that constraints our movement: if one wants to minimize the variance when pointing to a target, she must limit the speed of movement. The brain has its own internal model of the world and the motor state of the body (Wolpert and Ghahramani, 2000), and many ways of implementing the same movement (Todorov, 2000). These alternatives are not equivalent in all respects (duration, variance, etc.) so that the actual movement (or the actual class of movements) produced is singled out by the brain based on optimality principles. For instance in a pointing task, if the cost function (i.e. what needs to be minimized) is the distance to the target, this imposes a constraint on the maximum speed allotted to the movement.

The optimal motor control theory is therefore characterized by a view that the brain computes optimal solutions under constraints. To this purpose, this research field uses mathematical tools to provide a normative account of the motor control, i.e. what we should do (Todorov, 2006). A general idea might be useful is that to characterize optimality, one should define the constraints: what the

system can do, which relies on the mechanical properties of the motor system and what the system tries to do at best, which corresponds quantitatively to a cost function to minimize. A second idea is that the solution to this problem is not planned beforehand and then executed, but rather optimized on the fly based on the proprioceptive integration (Todorov and Jordan, 2002; Todorov, 2004). Key contributions from the motor control literature were dedicated to the investigation of how the motor system compensates for perturbations in the course of the action.

What follows from above is that if certain accuracy of movement is aimed, then there is a minimal bound to the movement duration. But what if actually neither the accuracy nor the duration is constrained as such but rather that the aim is to find the best balance between duration and accuracy? The cost function for the accuracy is, intuitively, the reduction of the error on the target location. But it is less clear what the cost function for the duration should be. It was proposed (Tanaka et al., 2006) that the duration should also be minimized. The rationale in this study involved general principles such as increased adaptive value, survival rate etc. to motivate the reduction of movement duration. In this case, when both duration and accuracy were let free, the Fitt's law also emerged, supporting the generality of this principle.

Cost functions in the motor control literature are often framed only in terms of pure cost, i.e. things that are minimized. In particular, they rarely include benefit in the constraints, i.e. things that are maximized. Following the more traditional utility perspective with both cost and benefit, it was recently proposed that motor control may also implement behavior motivated by the maximization of reward at the minimum expense level (Rigoux and Guigon, 2012). In addition, this latter model also included the temporal discounting of reward in the utility function. This study was in line with the attempt to investigate to what extent economic-like decision-making behaviors are similar to motor behaviors (Wu et al., 2009). Supporting that the two processes have much in common, Rigoux and Guigon proposed, and tested, an overarching optimal model to account for both decision (choices) and actions (motor implementation), revealing that both are underpinned by the same equation and parameters.

#### ***2.5.1.2 Foraging theories***

The theory of optimal foraging was presented earlier, in particular with the seminal theoretical work by the American evolutionary ecologist Eric Charnov. He suggested that for an animal foraging in a patchy environment, there is optimal duration the animal should stay in the patch. The way this optimal duration is estimated is based on the marginal value, which is the amount of food that an additional unit time in the patch provides. For many different reasons, the quantity of food available in the patch decreases over time during the exploitation, as shown on Figure 15. A more intuitive

way of representing the marginal value is to plot the derivatives of the food intake, which is high when the animal starts exploiting the patch and then gradually decreases.

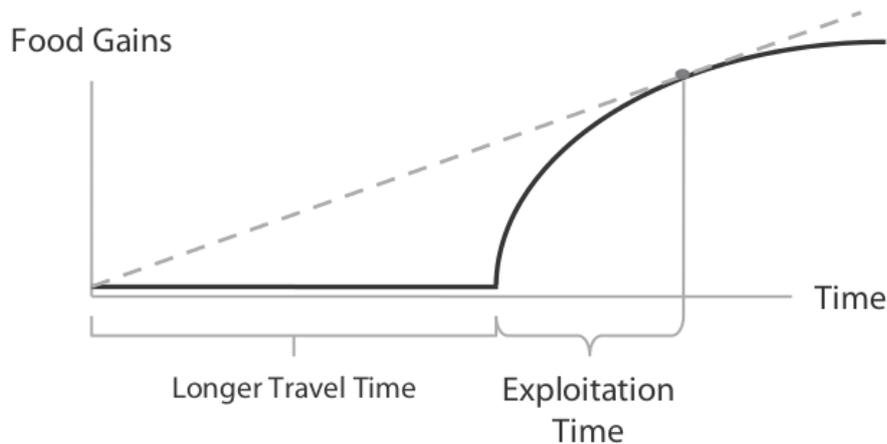


Figure 15 *Optimal foraging and the marginal value theorem.*

*The animal starved when travelling between the patches so that she cannot accumulate food in the meantime. When she arrives at the patch and start foraging, she accumulates food. Note that the curvature of the food intake increase is downward because the patch is progressively depleted from food. From (Stephens, 2008).*

Charnov's theoretical work suggested that the optimal time to stop exploiting the patch is when the marginal utility starts decreasing (Charnov, 1976). The reader can refer to the Figure 15 to get a visual understanding. The slope formed by connecting the origin and a point of the black curve (the quantity of food accumulated at a given time) corresponds strictly to the mean food intake rate. The highest slope for such a curve is achieved by the tangent plotted in dashed line. If the animal chooses to keep exploiting the patch longer than what is presented on the figure, the mean food intake rate is lower: she would have done worse than on the figure, the reason is because the patch is getting too much depleted. On the contrary, if she stops exploiting it earlier, the mean food intake rate is also lower, the reason is that with respect to the travel time, she did not benefit enough from the patch. Another way of restating the same facts is that, the marginal value is lower before and after the limit of the exploitation time sketched on the figure: before this point, an additional time unit of exploitation yields the animal more food than what she has done so far on average, and after this point, an additional unit time of exploitation yields less than what she has done so far on average: there is an optimal marginal value that guides when the predator should leave the patch.

The optimal foraging theory is not restricted to the problem of how time should be allocated. This problem also comprises the choice of the food type (the optimal diet), the choice of the food patches (the optimal patch), the allocation of time between patches, the patterns and speeds of movement (the intensity of foraging) and the strategic choice of the central place from which to forage (Pyke et al., 1977; Pyke, 1984). For the effort allocation problem we are interested in, the allocation of time,

and the intensity of effort are of particular interest. How to find the optimal policy? Charnov suggested that what should be optimized is the marginal utility, this is actually a subcase of a more general approach that can be summarized in three steps: (1) the choice of a currency to assess the efficiency of the foraging policy, for instance the energy spent and acquired (2) the choice of the appropriate cost-benefit function and (3) to find the optimum of this function (Pyke et al., 1977). Although the method is simple, its reliability critically depends on the choice of the currency and the cost-benefit function. How to choose these variable and function is not the matter of computer science, but more the issue of biology. Pyke summarized this, in the case of the foraging problem, as follows: 'there is no recipe for determining just what the currency and constraints should be in a particular situation, and it will always be the job of the naturalist to understand the biology of an animal sufficiently well to know which currency is being optimized' (Pyke et al., 1977). What these variable and function should be in the case of the effort allocation problem, like the fridge problem presented earlier, motivated this introduction.

There are simple situations. Let me introduce another problem from behavioral ecology that clarifies the matter: the diving problem. Aquatic mammal predators, for instance sperm whales, dive to hunt. This activity is strenuous and requires a very efficient metabolism to supply energy, which is achieved mainly through aerobic metabolism. Large oxygen storage can only be acquired at the water surface, so that there is a limitation to the time that the animal can spend diving and hunting: the diving yields food intake (the benefit) but is costly because it depletes the oxygen budget (the cost): the tradeoff between the two defines an optimal diving duration (Mori, 1999).

### *2.5.1.3 Sport literature and the athletic performance*

The sport literature has a particular interest in improving the performance of athletes. In many sports, like cycling, most of the tasks are about endurance so that part of the research is dedicated to a better understanding of the reasons why we stop the effort. This research has a strong physiological flavor and the problem is framed as task failure: why is it that, at some point, the task cannot be sustained anymore? Following the task failure perspective, some 'catastrophic models' were proposed. In these models, a common assumption is that there is a hard constraint on the effort production and a breaking point (Noakes, 2000; Abbiss and Laursen, 2005). This limit may be due to the cardiovascular system: the heart and lung cannot supply enough oxygen for the aerobic demand, so that to protect the body from anaerobiosis, the effort is cut down. In this model, the heart is the limiting factor and the cause of fatigue. Another limiting factor may be a shortage in muscle energy supply. Another limitation may be the declining recruitment of skeletal muscle fibers ensued by central and peripheral fatigue: the effort cessation occurs when the reduction of force production

does not meet the task requirement anymore. Other models involved the regulation of biomechanical parameters, in particular the thermoregulation: in order to limit the body temperature, of which muscle exercise is major source of increase, the effort must be stopped when all the other thermoregulatory mechanisms saturate. Noakes 2000 and Abbiss & Laursen 2005 provide a comprehensive overview of all these models.

All of these models have been criticized because they posit a limit to trigger the effort cessation. Instead, the doubt is cast by several physiological experiments that show that these limits are actually not reached, and therefore that they cannot be the cause of the task failure. For instance in the cardiorespiratory limitation model, the protected variable is the anaerobic metabolism of the skeletal muscles that should not be reached so that the effort is stopped when the cardiovascular production cannot meet this demand. This implies that under hypoxia, the demand of the muscle is met by a higher cardiac activity, so that at task failure, the cardiac activity should be higher in hypoxia than in normoxia. However, comparison of similar exercise under both conditions revealed on the contrary that the peak of cardiac activity at task failure was higher in normoxia (and hyperoxia) than in hypoxia (Noakes et al., 2001). These results suggest that instead of the anaerobiosis in the skeletal muscles, it is the anaerobiosis in the myocardia that is protected. But again, the physiology showed that athletes stop their effort before this limit is actually reached, suggesting that there is somehow a central governor that limits effort production in anticipation of any hard limit (Noakes, 2011). This anticipatory model, with a central governor, explains observations that cannot be accounted for by catastrophic models:

- The end spurt effect: when we find extra resources as we know that we are close to the end of the effort. The catastrophic models cannot explain this effect because they suppose that our resources gradually decrease as effort exertion is prolonged. On the contrary, the central governor can explain this effect: the governor anticipates that the system must be protected by inhibiting the effort production, if it realizes that the end of the effort is actually closer than expected, it can release the break and allocate more resources at the end.
- The homeostasis is preserved even under strenuous effort because the system is never driven to the limit.
- Central drugs affecting the central governor can affect performance, despite the absence of effect on the motor production per se.

- The rating of perceived exertion is more related to the relative duration of the exercise (e.g. to be at half or two third of the run) than the absolute intensity of the exercise, suggesting that it co-varies more with the planning of the effort schedule than the effort exertion.

The central governor model suffers from a lack of precision about how it might be implemented in the brain, and how it could incorporate psychological factors such as motivation. It is striking that reviews on the topic only mention psychological factors for the sake of completeness and because it seems rather intuitive, but do not provide in-depth description of the mechanisms. Besides this weakness of the model, there are at least three key ideas that should be kept in mind. First, the effort allocation is not passively enslaved to physiological limitations but benefits from anticipations. Second, this anticipatory control to regulate the homeostasis benefits from the body afferent sensory feedbacks during effort. Third, the central nervous system is the key conductor of the anticipation mechanism.

The central governor model is rather controversial. Several authors argued against this model because it may undermine the limitations ensued by physiological variables such as the cardiorespiratory function (Shephard, 2009), or because it undermines the relative weight of peripheral over central origins of fatigue (Weir et al., 2006). A reasonable position may be to acknowledge that physical fatigue arises from very different origins and is regulated also at many levels. Abbiss proposed a visual sketch of how all these sources might be combined (see Figure 16). At first sight, this figure may convey the idea that the concepts on physical fatigue are messy and not much informative in terms of predictions and reliability (notice the 'nature of the task' phrase positioned at every corner of the figure to remind the reader that the relative weight of each factor depends much on the task and parameters that are sometimes subtle). I prefer the optimistic view that a wealth of knowledge has been accumulated about fatigue across types of exercise, across people and even more, across species so that the picture, despite complex, is getting more and more understood.

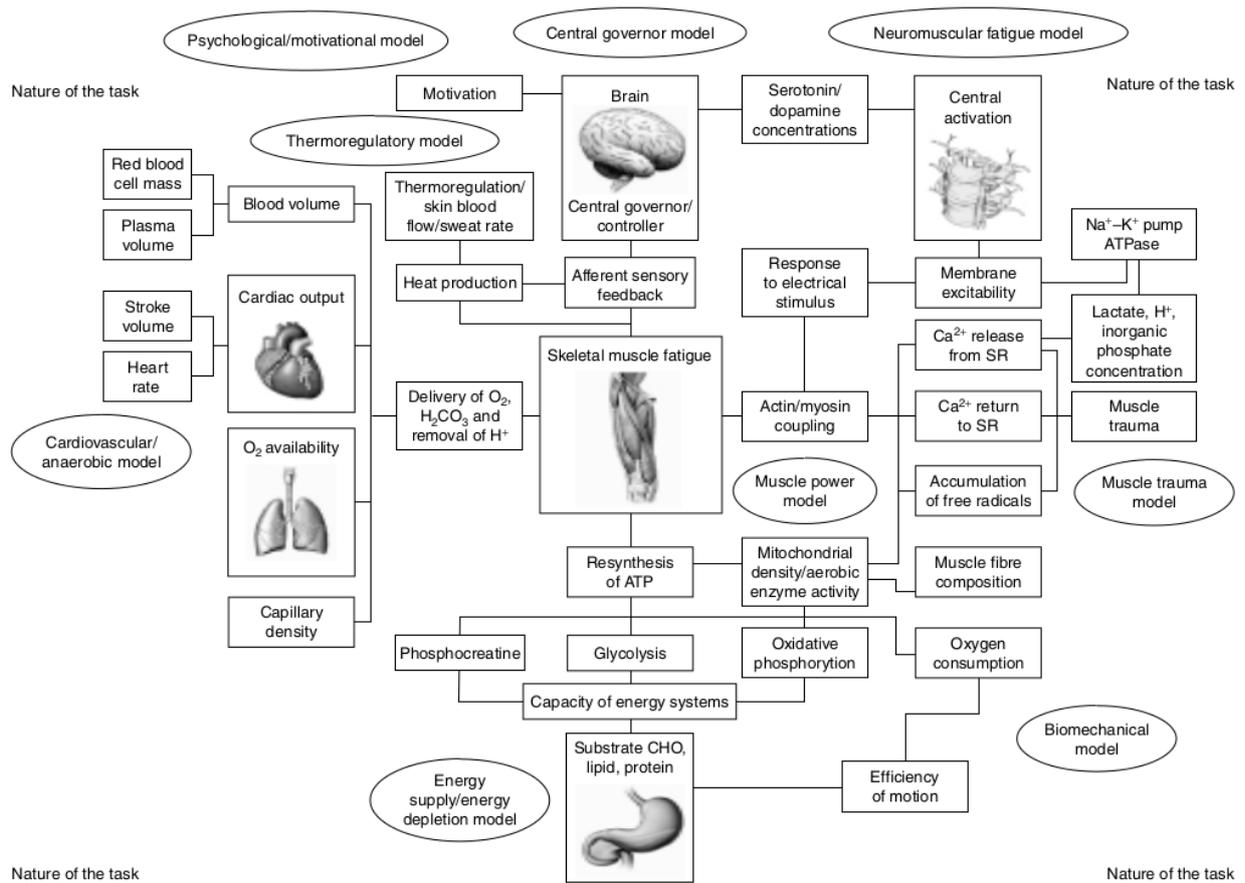


Figure 16 A multifactorial model of physical fatigue.

Abbreviations: ATP = adenosine triphosphate; ATPase = adenosine triphosphatase; CHO = carbohydrate; SR = sarcoplasmic reticulum. From (Abbiss and Laursen, 2005)

## 2.5.2 When to start the exertion?

### 2.5.2.1 The opportunity cost

The notion of opportunity cost was presented several times in this introduction. The formal definition of this notion in microeconomic theory is the value of the best alternative forgone, i.e. the highest utility that one could have got if she had made another choice (Cahuc and Zylberberg, 2004). Note that it makes sense only when alternatives are mutually exclusive, so that choosing one option corresponds to declining all the other options. The opportunity cost presses to reconsider the current choice when another alternative has a higher utility. To illustrate this notion and the way it impacts the timing of the exertion onset, let me follow up an example presented previously: the predator foraging in a patchy environment (see Figure 15). Let A and B be two options: both are strictly equivalent patches in terms of productivity, but it takes longer to go in B than to go in A. Say that the figure corresponds to patch B. The case of patch A corresponds to shifting the exploitation time to earlier latencies. The marginal value theorem by Charnov for foraging suggests that choosing A is

better than B. It can be rephrased equivalently in terms of opportunity cost: there is an opportunity cost in the travel duration that favors A over B.

The notion of opportunity cost was also introduced while addressing the invigoration of the response. Say that a rat faces a lever and that pressing this lever triggers the delivery of a reward (see Figure 17). The model by Niv and colleagues assumes that 1) the rat receives a reward (food pellet) when she presses a lever 2) that the cost to press the lever has two components, one that is fixed (there is a cost *per se* of pressing the lever, but this cost can depend on the state of the rat) and a delay-related cost (the cost is inversely proportional to the delay) and 3) the rat updates her policy (which action to choose and within which latency) following a reinforcement learning model to optimize the long-term utility (i.e. benefit minus cost), or in other words, the average utility over time (Niv et al., 2007).

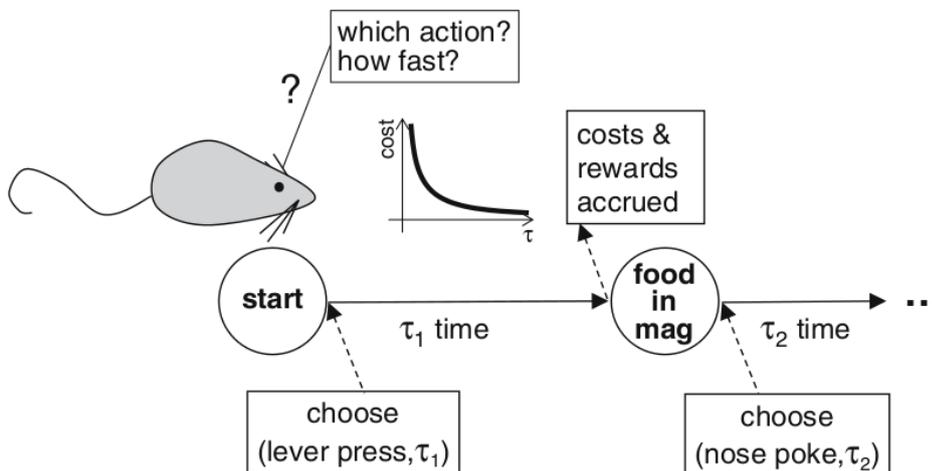


Figure 17 *Opportunity cost.*

From (Niv et al., 2007)

When the latency is very long, the average utility decreases, because the rate of reward delivery is low. When the latency is very short, the average utility also decreases, because the cost of acting fast is high and it is not compensated sufficiently by the increased reward delivery rate. For intermediate latencies, there is therefore an optimal utility. In other words, the onset timing of the exertion (the lever press) should be neither too delayed nor too brief. It is crucial in this model that the cost increases for shorter latencies; otherwise there would be no reason why the response should not be too brief. In the absence of such a latency-dependent cost, a reason why the delay should not be too short would be when there is also a reward associated to doing nothing, something like a leisure utility in the economic theory on labor supply (see page 87). It would have been interesting that Niv and colleagues had better argued why there is a latency-dependent cost instead of simply assuming it and what is the relation to temporal discounting of reward.

A factor that is worth considering with opportunity cost is whether there is time pressure in the task. The intuitive idea is that we should be more concerned with our average utility rate when the time allowed for the task is limited because not optimizing the utility at a given moment jeopardizes the utility maximization in the allotted time. It is less clear when there is no time pressure. There is behavioral evidence that people take this factor into account. For instance, Schweighofer and colleagues used an inter-temporal choice task in which the task duration was limited, so that choosing the later-but-higher reward also reduced the amount of subsequent trials. In this task, people adopted an exponential rather than a hyperbolic discounting of reward by delays and the discount rate was nearly optimal given the task duration (Schweighofer et al., 2006).

### ***2.5.2.2 The diving problem***

The diving problem was presented earlier: some aquatic mammals forage in the depth and need efficient aerobic metabolism, which depletes their oxygen budget, so that there is a limit on the time they can spend foraging (Mori, 1999). The opportunity cost framework also suits the description of this problem. In the case of the diving problem, it is very clear why the time spend at the water surface to replete the oxygen budget should be neither too long nor too short. If this duration is too long, there is an opportunity cost that corresponds to the food that would have been foraged if the mammal had dived. If this duration is too short, the oxygen budget is not sufficiently restored so that the diving will be limited to durations shorter to what the animal can have sustained if she had restored her oxygen budget completely. Given that there is a travel cost, i.e. a duration during which there is no food when the mammal dives to the depth where she can forage, a limited oxygen budget can make the dive useless if the animal does not even reach the depth. There is therefore an optimal duration for both diving and surface time. The diving problem is given here because it is a case in which the resources to manage are quite easy to identify so that this example can hopefully be inspirational for the physical effort allocation problem.

### **2.5.3 How to set the work-to-rest ratio?**

The two previous sections aimed at presenting theories that account for the duration of work or the duration of rest. However, duration of work and rest are generally not independent and common mechanisms can regulate both. This section aims at emphasizing this point.

#### ***2.5.3.1 Economics and the allocation of time***

Part of the economic literature is dedicated to how people choose how long to work. This problem is addressed by the so-called labor supply theory (Becker, 1965; Cahuc and Zylberberg, 2004). The prevailing idea is that people choose whether to participate to the labor force and to what extent they participate, i.e. how to allocate time between work and leisure. In that sense, the time

dedicated to work proceeds from a choice. The central assumption of this theory is that people seek to maximize consumption and leisure and that both are substitutable. In other words, equal levels of utility can be achieved by high consumption and low leisure duration, or low consumption and high leisure duration. The relation to work follows from the fact that work brings money and that money is needed for consumption and leisure. It is quite intuitive for consumption; regarding leisure, it follows from the fact that leisure is not just the opposite of paid work, but to work in a broader sense that include household production (raising children, housekeeping, etc.). The wage is therefore the counterpart of the price of goods and the opportunity cost of leisure.

Without further constraints and assumptions, the labor supply theory proposes a normative account of how the ratio between work and leisure should be set. The noticeable effect is that if the hourly wage increases, people tend to spend more time at work. However, the theory does not say how the ratio between work and leisure should be implemented in the allocation of time, i.e. whether it is better to work 4 days & rest 2 days, or to repeat twice 2 days of work & 1 day of rest. All that the theory says is that the wage rate of work (among other kinds of benefit and cost) are crucial to determine the ratio between leisure and work.

### *2.5.3.2 Behavioral ecology and the allocation of time*

The foraging theory is another framework that describes both how much time should be spent on foraging or not and why. The mechanisms regulating the durations of foraging and non-foraging time depend the same 'currency' like Pyke said (Pyke et al., 1977). This intuition is easy to grasp if one refers to Figure 15. With respect to the marginal value theorem (Charnov, 1976), two situations may be equivalent: a short travel time and a low yield rate in the patch or a long travel time and a high yield rate in the patch, because both result in similar mean food intake rates. The optimal travel time and the optimal yield rate of a patch actually depend on one another.

A key difference between the optimal foraging theory and the labor supply theory is that the former does not only set the ratio between foraging and non-foraging time but also sets the duration of each part. In other words and to take again the example of the diving problem, it is not equivalent to spend 2 hours at the surface and 4 hours diving and foraging vs. to repeat 1 hour at the surface and 2 hours diving. The constraints of the problem, e.g. the oxygen budget in the diving problem, have strong dynamical properties, such that the optimal policy is not a simple ratio between rest and work but actually precise durations of rest and work to alternate.

#### 2.5.4 Summary

- When allocating effort over time, we have to choose between effort and rest, knowing that there is a tradeoff between the two.
  - Effort may be more profitable than rest if it brings reward.
  - But effort is limited in time. The inability to sustain effort is unlikely to depend on a single physiological variable and leaves room for monitoring and anticipation.
  - There is a benefit to rest because it allows resource replenishment.
  - But if effort brings more benefit than rest, there is an opportunity cost to rest.
- How should we choose between effort and rest to optimize our behavior?
  - An optimization problem is characterized by a set of constraints and a quantity to maximize, which defines the utility function. Hence, to know what is optimal in effort allocation, we should define the set of constraints and a quantity to monitor.
  - The utility function and the constraints should be identified for their biological relevance.
  - The optimization is better made on the fly rather than rigidly computed in advance and then implemented, because it can overcome uncertainty.
- Effort and rest are not independent and both constrain the dynamic of the effort allocation problem.
  - If the utility of effort and rest depend on each other (like work and leisure in an economy), there should be an optimal ratio between the two.
  - Besides this ratio, if there is a dynamic of cost and benefit within effort and rest, there should also be a dynamic of effort allocation (i.e. alternation of effort and rest).



## 2.6 A proposal

None of the tasks presented in the introduction were completely satisfactory to address the physical effort allocation problem. Some tasks do operationalize the opposition between cost and benefit, but the cost is sometimes confounded with the delay, and effort is sometimes effort only in the minimal sense (e.g. (Croxson et al., 2009)). In addition, there is no choice in this latter task, so that what is studied is the neural correlate of the imposed action, not the behavior. The efforts used in the sport medicine literature, e.g. cycling on ergometer, are undoubtedly 'real' effort. However, the timescale of the exertion is the hour, which is not compatible with the variations of experimental conditions across successive trials: this is too long. Additionally, this kind of effort is not compatible with neuroimaging methods such as functional MRI and magnetoencephalography. Finally, many effort-based decision-making paradigms divide the key decisions into discontinuous steps. For instance in the T-maze task there is a crucial decision to turn left or right in the maze, but nothing in this paradigm corresponds to the problem of how the effort allocation might unfold over time.

Second, the issue of this dissertation is to understand why and how we alternate effort and rest during the course of long-run actions. None of the existing models in the literature were completely satisfactory either.

Therefore, there are two challenges in this work: to define an experimental paradigm and to define a theoretical model to account for the behavior in this task.

### 2.6.1 The scope

I started the introduction with the example of moving a fridge to illustrate the physical effort allocation problem. To better understand this problem, I tried to look at related or similar topics, like how aquatic mammals dive and the labor supply economic theory. These examples were numerous and diverse. In addition, there are many ways of instantiating experimentally the effort allocation problem, with paradigms that could better fit in the laboratory than moving a fridge. I now highlight the key features that we retained to implement effort allocation in the experimental context.

First, the benefit associated to the effort is money, because money is easy to quantify and to manipulate, because there are limited satiety effects with money compared to primary rewards and because it is a powerful incentive used in many experimental settings with healthy humans.

Second, the time scale of the effort allocation problem must fit into the routine duration of behavioral experiment. To meet the requirement of the amount of replicates required for statistical approach, the timescale of this effort allocation should fit in several seconds. The kind of effort selected to this purpose was the hand grip effort, which has the substantial advantage of benefiting

from previous experience in the host laboratory. Reduction of force production in the hand grip occurs within seconds, so that the timescale of the effort allocation problem corresponds to the mentioned requirements. Furthermore, this kind of effort is compatible with most neuroimaging methods.

Last, we used a fix reward rate during the effort. This rate might be subjected to experimental manipulation during the experiment, but during a given effort, the rate is fixed so that the duration of the effort strictly controls the amount of reward earned. This rule shortcuts several problems such as the delay discounting of the reward delivery.

It is much shorter to say what kind of effort allocation problem is implemented in our experimental paradigm than to list what this paradigm is unlike to. The above-mentioned choices on the types of effort and incentive are not meant to dwarf the effort allocation problem, but rather to operationalize it. I will discuss potential limitations to this approach the final discussion.

### 2.6.2 The model

The purpose of the model is to account for the timing of the decision to stop and to resume the effort. The problem to solve is therefore essentially dynamical. Given the concepts presented in the introduction, we can assume that these decisions follow rational explanations, i.e. that we do not decide when to stop or resume the effort randomly but so as to optimize our behavior. That effort allocation behavior is strategic is an assumption that will be tested experimentally. What is the dynamic of cost and benefit in the effort allocation task? By design, the dynamic of benefit is stationary over time: there is fixed reward rate during effort and no reward during rest. By contrast, the dynamic of cost is more complex: due to fatigue the cost entailed by producing a time unit of effort increases along the effort duration and due to recovery it decreases along the rest duration.

Given that, on the one hand, the decisions to explain are timings of effort and rest, and on the other hand that the effort cost varies over time, a simple model to start with would be an accumulation-dissipation process. This model relies on three critical assumptions (see Figure 18). First, there are accumulation and dissipation: the dynamics should be in opposed directions between effort and rest. Second, the fluctuations are bounded so that there is a limit that governs effort duration. Third, there is a similar bound to govern effort resumption.

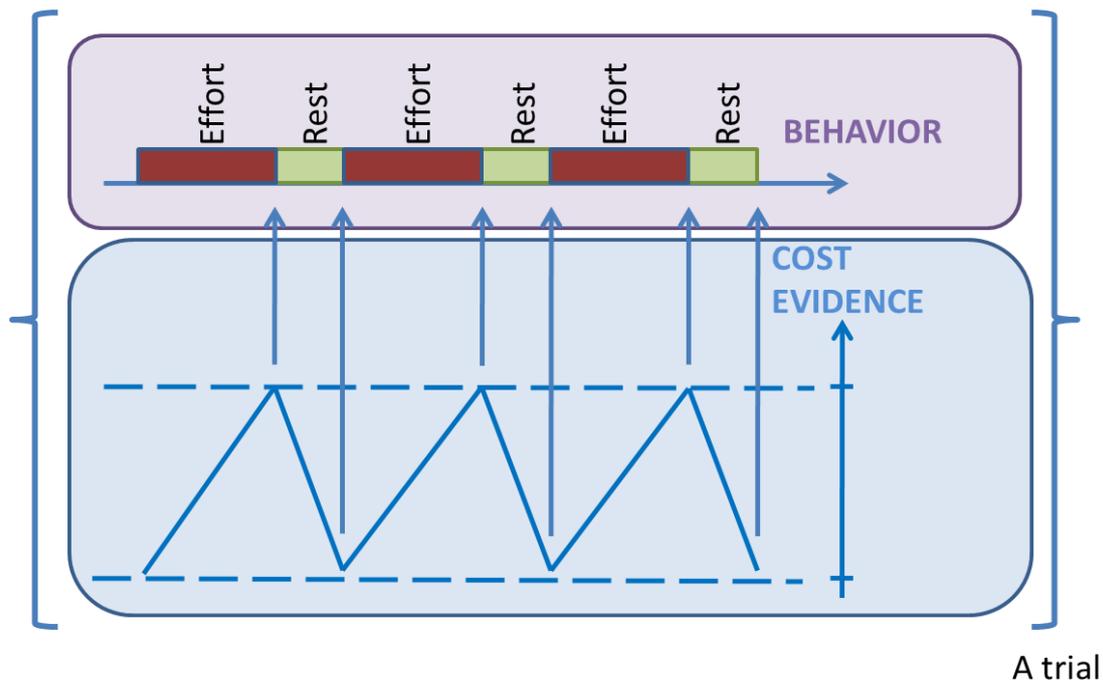


Figure 18 An accumulation model to account for temporal allocation of effort.

There are several advantages to this accumulation approach. First, this way of putting things is intuitive and simple. The accumulation-to-bound process triggering effort cessation corresponds exactly to the intuition by William James presented at the beginning of the introduction. The model simply translates in computational terms the pervasive idea that we monitor effort cost and stop the effort when it is excessive. What we add with the present model compared to James' intuition is that a similar mechanism could occur at rest: effort cost decreases during rest due to recovery and we resume effort when this cost is sufficiently low. This is an additional advantage: the same variation-to-bound mechanism is used during effort and rest, which is parsimonious. Second, this model is biologically plausible. The accumulation model is by essence related to the decision timing, which is the key behavioral dependent variable to explain. There are numerous reports in neuroscience of accumulation-to-bound mechanisms triggering decisions, and there are several candidate signals in the brain reflecting effort cost that could serve as input to the model, like the proprioceptive and nociceptive signals. Third, this model is economically advantageous. Indeed, the effort is stopped when cost is too high, which ensures that the benefit is earned with a limited cost. In addition, the effort is not resumed immediately after cessation but later so that the following effort is initiated at a more profitable cost. Through these two mechanisms, excessive physical costs are prevented, which is safer for the body, but also economically reasonable. Does this heuristic strictly correspond to economic utility maximization? We could compute off line the optimal behavior; however, this would require that we know exactly the cost function. This function depends on precise physiological properties that are beyond our knowledge; defining such a function suffers in particular from a lack

of data related to recovery. In addition, this problem can be solved off-line only if the relevant information is completely available, which is often not the case. For example in the fridge problem, I am not an expert of this situation: I do not know how hard it is so that planning beforehand is impossible. By contrast, the proposed model leaves room for on-line estimation.

What would be the variable accumulated in this model? How is it affected by motivation or by the task difficulty? If one follows William James' intuition, the accumulated variable is the effort cost. In this case, the bound to stop the effort may be shifted to allow more effort for higher incentives. An alternative model is that the accumulated variable is directly the effort utility. In this case, the dynamic would be the opposite direction of that presented in Figure 18: the effort utility decreases during exertion due to fatigue. In addition, the bound to stop this reduction of utility would be constant and correspond to when effort is no longer the best option to follow. Other variables with waning and waxing dynamics could be suggested. What matters is that computationally, the nature of the accumulated variable should be reflected in how the dynamic is affected by task parameters such as the effort reward rate and the effort difficulty, in particular, whether and how the bounds or the slopes of this variable are affected. The nature of this variable also corresponds to distinct psychological interpretations and hence, potentially distinct neural underpinning.

I am sorry that I have to kill the suspense now, but it turned out that the experimental results favored that the accumulated variable is related to the effort cost. I prefer to avoid unnecessary confusion and not to change the name of this variable to account for the incremental experimental refinements. This variable is given a single name all thorough this manuscript: cost-evidence.

### **2.6.3 Outline of the work**

As outlined above, the proposed cost-evidence accumulation model for the effort allocation problem makes critical assumptions that need to be tested experimentally and leaves several options open that could be constrained by behavioral and neuroimaging results.

I chose to present the results of my work as an incremental construction. The reason is that the work truly started with an intuition that an accumulation process could underpin the effort allocation problem. This intuition was progressively refined based on experimental results, from both behavioral and neuroimaging evidence. These results unraveled some constraints on this model and also raised other issues so that the results of some experiments appealed other experiments to be carried out. The research path was not linear; the behavioral and neuroimaging results were actually intermingled. For the sake of clarity I will nonetheless split the behavioral and neural results and try to show how this experimental work gradually participated in refining the description of the model.

### ***2.6.3.1 Behavior: what the effort allocation is adapted to***

In the first part of the results, I focus on behavioral experiments that aimed at understanding what the effort allocation is adapted to. This section checks some assumptions that are critical to the model. A first assumption is that the decisions to stop or resume the effort are made depending on the level of cost-evidence: the behavior is adapted on-line to the cost-evidence level. A second assumption is that the effort allocation behavior is strategic, so that this on-line adaptation is not simply a reactive mechanism, but also implements strategic control to optimize benefit against cost. With respect to cost, a distinction will be made about the true effort cost and the expectation on the effort cost. The way the behavior is adapted to these two estimates argues in favor of the fact that the accumulated variable is related to a physical cost. Because I argue that the behavior can be accounted for by an accumulation model, this model, in particular the slopes and bounds, will be fitted on the behavior to show that the effects of the incentive, the experienced effort cost and the expected effort cost have distinct computational roles. This computational view of the behavior facilitates the description, from a psychological and cognitive perspective, of the mechanism at play. Finally, I raise the issue of how this computational variable relates to introspection.

### ***2.6.3.2 Neuroimaging: functional anatomy of the effort allocation***

In the second part of the results, I start with the search in neuroimaging data for a correlate of the computational variable inferred from the behavior. Such a correlate would support the idea that the computational variable is implemented in the brain. This implementation provides additional evidence (although not definite) that this variable could underpin the effort allocation behavior and the anatomy of this correlate may suggest how this mechanism is implemented in the brain. In other words, the aim is to give a functional description of the process. The second aim of the neuroimaging studies is that, in the strong interpretation, they potentially give a full access to the accumulated signal in the brain, whereas the behavior only corresponds to what happens when this variable reaches extreme values: maximal values correspond to the transition from effort to rest and minimal values correspond to the transition from rest to effort. This full access to the accumulated variable will be the opportunity to constraint degrees of freedom in the model that cannot be constrained by pure behavioral data. The point is that behavioral data predict that the bounds of the signal are affected by some experimental factors, but all they can say is about the distance between the bounds, so that changing the upper or the lower play symmetrical results, despite distinct psychological interpretations.

In this second part of the results, I also focus on the rest periods that play a critical role in the effort allocation process. These rest periods implement in particular an optimization of the behavior to

reduce the opportunity cost. I will show that the motor beta synchrony, a signature in the frequency domain of motor control (13-30 Hz) is a likely candidate to translate the motivation of opportunity cost reduction into a behavioral adaptation: the faster resumption of effort.

Quite intriguingly, the data from the neuroimaging studies reveal a negative finding: no sustained effect of the incentive was found in the data, despite a sustained effect on the behavior. The absence of effect is certainly not evidence that there is no effect; however, it lets open the question how the sustained effect of the incentives might be implemented to affect the behavior. I suggest that the incentive effect may be reflected in an increased arousal, or increased activation state, based on results from the cardiac frequency.

### ***2.6.3.3 Manipulation of the brain and effort allocation***

In the last section of the experimental section, I address three questions related to effort allocation. The results presented in the two other sections of the results suggest that: incentive motivation is one (out of several) determinants of effort allocation and that the theoretical variable that accounts for effort allocation behavior is related to the effort cost. The questions addressed are therefore: Can the role of benefit in the effort allocation process be modified by altering the motivational state of the subject? Effort cost could be signaled through nociception, hence: Do analgesics change the effort allocation? Effort cost could also be signaled by fatigue-related neuromodulators such as serotonin, hence: does serotonergic manipulation affect the effort allocation? In four separate studies, hypnotic suggestion was used to modulate the motivational state of subjects or to induce analgesia; analgesia was also induced using paracetamol and last, the serotonergic manipulation was induced by a selective-serotonin reuptake inhibitor (Escitalopram). Subjects in all groups were healthy young adults.

### 3 Experimental work



## 3.1 How is effort allocated over time? Evidence from behavioral data

### 3.1.1 Decisions to stop and resume effort are based on the on-line level of cost-evidence

#### 3.1.1.1 Introduction

The aim of this first set of three studies was to check whether the behavior is adapted to the level of cost-evidence on the fly. To this end, we adopted a dose-response design. The pattern of effort production was imposed to manipulate the cost-evidence level and to observe the subsequent subject's behavior. We manipulated cost-evidence by varying only the durations of effort and rest, since 1) given the cost-evidence model, manipulating durations is a way to control the cost-evidence level, 2) the duration manipulation is symmetric between effort and rest, whereas the effort difficulty for instance as no intrinsic counterpart during rest.

More specifically, this first set of studies should test the four basic assumptions of the cost-evidence accumulation model. The reader can refer to Figure 18 for a graphical view. First, *ceteris paribus*, longer hand grip effort duration should increase the cost-evidence level. Second, the cost-evidence accumulation is bounded so that there is a limit to effort production. Third, the cost-evidence level should be dissipated during rest so that there is a progressive recovery and fourth, this recovery is also bounded, so that long rest durations do not dissipate cost-evidence further than intermediate rest durations.

Different sets of subjects participated in three tasks: first we observed the durations of effort produced after increasing cost-evidence with an imposed effort (Task 1), second, we observed the durations of effort produced after dissipating cost-evidence with an imposed rest (Task 2) and last, we observed when subjects resume exertion and ultimately stop after increasing cost-evidence with an imposed effort (Task 3).

#### 3.1.1.2 Results

In all three tasks, trials comprised two parts: a first part with an imposed pattern of effort and rest duration to manipulate the level of cost evidence, and then a second part during which participants were free to behave. In the two first experiments, we observed the effort duration, and in the third task, participants were allowed to adjust rest duration before we observe the effort they produced freely. These situations are forms of effort allocation. The aim was to observe whether people allocate effort depending on the on-line level of cost-evidence, i.e. whether the effort and rest durations imposed constrain the effort and rest duration produced subsequently. Note that in the three tasks, there was an incentive to produce effort in the free effort of each trial because participants received a reward proportional to the time of free exertion.

A summary of the design and main results across the three tasks is presented in Figure 19.

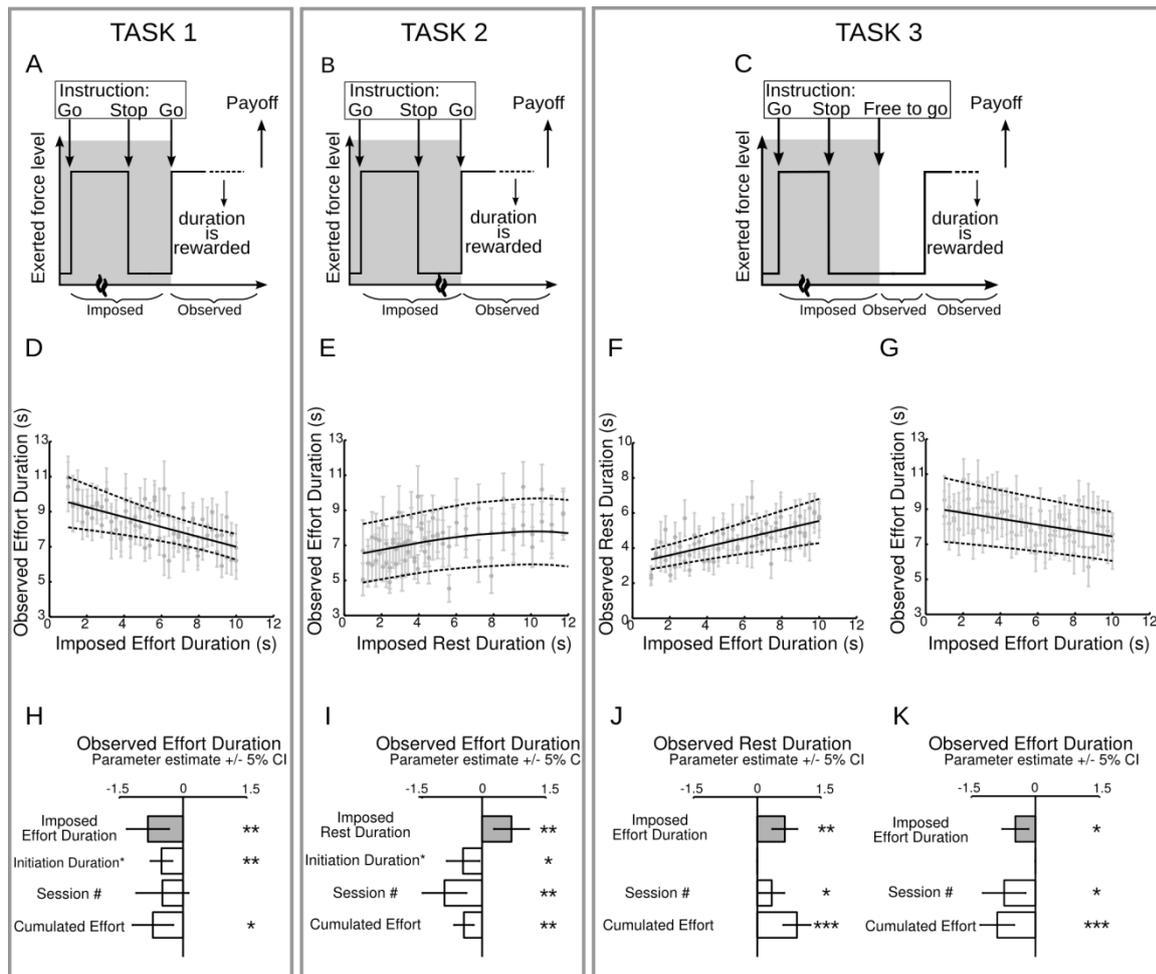


Figure 19 Local adaptation of cost-evidence.

Three tasks are presented column-wise.

Top row (A, B, C) depicts the paradigms: gray shading denotes that the participants' actions are imposed; no shading denotes that they are free. Each plot sketches the exerted force level over time. The wave-like stroke on the horizontal axis denotes that the effort (A, C) or rest (B) durations were varied experimentally. Note that the imposed rest (A) and effort (B) duration were constant across trials.

Middle row (D, E, F, G) presents the group mean durations +/- s.e.m. Two points are plotted for a given horizontal position, corresponding to each hand. The black line is the group average of the model fit estimated at the subject-level; dash lines size the 5% CI of the average.

Bottom row (H, I, J, K) presents the statistical results with the effect sizes, with factors of interest highlighted in grey. The Initiation Duration\* denotes that the variance due to the cumulated effort and session # was regressed out prior to estimation. P-value (bilateral test): \* $<0.05$ , \*\* $<0.005$ , \*\*\* $<0.0005$ .

### 3.1.1.2.1 Task 1: Cost-evidence limit effort production

In the first task, participants were asked, in each trial, to exert a constant force level (60% of their maximal force) with a duration that was varied experimentally and imposed by a stop signal without prior notice. Then they rest for two seconds, and finally they were asked to resume effort again at a

go signal. The duration of this last effort was not instructed: subjects decided how long to exert the effort; knowing that this duration determined proportionally their payoff (see Figure 19A). To ensure that rest duration was well controlled, we checked that initiation duration (the duration between the go signal and the effort initiation) was not impacted by the imposed effort duration (weak trend:  $p = 0.09$ ), and was not affected by the effort duration previously cumulated in the current session ( $p = 0.76$ ) nor by the session number (0.25).

The critical prediction of the cost-evidence model is that the observed effort duration should decrease after longer imposed effort duration. This effect was significant ( $p = 0.0037$ ), see Figure 19D & H.

In addition to this critical effect, we also noted that the behavior was impacted by other factors, corresponding to different time scales. The observed effort duration was reduced by the past effort duration cumulated in the current session ('Cumulated effort',  $p = 0.008$ ) but did not change over sessions ('Session #',  $p = 0.11$ ). This could reflect a fatigue effect that accumulated over trials but remained low between sessions, possibly because hands were switched between sessions. We also estimated in the same model the effect of the initiation duration on the observed effort duration and found a significant effect ('Initiation duration\*',  $p = 0.001$ ). Even if the effects of the imposed effort duration, the previously cumulated effort and the session number were marginal, we removed this source of variance from the initiation duration regressor to isolate its pure effect. The negative correlation found could reflect trial-to-trial fluctuations of fatigue on top of (and orthogonal to) the experimental manipulation: when less fatigued on some trials, participants initiate the effort faster and sustain it longer.

We also tested whether for short constrained effort durations, the (constant) rest allowed could reset the effect of the cost accumulated during the imposed effort, so that there is a minimal effort duration below which the constrained effort duration has no impact on the observed effort duration. To test this, we compared three models of observed effort durations predicted by the imposed effort duration: 1) a linear effect (no saturation), and two models of saturation at short imposed effort durations: 2) a linear effect bounded by an upper plateau, 3) an exponential asymptotic plateau. Bayesian model selection was used to discriminate the best model. Family comparison revealed that the model without saturation was far better than the models with a saturation effect (model #1 vs. #2 & #3,  $x_p=0.96$ ), arguing against any saturation effect. This could mean that the imposed rest (2s) is not sufficient to clear fatigue accumulated even from short effort duration.

### 3.1.1.2.2 Task 2: Rest progressively dissipates cost-evidence

The second task was quite similar to Task 1, except that constant effort and variable rest durations were used instead of variable effort and constant rest durations (see Figure 19B). Again, as we wanted to strictly control rest durations, we checked that delays at the go signal were not affected by the imposed rest duration ( $p = 0.10$ ). There were neither affected by the session number (0.46) but slightly by the past cumulated effort duration ( $p = 0.03$ ).

The critical prediction of the cost-evidence model is that the observed effort duration should increase after longer imposed rest durations. This effect was observed ( $p = 0.0035$ ), see Figure 19 E & H.

In addition to this critical effect, we also noted, as in Task 1, that the behavior was impacted by other factors, corresponding to different time scales. The observed effort duration was progressively reduced over sessions ( $p = 0.003$ ) and along with the past effort duration cumulated in the current session ( $p = 0.002$ ). This probably reflects an accumulation of fatigue both between and within sessions. We also estimated in the same model the effect of the initiation duration on observed effort durations and found a significant effect ( $p = 0.026$ ). This negative correlation could reflect additional trial-to-trial fluctuations of fatigue as in Task 1.

We also wanted to test whether long imposed rest durations could reset completely the effect of the cost accumulated during the imposed effort, so that longer rest durations would not increase further the observed effort duration. To test this, we compared three models of observed effort duration based on the imposed rest durations: 1) a linear effect (no saturation), and two models of saturation at high imposed rest durations: 2) a linear effect bounded by an upper plateau, 3) an exponential asymptotic plateau. Bayesian model selection was used to discriminate the best model. The saturation family (model #2 and #3 vs. #1) had the highest expected frequency (0.79) and exceedance probability ( $x_p=0.94$ ). Comparing directly model #2 and #3 revealed that the asymptotic saturation was more likely than the linear plateau ( $x_p=0.98$ ).

### 3.1.1.2.3 Task 3: Rest is adapted to cost-evidence levels

This task was quite similar to Task 2, except that participants were not asked to resume effort immediately at the go signal, but when they were willing to, so that there were two dependent variables of interest: the observed rest duration and the subsequently observed effort duration (see Figure 19C).

The critical prediction of the cost-evidence model is that participants should not resume effort until the cost-evidence is cleared to a low level. Hence, cost-evidence being increased by longer effort,

participants should increase rest duration after longer imposed effort durations. This effect was significantly observed ( $p = 5 \cdot 10^{-4}$ ), see Figure 19F & J.

In addition to the critical effect predicted for rest duration from the cost evidence model, rest was also affected by factors at other time scales. The observed rest duration was progressively increased over sessions ( $p = 0.028$ ) and along with the past effort duration cumulated within sessions ( $p = 5 \cdot 10^{-5}$ ), probably reflecting that more rest is allocated to compensate the accumulation of fatigue that builds up over the experiment, as in Tasks 1 & 2.

One may have speculated that the rest participants allocate themselves is sufficient to reset completely the fatigue effect ensued by the preceding imposed effort. This was not the case as the observed effort duration was significantly decreased by longer imposed effort duration ( $p = 0.006$  - see Figure 19G & K). Note that the rest duration necessary to fully compensate the 7s of effort imposed in Task 2 was, according to the models with plateau fitted for each subject,  $8.8s \pm 0.7$  s.e.m. According to the linear fit in Task 3 (Figure 19F), subjects rested on average for  $4.8s \pm 0.4$  s.e.m. after 7s of effort. The comparison of fit between tasks suggests that subjects may not rest long enough to compensate completely the imposed effort, which could be the reason why the imposed effort duration has an impact on the observed effort duration. In addition, the absence of complete reset is also in line with the reduction of effort duration over sessions ( $p = 0.008$ ) and within session ( $p = 4 \cdot 10^{-4}$ ), that is likely to reflect the accumulation of fatigue that is never completely cleared, like in Task 1&2.

The rest participants allocate themselves could attenuate the effect of the accumulated cost-evidence, i.e. a partial rather than a complete compensation. To estimate this effect, we compared Task 1 (imposed constant rest) and Task 3 (free rest). The accumulated cost-evidence effect was estimated as the slope between the imposed effort duration and the observed effort duration. We found that the slope was shallower when subjects could rest freely between the two efforts (Task 1, slope  $-5.68 \pm 1.5$  s.e.m; Task 3, slope =  $-3.6 \pm 1.0$  s.e.m), however the difference did not reach significance ( $p=0.28$ , bilateral t-test).

There is another interesting comparison between Task 1 and Task 3. In Task 3, after 2s of imposed effort, participants took a rest longer than 2s. This could explained that in Task 1, 2s of rest were not enough to completely clear cost evidence induced by 2s of effort, so that no saturation effect was observed for short effort duration.

### *3.1.1.3 Discussion*

To summarize the findings of these three tasks, it seems that the effort allocation is adapted the fly to effort-induced level of cost evidence. We manipulated this level by varying the effort and rest duration and showed that people exert shorter effort when cost-evidence is previously increased by longer effort and they exert longer effort when cost-evidence is previously decreased by longer rest. Our data revealed that cost-evidence imposes a higher bound on effort production but also has a lower bound that corresponds to the maximal improvement rest can bring. We also showed that when people are allowed to adapt rest duration, they lengthen rest to compensate the accumulated level of cost-evidence.

This set of results demonstrates that over several seconds, the kind of effort used induces local constraint on the effort allocation. The results can be reframed in terms of the cost-evidence accumulation model: on the time scale of several seconds, the hand grip effort or rest manipulated here impacted the level of cost-evidence accumulated or dissipated so that the subsequent effort or rest depended on this level. This suggests that this kind of experimental setting is suitable to test the effort allocation problem at a short time scale, compatible with the duration of experiments in the laboratory.

Before looking ahead to new experiments, it should be stressed that this set of studies also reveals two other interesting findings. First is the fact that, on top of the 'local' variations of cost-evidence level that constrain the effort allocation, there is also a global trend that progressively impairs effort production. By definition of fatigue, we can rephrase this observation: a fatigue effect builds up progressively along the task. It seems that in these tasks, the recovery never reset the fatigue effect completely. A second additional finding is that, on top of the local effect of cost-evidence level manipulated experimentally and the global fatigue, there is another factor that is not controlled experimentally. The effort initiation speed co-varies with the effort duration, such that faster initiations are associated to longer efforts. Crucially, this correlation was estimated independently from the factors manipulated (the effort or rest duration imposed, the cumulated effort duration and the session number). This suggests that there are fluctuations in effort production that are not accounted for by other task parameters.

In this set of tasks, the behavior was adapted to constraints manipulated locally that determined the cost-evidence level. This is the basic mechanism of the cost-evidence model of effort allocation. The second set of experiments adds new parameters in the problem: the effort difficulty and the reward rate associated to effort to assess whether the effort allocation behavior also integrates strategic effects.

### 3.1.1.4 Methods

#### 3.1.1.4.1 Set up

We used homemade power grips composed of two wood cylinders compressing an air tube when squeezed. The tube was connected to a transducer converting air pressure into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to the force exerted. The signal was amplified and digitized by homemade device. The digitized signal was read by a Matlab program (The MathWorks Inc., USA).

#### 3.1.1.4.2 Pre-processing

The effort onset was determined as the first sample exceeding 20% of the participant maximal force.

#### 3.1.1.4.3 Maximal force estimates

For all tasks, we measured the maximal force for each hand before starting the task, following published guidelines (Gandevia, 2001).

Participants were verbally encouraged to squeeze continuously as hard as they could, until a growing line displayed on a computer screen reached a target. The growing rate of the bar was not indexed on the participant exerted force level but constant to last 5s. Maximal force was set to the average of data points exceeding the median. Then subjects were provided a real-time feedback about the force produced on the handgrip, which appeared as a fluid level moving up and down within a thermometer, the maximal force being indicated as a horizontal bar at the top. Subjects were asked to try outreaching the bar and state whether it truly corresponded to their maximal force. If not, the calibration procedure was repeated.

#### 3.1.1.4.4 Behavioral tasks

All tasks were presented on a computer screen, and were programmed with Matlab using Psychtoolbox (<http://psychtoolbox.org>).

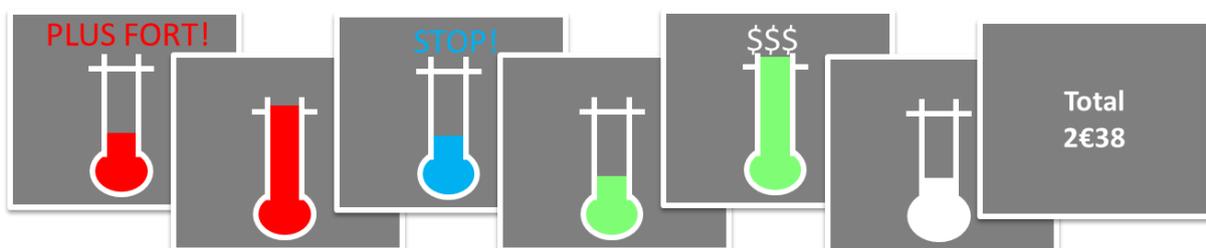


Figure 20 Screens displayed to participants.

The display was quite similar for all tasks. The exerted force level was always displayed as a fluid moving up in a thermometer, the target level on the top of the thermometer indicated 60% of the participant maximal strength. The color of the fluid in the thermometer indicated what to do: red for

constrained effort, blue for rest ('STOP!' was also display above the thermometer), green when the participant should initiate the effort, either immediately in Task 1&2, or, when she wants to in Task 3. When the participant stopped squeezing, more precisely at the first force sample under 50% of their maximal force, the color turned to white to indicate that the participant should rest until the following trial. For all tasks, when the force level was under the target level (60%) during the imposed effort, 'PLUS FORT' (meaning 'harder') was displayed above the thermometer. For Tasks 1 & 2, if the participant initiated the trial too late (more than 1s after the color change), the color turned to white and the message 'VOUS AVEZ APPUYE TROP TARD' (meaning 'you squeezed too late') was displayed. In all three tasks, when the participant was free to choose the final effort duration (green color) and when the force level was above the target level, a flickering dollar symbol was displayed to indicate that the effort was rewarded.

For each trial, the payoff was proportional (with a fixed rate across trials) to the time spent above the target force level during the effort of observed duration. The trial payoff and the cumulated payoff were displayed at the end of each trial.

**Task 1 (variable imposed effort, constant imposed rest, free effort).** Each trial comprised the following event: imposed effort (at 60% of maximal force), imposed rest (2s), go signal to initiate an effort of observed duration (20s allowed), feedback (2s), inter-trial interval (2s). The imposed effort durations were 36 points equally spaced between 1s and 10s. The same 36 durations were presented to the left and right hand in the same randomized order. The randomization was as follow: for each participant, the sequence 1, 2, 3, 4 was randomized into  $v_1, v_2, v_3, v_4$ . For session  $i$ , effort durations were picked up every 4 points in the sorted 36 effort duration values, starting at sample  $v_i$ . This procedure ensures that over subjects, the mean effort duration used in a session was constant.

**Task 2 (constant imposed effort, variable imposed rest, free effort).** Each trial comprised the following events: imposed effort (7s at 60% of maximal force), imposed rest, go signal to initiate an effort of observed duration (20s allowed), feedback (2s), and no inter-trial interval. Imposed rest durations were computed as follow: we simulated a mixture of Gaussians (10000 points):  $\frac{3}{4}$  of points were from  $N(3, 2)$ ,  $\frac{1}{4}$  of points were from  $N(10, 2)$  ( $N(m, \sigma)$  denotes a Gaussian distribution with  $m$  mean and  $\sigma$  standard deviation), this distribution was cut off to retain values higher than 1s, finally we picked the 37 equally spaced percentiles of this distribution and discarded the last percentile to end up with 36 values. This procedure was used to acquire more data at for intermediate rest durations than for long durations, to better characterize the expected increase in force produced than the plateau. The same 36 durations were presented to the left and right hand in the same randomized order. The randomization technique for durations was the same as for Task 1.

**Task 3 (variable imposed effort, imposed rest, free to rest and free effort).** Each trial comprised the following event: imposed effort (at 60% of maximal force), imposed rest (2s), a signal to indicate that the participant can initiate an effort when she wants, hence a rest and effort of free durations (20s allowed for observed effort and rest in total), feedback (2s), inter-trial interval (2s). The constrained effort durations were 36 points equally spaced between 1s and 10s. The randomization technique for duration was the same as for Task 1.

#### 3.1.1.4.5 Participants

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. All subjects were recruited via email within an academic database and gave informed consent prior to participating in the study. Participants were between 20 and 39 years old, with no self-reported psychiatric or neurological history, no self-reported current psycho-active substance consumption, no restriction on handedness. Participants were paid in cash at the end of the experiment. The payoff was partitioned into a fixed amount and additional amount that corresponded to the money won during the task (respectively 'Fixed' and 'Var' in the table). For Task 3, one participant was excluded because of calibration issues and another for cheating (repeated, direct manipulation of the air tube). 'Exp.' refers to the author who collected the data (FM: Florent Meyniel).

Task	Exp.	Period	N included	N male	N excluded	Mean age +/- s.e.m.	Fixed (€)	Var. (€)	Var. (€) range
Task 1	FM	03/2012	12	2	0	22.7 +/- 0.8	10	10.1	7 - 13
Task 2	FM	03/2012	12	0	0	21.9 +/- 0.4	10	9.6	4 - 15
Task 3	FM	03/2012	12	4	2	21.7 +/- 0.7	10	10.3	6 - 15

#### 3.1.1.4.6 Statistical analysis

Effort and rest durations were analyzed using multiple linear regressions. For Task 1, the 4 explanatory variables of the observed effort duration were the imposed effort duration of the current trial, two additional factors (the session number, and the effort, imposed and observed, cumulated session-wise), and the residual initiation duration of the effort, that is, the duration between the go signal and the actual effort onset detected. To ensure that explanatory variables are orthogonal, we took the residual initiation durations after regressing out the imposed effort duration and the two additional factors. The parameter significance was estimated with a two-tailed t-test at the group level. For Task 2, the same 4 explanatory variables were used, except that the manipulated

factor was the rest duration. For Task 3, the observed effort and rest duration were analyzed with the same model as for Task 1, except that there was no residual initiation duration regressor in the model.

#### 3.1.1.4.7 Bayesian model selection

To perform model selection, models were first estimated for each subject using a variational Bayes approach under the Laplace approximation (Friston et al., 2007; Daunizeau et al., 2009), using a toolbox by Jean Daunizeau (available at <http://sites.google.com/site/jeandaunizeauswebsite/>). This algorithm not only estimates linear and non-linear models but also estimates their evidence based on a free-energy approximation (Friston et al., 2007). The evidence of a model is the probability of observing the data given this model. This probability corresponds to the marginal likelihood, which is the integral over the parameter space of the likelihood of the parameterized model weighted by the prior on its parameters. This probability increases with the likelihood (which is the accuracy of the fit) and is penalized by the integration over the parameter space (which is the complexity of the model). The model evidence thus represents a trade-off between accuracy and complexity and can guide model selection (Stephan et al., 2009). Model selection was performed with a random-effect analysis at the group level to estimate the model exceedance probability based on the log-evidences of each subject for each model, using the Gibbs sampling implemented in SPM8 (Statistical Parametric mapping, Wellcome Department of Imaging Neuroscience, London, UK) (Stephan et al., 2009). The exceedance probability of a given model within a set of models is the probability that this model is more likely than any other models of the set, given the data from all subjects. Family-level inference was conducted similarly to model-level inference after defining a partition within the model space as described in Penny et al. (2010) and implemented in SPM8.

### 3.1.2 Strategic effects in cost-evidence based effort allocation: when the effort difficulty and the incentive level matter

#### 3.1.2.1 Introduction

In the general introduction to this dissertation, the opposite effect of cost and benefit on effort production (and on choice in general) was stressed. The aim of the new set of studies presented here is to address how two determinants of cost and benefit affect effort allocation: the difficulty of the effort produced and the reward that is earned through exertion. In other words, the question is whether the effort allocation behavior is strategic: does it maximize benefit against cost?

The previous and the current set of studies were designed to reveal, respectively, the mechanical and low-level constraints of effort allocation on the one hand and the higher-order, more strategic effects on the other hand. It does not imply that subjects had no strategic monitoring in the first set of

studies. They probably had since they were working for money: the reward earned along observed efforts was commensurate to their duration; however the money at stake or the effort difficulty were not varied experimentally to investigate this strategic monitoring. Such experimental manipulations are now introduced. I believe the distinction between mechanical constraints and strategic effects is worth to understand the effort allocation process. It seems rather consensual that the reward aspect of effort could pertain to strategic determinants (by definition) of the effort allocation problem. However, it is not so clear to which level the effort difficulty should be associated to. Indeed, the effort difficulty is a parameter of the effort biomechanics, but it is also a cost and hence, could be taken into account at a strategic level. Disentangling these two effects of effort difficulty is also the aim of this new set of studies.

### **3.1.2.2 Results**

Three effort allocation tasks were designed. They are all variations around the same common rules. Contrary to the first set of studies, here, participants were completely free to allocate their effort: there was no imposed rest or effort duration. The experiment was split in sessions of several trials. Each trial started by revealing the monetary incentive with a coin image displayed. Then subjects had 30s to win as much money as possible. They knew that the payoff was proportional to both the incentive and the time spent above the target bar displayed at the top of the thermometer, always at the same position. The force needed to reach this bar (a percentage of subject's maximal force, between 70 and 90%), i.e. trial difficulty, varied between trials. Subjects were provided with online feedback on both the exerted force (with a fluid level moving up and down within a thermometer) and the trial-wise cumulative payoff (with a counter displayed above the thermometer). The counter was only started when fluid level was above the target bar, with a rate proportional to the current incentive. Each trial ended with a 2s display of the session-wise cumulative payoff.

#### **3.1.2.2.1 Implicit effort allocation task**

The first task is illustrated in Figure 21A. This version is termed 'implicit', because there is no explicit cue on the screen about the difficult level. The first key observation is that participants alternated effort and rest within trials, see Figure 21B for an example. In principle, they could have spent their time squeezing continuously the grip above the target force level to accumulate the maximal amount of money. However, the task was designed with a trial duration that is long enough and target force levels that are sufficiently high so that subjects are unable to squeeze continuously the grip. Rather, they chunked their effort production in several effort periods per trial (mean amount of effort period per trial: 2.95 +/- 0.2 s.e.m. across subjects). In other words, the task makes it possible to investigate the effort allocation problem in the laboratory. Moreover, this effort allocation is implemented

within trial lasting 30s, so that trials can be repeated and factors varied across trials to investigate how they impact on effort allocation. Three levels of difficulty and three levels of incentive were crossed, so that the task comprised in total nine conditions, varied from trial to trial and repeated 8 times for each participants.

The duration of every single effort and rest period (not the total over the trial) was extracted to assess incentive and difficulty effects. Incentives significantly affected the duration of both effort ( $F_{2,72}=25.3$ ;  $p=1.6 \cdot 10^{-6}$ ) and rest ( $F_{2,72}=25.2$ ;  $p=4.9 \cdot 10^{-7}$ ) periods, whereas difficulty only affected effort duration ( $F_{2,72}=42.8$ ,  $p=2.8 \cdot 10^{-12}$ ), not rest duration ( $F_{2,72}=0.1$ ,  $p=0.86$ ). There was no significant *incentive x difficulty* interaction, neither for effort ( $F_{4,144}=1.7$ ,  $p=0.18$ ) or rest ( $F_{4,144}=1.4$ ,  $p=0.26$ ) duration. Thus, subjects spent more time squeezing and less time resting for higher incentives, and less time squeezing with higher difficulty (Figure 21C). This pattern of result show that participants tried to maximize their payoff while minimizing the effort in the task. This observation is crucial because it suggests that the effort allocation is not random but subserves a goal: it is strategic.

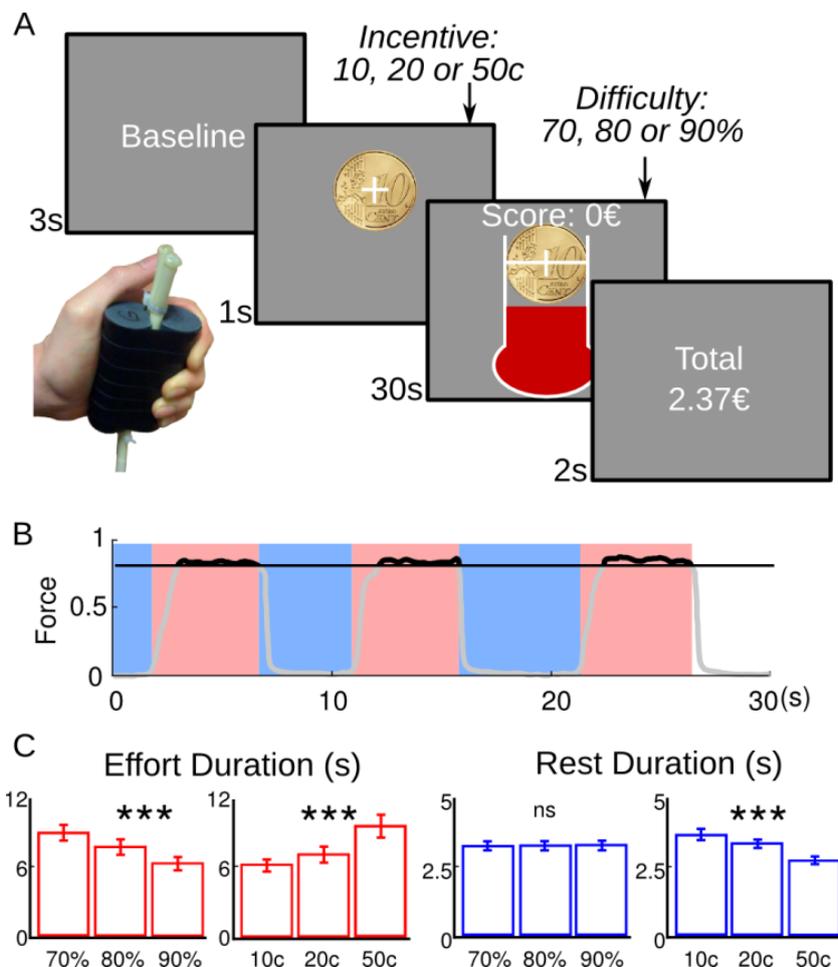


Figure 21 *Implicit task and results.*

A: The illustrated screenshots were successively presented in every trial. When the thermometer image was displayed, participants could squeeze a handgrip in order to win money. Subjects were provided with online feedback on force level and

*cumulative payoff. The payoff was only increased when force level was above the target bar, at a constant rate proportional to the monetary incentive. Two factors were manipulated over trials: the incentive (10, 20 or 50 cents), which was explicitly indicated as a coin image, and the difficulty, i.e. the force required to reach the target bar (70, 80 or 90% of maximal force), which remained implicit. The last screen indicated the money won so far, summed over all preceding trials.*

*B: Example recording of the force level produced during one trial. Three rest (blue shading) and effort (red shading) epochs could be defined. Force production was only rewarded when above the target threshold (here, 80% of maximal force), i.e. when plotted in black (not gray) on the graph.*

*C: Average data sorted by incentive and difficulty levels. Bars are mean effort and rest epoch durations and error bars the inter-subject standard errors. Significance of repeated-measure ANOVA main effects: \*\*\*  $p < 0.0005$ , \*\*  $p < 0.005$ , \*  $p < 0.05$ .*

The absence of effect of the effort difficulty on rest duration is surprising. If we take a simple economic perspective on the task, there is an opportunity cost of rest: not benefiting from the effort utility. Indeed, the time allotted for the task is limited; therefore increasing rest duration leaves less room for effort, through which reward is earned. The effort utility increases with the incentive (benefit) and decreases with the difficulty (cost). The opportunity cost therefore increases with the benefit and decreases with the difficulty. Participants limited this opportunity cost by reducing rest duration for higher incentives, why did not they take into account the difficulty level, which also determines the opportunity cost? To know whether this absence of effect is economically optimal, we would need to know the actual opportunity cost, hence the dynamic of cost during effort and rest. We do not. Nonetheless, incentive and difficulty have opposite effects on effort duration but not on rest; this asymmetry is at least surprising.

Two other versions of the effort allocation task are now introduced to better understand why rest durations are not modulated by the effort difficulty, although it could be strategic. To allow an easy and summarized comparison of the effect size across the tasks, the effort and rest duration were analyzed with separate linear regression models. Rest duration was shorter for higher incentives ( $p = 2.0 \cdot 10^{-5}$ ) but was not modulated by the experienced difficulty level ( $p = 0.32$ ). Effort duration was both longer for higher incentives ( $p = 8.1 \cdot 10^{-7}$ ) and shorter for higher experienced difficulty levels ( $p = 1.6 \cdot 10^{-10}$ ) (see Figure 23A-B). Interactions were included in the model, but incentive x difficulty interactions were not significant for neither effort nor rest durations (all  $p > 0.084$ ).

#### 3.1.2.2.2 Explicit effort allocation task

A possible explanation why the effort difficulty affected effort duration, but not rest is the following. The effort difficulty is a motor parameter that impacts the motor system, for instance by increasing the level of fatigue more steeply when it is more difficult. If the decision to stop is made when a given fatigue level is reached, all efforts should end with an equal fatigue level, which occurs at shorter latencies when the effort is more difficult. It is therefore possible that the effort difficulty

impacts effort duration, but that the estimate of the difficulty is not integrated (because not available, or not salient enough) in the strategic regulation of rest duration. On the contrary, the incentive level that is presented clearly in the task, impacts rest duration. A simple prediction of this explanation is that, if the effort difficulty is given explicitly like the incentive, it could impact rest durations.

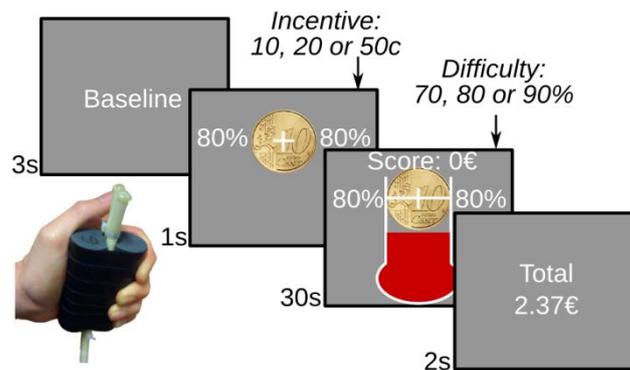


Figure 22 *The explicit task.*

*The only variation introduced compared to the Implicit task was that the effort difficulty level was indicated on the screen.*

To test whether making the difficulty level explicit for the subjects would change their effort allocation behavior, an explicit effort allocation task was designed. This task was exactly like the implicit task, except that the difficulty level was indicated on the screen as a percentage of the maximal force. The true, experienced difficulty and the cued difficulty levels were fully congruent (see Figure 22).

As in the Implicit Task, effort durations were both longer for higher incentives ( $p = 1.1 \cdot 10^{-3}$ ) and shorter for higher difficulty levels ( $p = 6.0 \cdot 10^{-6}$ ). Rest durations were also shorter for higher incentives ( $p = 9.7 \cdot 10^{-4}$ ). Contrary to difficulty levels in the Implicit task that were experienced but not cued and had no effect on rest durations, in the Explicit task, higher difficulty levels were both experienced and cued, and they increased rest durations ( $p = 1.6 \cdot 10^{-3}$ ). The difference in standardized effect sizes between the tasks was significant ( $p = 1.2 \cdot 10^{-4}$ ), see Figure 23C-D.

Interactions between incentive and difficulty in the explicit task were included in the model, but they were not significant for neither effort nor rest durations (all  $p > 0.1$ ).

The comparison of the explicit and implicit tasks suggests that cueing the difficulty level affects the behavior. It is as if the difficulty experienced during effort were not sufficient to impact rest duration, which is achieved only with further information. It suggests that the experienced difficulty and the expected difficulty could play distinct roles on effort and rest duration. To test this idea, these two aspects of the effort difficulty were dissociated in the last effort allocation task.

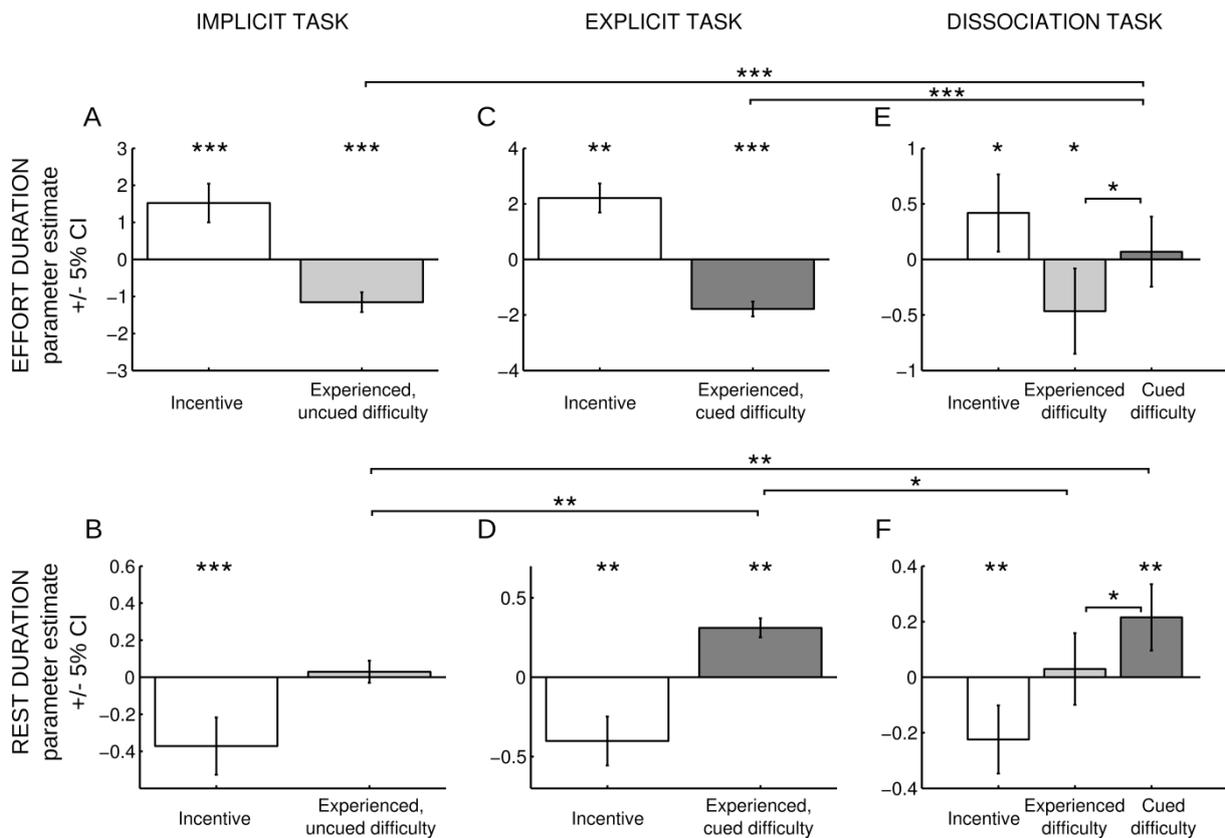


Figure 23 Summary of the three effort allocation tasks.

Three sets of participants performed three slightly different versions of the effort allocation task. The paradigm of the Implicit task is presented in Figure 21A. The only variation introduced in the Explicit task was that the effort difficulty was written on the screen (70%, 80%, 90%) along with the incentive level, revealed as a coin image. The Dissociation task looked like the Explicit task, the design however was different: the difficulty level reported on the screen was actually not predictive of the true difficulty level and both variables were crossed into a factorial design. Results of linear regression analysis are presented column-wise for each task. Coefficient weights were tested with bilateral t-test, p-values: \* < 0.05, \*\* < 0.005, \*\*\* < 0.0005.

### 3.1.2.2.3 Dissociation task effort allocation task

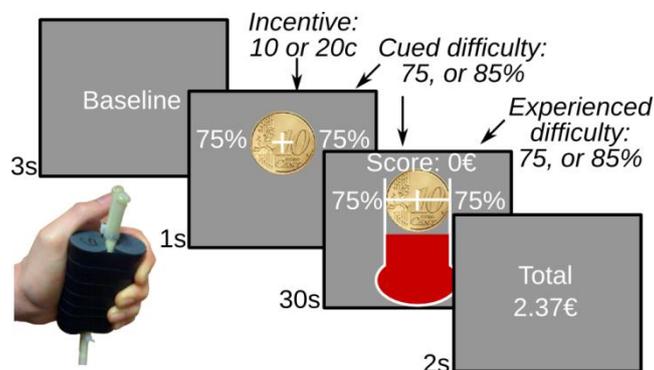


Figure 24 The Dissociation task.

*Compared to the explicit task, only two levels were used for difficulty and incentive, and the difficulty factor was split into two independent components: the cued difficulty level that was indicated on the screen and the true, experienced difficulty level that was manipulated independently from the cue.*

In this task, the levels of experienced and cued difficulty were manipulated independently. The incentive levels, the cued difficulty levels and the experienced difficulty levels were fully crossed in a factorial design. To achieve a reasonable amount of replicates, the amount of levels was reduced to two (see Figure 24).

As in the Implicit and Explicit tasks, higher incentive levels increased effort durations ( $p = 0.022$ ) and shortened rest durations ( $p=1.5 \cdot 10^{-3}$ ). Rest durations were affected by the cued difficulty level ( $p = 1.7 \cdot 10^{-3}$ ) but not by the experienced difficulty level ( $p = 0.63$ ). The difference in effect size was significant ( $p = 0.045$ ) (see Figure 23F). We also compared effects on rest durations between tasks: the comparison of the cued difficulty effect in the Dissociation vs. the difficulty effect in the Implicit task ( $p=0.002$ ), and the experienced effect in the Dissociation task vs. the difficulty effect in the Explicit task ( $p=0.008$ ) both revealed significant differences. Within-task and between-task comparisons both support dissociation between the cued and experienced difficulty levels on rest durations. Effort durations were affected by the experienced difficulty level ( $p = 0.021$ ) but not by the explicit difficulty level ( $p = 0.64$ ). The difference in effect size was significant ( $p = 0.050$ ) (see Figure 23E). We also compared effects on effort durations between tasks: the comparison of the cued difficulty effect in the Dissociation vs. the difficulty effect in the Implicit task ( $p=4.3 \cdot 10^{-7}$ ), and the Explicit task ( $p=2.3 \cdot 10^{-6}$ ) both revealed significant differences. Within-task and between-task comparisons support dissociation between the cued and the experienced difficulty levels on effort durations.

As the dissociation p-values are near the 0.05 type I error rate, we conducted a permutation test to ensure the reliability of the parametric t-distribution in our small sample. The permutation-based t-distribution actually yielded the same result for p-values up to the 3<sup>rd</sup> decimal.

Second and third order interactions terms between the incentive level, the cued and experienced difficulty level were included in the model, but for none was significant for neither rest nor effort durations (all  $p>0.18$ ).

We also checked that there was no interaction between the cued difficulty level effects and time, which could potentially reflect that the cue effect (with does not predict the true difficulty) is progressively ignored. Time was modeled at three nested time scales (rest or effort period position

within a trial, trial position within a session, and session number) and two-way interactions with the cued difficulty levels estimated. None was significant (all  $p > 0.25$ ).

The shortening of rest duration was economically strategic in this task because effort was rewarded. Decreasing rest durations left more room for effort. Indeed, there were more effort periods in the low than in the high cued difficulty condition (paired-difference  $5 \pm 2.4$  s.e.m., Wilcoxon  $p = 0.017$ ) and overall, shorter effort durations cumulated per trial for higher cued difficulty levels (beta =  $-0.29 \pm 0.13$  s.e.m., bilateral t-test  $p = 0.038$ , same regression analysis as for effort and rest duration).

### **3.1.2.3 Discussion**

To summarize, the tasks designed and presented here make it possible to operationalize and test the notion of effort allocation in the laboratory, at the time scale of a trial (30s). Thanks to this short time scale, factors like the incentive level and the effort difficulty can be varied from trial to trial. The main finding is that the effort allocation process is adaptive: more effort is produced when the utility is higher. More precisely, there are two leverages to increase effort production in these tasks: lengthening effort periods or shortening rest periods. The incentive had a global effect across the three tasks: it lengthened effort period and shortened rest period. The effort difficulty had a more subtle effect. Without explicit information on the difficulty level, the effort difficulty affected only the effort periods that were shorter when more difficult. This effect left rest durations unaffected. The effort allocation was incompletely optimized due to this absence of effect because increasing rest duration when the effort is more difficult would further improve the utility. This effect was recovered by introducing explicit information on the difficulty level. The dissociation between the experienced difficulty that limits effort production and the expected difficulty that subserves the optimization of rest durations was demonstrated by manipulating these factors independently.

Therefore, the incentive affected both effort and rest duration and the experienced and expected difficulty affected respectively effort and rest duration. These results refine our understanding of the effort allocation process. They suggest that the effort allocation behavior is strategic: it aims at maximizing benefit against cost, but they also suggest that there might be different levels of monitoring. There would be a low level through which experienced difficulty affects motor production and a more strategic level through which the incentive affects both effort and rest and the expected difficulty affects rest durations. These levels could correspond to different kinds of constraints the effort allocation is adapted to: utility maximization and biomechanics. These results also suggest that the behavioral monitoring is not fully planned beforehand based on a cost benefit ratio, but leaves room for on-line adaptation to the experienced cost.

This set of results answers some issues that were left open in the general introduction to this dissertation (see page 92). The primary determinant of effort allocation may be related to the dynamics of cost. Second, on top of this experienced cost dynamics, the determinants of the utility such as the money at stake or the expected difficulty would further modulate the effort allocation. The computational view of this process, i.e. how the factors presented here (experienced effort difficulty, expected effort difficulty, incentive) affect the parameters of the cost-evidence model could further refine the understanding of this mechanism.

### **3.1.2.4 Methods**

#### **3.1.2.4.1 Set up**

We used homemade power grips composed of two plastic or wood cylinders compressing an air tube when squeezed. The tube was connected to a transducer converting air pressure into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to the force exerted. The signal was amplified and digitized by a signal conditioner (CED 1401, Cambridge electronic design, UK) read by a Matlab program (The MathWorks Inc., USA).

#### **3.1.2.4.2 Pre-processing**

Effort onsets and offsets were identified off-line with an algorithm using the same two criteria for all conditions: 1) force temporal derivative higher than one standard deviation and 2) force level lower (for effort onset) or higher (for effort offset) than half the maximal force. The first rest period started with coin presentation and the subsequent effort and rest periods were defined by force onsets and offsets.

#### **3.1.2.4.3 Maximal force estimates**

For all tasks, we measured the maximal force for each hand before starting the task, following published guidelines (Gandevia, 2001).

Participants were verbally encouraged to squeeze continuously as hard as they could, until a growing line displayed on a computer screen reached a target. The growing rate was proportional to the force produced to motivate subjects to squeeze hard. Maximal force was set to the average of data points over the last half of the squeezing period exceeding the median. Then subjects were provided a real-time feedback about the force produced on the handgrip, which appeared as a fluid level moving up and down within a thermometer, the maximal force being indicated as a horizontal bar at the top. Subjects were asked to try outreaching the bar and state whether it truly corresponded to their maximal force. If not, the calibration procedure was repeated.

#### 3.1.2.4.4 Behavioral tasks

All tasks were presented on a computer screen, and were programmed with Matlab using Cogent 2000 (Wellcome Department of Imaging Neuroscience, London, UK) for the Implicit and Explicit task, and using Psychtoolbox (<http://psychtoolbox.org>) for the Dissociation task.

**The Implicit task.** Task sessions included 9 trials corresponding to the 9 cells of the factorial design (3 incentive x 3 difficulty conditions), which were presented in a random order. Subjects performed 8 sessions in total, switching hands as instructed between sessions to avoid muscle exhaustion. After baseline measurement of the pressure at rest, each trial started by revealing the monetary incentive with a coin image (10, 20 or 50 cents) displayed for 1s. Then subjects had 30s to win as much money as possible. They knew that the payoff was proportional to both the incentive and the time spent above the target bar, which was always placed at the same height in the thermometer. The force needed to reach the bar (70, 80 or 90% of subject's maximal force), i.e. trial difficulty, was not indicated to subjects. Subjects only knew that task difficulty would vary across trials. They were provided with online feedback on both the exerted force (with a fluid level moving up and down within a thermometer) and the trial-wise cumulated payoff (with a counter displayed above the thermometer) The fluid had the same luminance as the background to avoid confounding force level with basic visual features. Each trial ended with a 2s display of the session-wise cumulated payoff.

The only difference from the Implicit to the **Explicit task** is that the difficulty level was displayed on the right and left on the coin, as percentages (70%, 80%, 90%).

In the **Dissociation task**, incentives (10c, 20c), cued difficulty levels (75%, 85%) and experienced difficulty levels (75%, 85%) were associated into a factorial design comprising 8 cells. Cued difficulty levels were indicated on the screen as in the Explicit task and were congruent with the experienced difficulty levels (the difficulty level actually used in the trial) only in half on the trials. The experiment was divided into 8 sessions of 8 trials each presenting all conditions in a randomized order. The randomization avoided the presentation of identical cues (incentive level and cued difficulty level) in two consecutive trials. Apart from the potential mismatch between the cued and experienced difficulty levels, the trial settings were identical to those of the explicit task.

#### 3.1.2.4.5 Participants

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. All subjects were recruited via email within an academic database and gave informed consent prior to participating in the study. Participants were restricted to right-handed only for the Implicit task, with age between 20 and 39 years old, no self-reported psychiatric or neurological history and no self-reported current

psycho-active substance consumption. In all tasks, but the Implicit task, participants were paid in cash at the end of the experiment. The payoff was partitioned into a fixed amount and additional amount that corresponded to the money won during the task (respectively ‘Fixed’ and ‘Var’ in the table). The task-specific information is as follows. For the Implicit task, half of the subjects performed the experiment in an MRI scanner, the other half under a MEG helmet. One subject in the MRI group was excluded from the analysis because of calibration issues. The payoff was eventually rounded up to a fixed amount (100€) credited by bank transfer. For the Dissociation task, one participant was excluded due to an instruction issue: she could not understand the meaning of the percentage displayed on the screen, which represented the difficulty levels. Two other participants were excluded due to calibration issues. ‘Exp.’ refers to the author who collected the data (FM: Florent Meyniel, LS: Lou Safra)

Task	Exp.	Period	N included	N male	N excluded	Mean age +/- s.e.m.	Fixed (€)	Var. (€)	Var. (€) range
Implicit Task	FM	04-05/2010	38	16	1	24.2 +/- 0.65	50	31.6	15 – 48
Explicit Task	FM	03/2011	14	10	0	23.7 +/- 0.4	15	13	8 – 19
Dissociation Task	LS	10/2011	15	5	3	25.4 +/- 0.8	10	6	3 – 10

#### 3.1.2.4.6 Statistical analysis

The effort and rest durations were submitted to a multiple regression analysis. The explanatory variables of the linear model comprised the manipulated factors (incentive level, difficulty level, for the implicit and explicit task; incentive level, cued difficulty level and experienced difficulty level for the dissociation task), temporal factors (the session number, the trial position within a session and the effort or rest position within a trial), and interaction terms (the two-way interactions between the manipulated factors and the temporal factors, and the two-way interaction between the manipulated factors, which was extended to a third-way interaction between the three manipulated factors of the dissociation tasks). All the explanatory variables were z-scored for the regression.

The significance of the parameter estimates was assessed using a random effect at the group level with a two-tailed t-test. Dissociation between cued and experienced difficulty levels in the dissociation task was estimated using a two-tailed paired t-test on the parameter estimates. For the non-parametric t-test, we estimated the null t-distribution using all possible permutations ( $n=2^{15}$ )

between the cued and experienced labels, and estimated the p-value for a more extreme statistics (bilateral test).

The results of the ANOVA reported for the Implicit task were derived as follows. Effort and rest period durations were separately analyzed using a repeated-measure ANOVA (R software, John Fox CAR library), with incentive and difficulty as factors of interest. The p-values reported for these repeated-measure ANOVA integrate the conservative Greenhouse-Geisser correction.

### **3.1.3 Computational account of effort allocation**

#### **3.1.3.1 Introduction**

So far, the cost-evidence model of effort allocation was presented for its basic properties: effort and rest are subjected to opposite accumulation dynamics and high (respectively low) cost-evidence levels trigger the decision to stop the effort and to resume it. The first set of studies revealed that the durations of effort and rest constrain the subsequent effort or rest durations, which in terms of the accumulation model, can be rephrased as: the manipulation of effort and rest shifts the cost-evidence toward the upper or lower bound, which constrains the subsequent effort duration within which the upper bound is met, or conversely, the lower bound for rest. The second set of studies revealed that the effort difficulty and the monetary benefit associated to the effort impacted behavior. From a computational perspective, can this effect be pinpointed onto specific model parameters? This computational account of the factors would refine our interpretation of the mechanism underpinning the effort allocation. In particular, it would refine our understanding of the bounds and the slopes of the model.

#### **3.1.3.2 Results**

The accumulation model presented in the general introduction (see p 92) relates the level of the accumulated cost-evidence level to the decision to stop the effort or the resume it. As a default assumption, the accumulation and dissipation of cost evidence are linear, so that the effort duration is  $Te = A/Se$  and rest duration is  $Tr = A/Sr$ , where  $A$  is the distance between the 'stop' and 'go' bound,  $Se$  the accumulation slope during effort and  $Sr$  the accumulation slope during rest. A graphical representation of the model aligned on an example trial is shown in Figure 25.

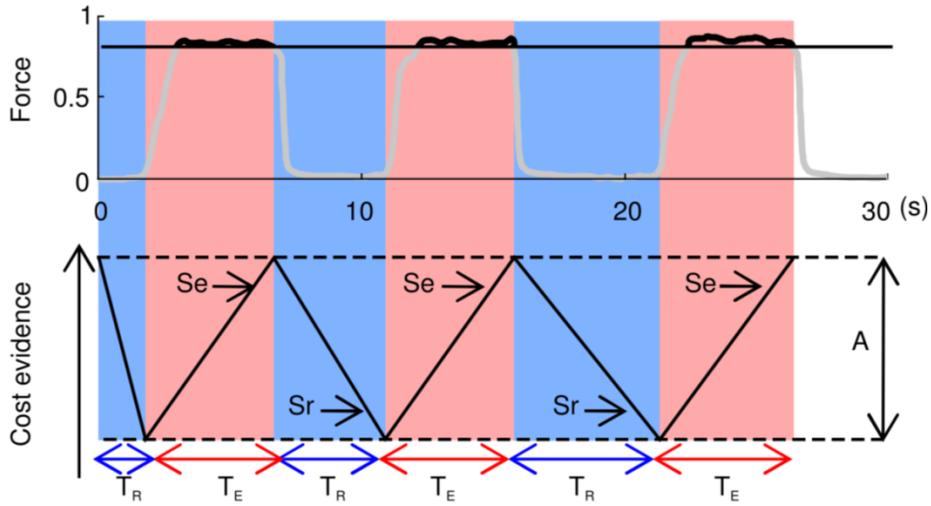


Figure 25 Cost-evidence accumulation model for effort allocation.

Example recording of the force exerted during one trial. Three effort (red shading) and three rest (blue shading) epochs could be defined. Force level is shown in black (not gray) when rewarded, i.e. when above the target level (here, 80% of maximal force). The theoretical cost evidence that was hypothesized to underpin effort production is shown below. A: amplitude between bounds; Sr: dissipation slope during rest; Se: accumulation slope during effort; Te: effort time; Tr: rest time.

How are the model parameters (the accumulation slope  $Se$ , the dissipation slope  $Sr$  and the amplitude between the bounds  $A$ ) modulated by the experimental factors? In the implicit task, there are two factors: the incentive level  $I$  and the difficulty level  $D$ . For each parameter  $X$ , we assumed that variations around the mean value  $X_m$  could be explained by a weighted sum of the incentive  $I$  and the difficulty  $D$  (all factors are z-scored), with respective weight  $X_I, X_D$ , so that the model is:

$$Te = \frac{A}{Se}; Tr = \frac{A}{Sr}; \text{with: } \begin{cases} A = A_m + A_I I + A_D D \\ Se = Se_m + Se_I I + Se_D D \\ Sr = Sr_m + Sr_I I + Sr_D D \end{cases} \quad (1)$$

In principle, all parameters could be modulated by all factors. However, it is possible that simpler models with fewer modulations could still well account for the data. Simpler models can be derived from the previous equation by excluding some modulations. There are three parameters each of which can be modulated or not by two factors, resulting in  $4 \times 4 \times 4 = 64$  models. Among those models, some are a priori not worth considering because they cannot account for the behavioral effect identified in the task, so that the space of the model search was restricted to plausible models, i.e. models that can account for the observed behavior (see methods and Table 2). The restricted search space comprised 24 models. These models differ in their degrees of freedom and in how well they can fit data. We identified the best tradeoff between simplicity and fit with a Bayesian model selection. The model with the higher expected frequency had an exceedence probability  $x_p = 0.90$ . In other words, there is 90% confidence that this model is more frequent than the other models of the set (i.e. the 23 other models) in the general population based on the data observed in our group of

subjects (see page 108 for details on Bayesian model selection and statistics). This model is the following:

$$Te = \frac{A}{Se}; Tr = \frac{A}{Sr}; \text{with: } \begin{cases} A = A_m + A_I I \\ Se = Se_m + Se_D D \\ Sr = Sr_m + Sr_I I \end{cases} \quad (2)$$

The model selection and the model fit suggest the following computational effects in the implicit task: higher incentive levels increase the distance between the bounds and steepen the cost-evidence dissipation during rest; higher effort difficulty levels steepen the cost-evidence accumulation during effort.

In the Dissociation task, there is an additional parameter since the factor difficulty is split into a cue difficulty level  $Dc$  and an experienced difficulty level  $De$ . The model equations (1) can be extended in this case:

$$Te = \frac{A}{Se}; Tr = \frac{A}{Sr}; \text{with: } \begin{cases} A = A_m + A_I I + A_{De} De + A_{Dc} Dc \\ Se = Se_m + Se_I I + Se_{De} De + Se_{Dc} Dc \\ Sr = Sr_m + Sr_I I + Sr_{De} De + Sr_{Dc} Dc \end{cases} \quad (3)$$

Note that in the explicit task, the cued difficult and the experience difficulty have the same values, so that the weights of these two factors cannot be estimated in a linear model. The model equations for the explicit task are therefore like in (1).

As done previously for the implicit task, we tried to identify whether a simpler model would provide a good account of the data. Simpler models were derived from the full model that comprises all modulations of every factor, and the search space was restricted to plausible model a priori. 144 linear models were identified for the Dissociation Task and 16 for the Explicit Task (see Table 3 and Table 4). Yet, these models resulted in poor fit and unclear discrimination for a best model in the Bayesian model selection. We reasoned that the effect of cued difficulty levels might be better captured by hyperbolic rather than linear discount of the incentive as it is frequently modeled in the literature for delay discounting (Peters and Büchel, 2011) but also for effort (Prévost et al., 2010). This ‘hyperbolic’ model, in the full formulation, is:

$$Te = \frac{A}{Se}; Tr = \frac{A}{Sr}; \text{with: } \begin{cases} A = A_m + \frac{A_I I}{1 + A_{Dc} Dc} + A_{De} De \\ Se = Se_m + \frac{Se_I I}{1 + Se_{Dc} Dc} + Se_{De} De \\ Sr = Sr_m + \frac{Sr_I I}{1 + Sr_{Dc} Dc} + Sr_{De} De \end{cases} \quad (4)$$

We identified the hyperbolic counterparts of our linear models, which resulted in a search space comprising 78 additional models for the Dissociation and Explicit tasks (see Methods, Table 5). Note that there is a one to one mapping between the hyperbolic models of the Explicit and Dissociation

tasks despite cued and experienced difficulties are confounded in the Explicit Task. Indeed, there is no redundancy in the hyperbolic case, as opposed to the linear case in which two identical regressors cannot be estimated unambiguously.

Comparing families of models revealed that there was far more evidence in favor of a hyperbolic rather than linear discount of incentives by the cued difficulty levels in both the Explicit and Dissociation tasks (all  $x_p > 0.999$ ). Among the 78 hyperbolic models, a best model was identified with  $x_p = 0.90$  in the dissociation task and with  $x_p = 0.82$  in the explicit task.

Crucially, the best hyperbolic model identified independently in the Explicit and Dissociation tasks was the same model:

$$Te = \frac{A}{se}; Tr = \frac{A}{sr}; \text{with: } \begin{cases} A = A_m + A_I I \\ Se = Se_m + Se_{De} De \\ Sr = Sr_m + \frac{Sr_I I}{1 + Sr_{Dc} Dc} \end{cases} \quad (5)$$

This model also corresponds to the best model identified in the Implicit task as there is a correspondence between the two when considering that the cued difficulty factor is null in the Implicit task.

Overall, the Bayesian model selection enabled to pinpoint the computational effect of the factors manipulated in the three tasks. A graphical summary of these effects in the evidence-based accumulation model for effort allocation is presented in Figure 26.

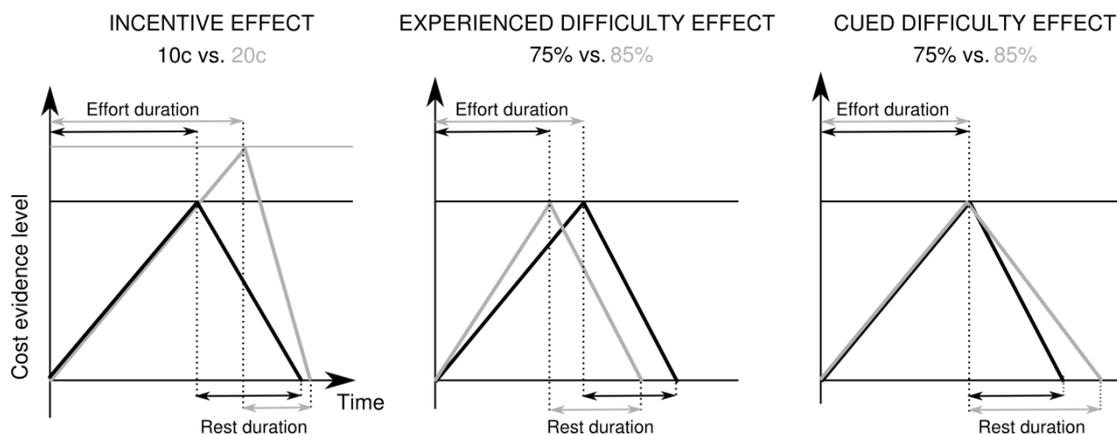


Figure 26 Cost-evidence accumulation model: summary of the factor effects.

The diagrams illustrate how the experimental factors (monetary incentive, experienced difficulty, cued difficulty) affect the accumulation and dissipation of cost evidence. Bayesian model comparison reveals that the computational effects of these factors are distinct: higher incentives increase the amplitude between bounds and the dissipation rate of cost evidence; higher experienced-difficulty levels steepen the accumulation rate of cost evidence; higher anticipated-difficulty levels shallow the dissipation rate of cost evidence.

### 3.1.3.3 Discussion

We can summarize the computational effects of the task factors (see Figure 26) as follows: the incentives pushed the limit of cost-evidence back, the experienced difficulty steepened cost-evidence accumulation during the effort, and the dissipation rate of cost-evidence was modulated by the effort expected utility: the rate was steeper for higher incentive levels and more shallow for higher cued difficulty levels.

During effort there are two components modulating the dynamic of cost-evidence: an accumulation modulated by the difficulty level and a bound modulated by the incentives. This dynamic is similar to what we expect from the dynamics of the physiological cost and the monetary benefit. By design, the dynamic of benefit is stationary over time: each time unit of effort brings the same reward in a given trial, there is a fixed reward rate that depends on the incentive level. By contrast, the dynamic of cost is more complex: due to fatigue, the cost entailed by producing a time unit of effort increases with effort duration. This increase is steeper for higher difficulty levels because producing more intense effort increases fatigue more rapidly. Given the similarity between the fitted dynamic of cost-evidence and the cost-benefit economic perspective, it is tempting to interpret the accumulation process as maximization of effort utility. The utility cumulated during effort corresponds to the difference between the cumulated benefit and cost. These cumulated values are related to their instantaneous counterpart: the cumulated value is the integration over instantaneous values, and conversely, the instantaneous value is the derivative of the cumulated value (see Figure 27). The optimal effort duration is the one for which the cumulated utility is maximal. This maximization problem can be solved on-line by monitoring at each time step the instantaneous cost and benefit: see Figure 27, the two curves meet at the optimal effort duration. This interpretation is compelling because first, it corresponds to the intuition that there is on-line monitoring during effort allocation and second, it is economically sensible.

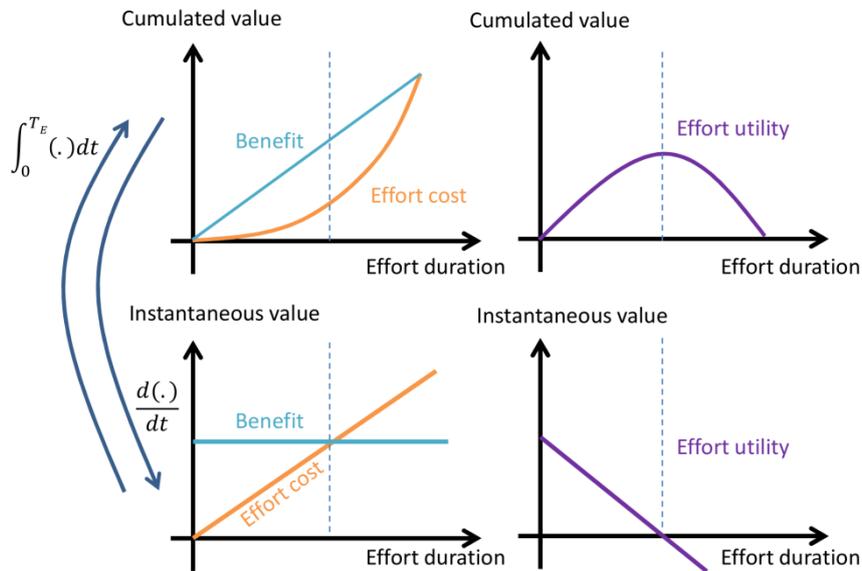


Figure 27 Economic view of the effort.

However, this parallel between the economic view and the cost-evidence accumulation model rapidly comes to an end. First, if participants decide to do the effort or not based on the instantaneous utility value, it is hard to explain why they do not resume effort immediately after cessation. Indeed, given that the effort cost is cleared during rest, the instantaneous effort utility would return to positive values shortly after the effort cessation. Instead, participants took substantial rest durations. Second, what should be maximized is not the cumulated utility of single efforts, but the cumulated utility over efforts and rests across the entire experiment. This long term strategy implies that efforts should not be too short, because longer efforts can bring more utility; and efforts should neither be too long because it could impair the ability to perform efforts in the future (and hence, to accumulate benefit). Accordingly, rests should not be too short, to prevent that efforts are produced only in a high cost range; but they should neither be too long because longer rests leave less time for efforts, which are rewarded. To solve these tradeoffs, the effort and rest durations should be optimized relatively to a utility function encompassing the whole experiment. It is difficult for us, as scientists, to solve this problem because we do not know exactly the dynamics of cost during effort and rest and how they are affected by the difficulty level. There is in particular a lack of information on the dynamic of recovery. It is also difficult for the brain to solve this problem, because it would imply that 1) the brain can solve the utility maximization problem, 2) that she knows cost functions for all kinds of effort the body can produce and 3) that she has all the relevant information beforehand, which is not the case when the difficulty level is implicit for instance.

Can we find an easy workaround? There are two lessons in the above analysis: first, it is sensible to have on-line monitoring, because it corresponds to intuitions we have in this problem and because it

enables to take into account unpredictable dynamics; second, it is also sensible to optimize the behavior on the long run, and not blindly at each time step. It seems that monitoring the instantaneous cost and setting an upper bound on the accumulation during exertion and a lower bound on the dissipation during rest has these advantages: costs are then limited during effort and adapted to the benefit through the upper bound; they are also limited through the lower bound that controls that effort is not confined to high cost range and that the opportunity cost of rest is also limited.

There are several plausible candidates cost signals for the brain to implement such a heuristic. From a body perspective, such signals could be seen as predicting the risk of damage in muscles and other organs ensued by prolonging the exercise or ensued by initiating an effort during rest. The accumulation-to-bound mechanism during effort would protect the system against harm and the dissipation-to-bound mechanism during rest would be like a refractory period to ensure that the body is sufficiently prepared for new efforts. During exertion, this signal would predict potential damage based on the cost currently experienced. During rest, there is less information so that this signal could rely more on expectations. This interpretation is an attempt to make sense of the dissociated effects of expected and experienced difficulty levels between effort and rest. The effect of the incentive on the dissipation could have two meanings: first, that there is an active recovery modulated by the incentive so that the risk of damage would truly decrease faster; second, that this alarm signal relies more on expectation during rest, which includes both the expected difficulty and the motivation to resume the effort shortly.

### 3.1.3.4 Methods

#### 3.1.3.4.1 Specification of the model space

The aim of the model space definition is to define a class of models that can a priori produce the results we are interested in explaining. These models were then submitted to a Bayesian model selection to identify the best model among all the plausible models. The model space was defined by simplifying a full model. The implicit task can serve here as an example, the specificity of the linear and hyperbolic models for the explicit and dissociation tasks are pointed later on. The accumulation is linear; the effort duration  $Te$  (respectively  $Tr$  for rest duration) is thus the ratio between the amplitude  $A$  between the bounds and the accumulation slope  $Se$  (respectively  $Sr$  during rest). In the full model, the accumulation variables  $A$ ,  $Se$ ,  $Sr$  have mean values  $(A_m, Se_m, Sr_m)$  and variations around this mean value are modeled linearly across trials by the incentive level  $I$  and the difficulty level  $D$ , with weights  $A_I, Se_I, Sr_I, A_D, Se_D, Sr_D$ , see equation (1).

Simpler models were designed by considering that one or more weights were null. There are 6 weights that can be set to 0 or not, resulting in  $2^6 = 64$  models. However, some of these models are not worth considering as they cannot account for the effect we want to explain. The most extreme case is when all weights are null, such a model cannot produce any modulation by the incentive and difficulty levels observed in the data. We restricted the model space to 24 models; models were excluded when they could not produce at least one of the significant results reported in Figure 21. Note that predicting a potential effect whereas no significant effect was found in the data was not a criterion for rejection. The models are presented in Table 2.

The same scheme was repeated for linear models of the explicit task, leading to a search space of 16 models. For the dissociation task, another modulator was included as there were two types (rather than a unique type) of difficulty. The full model has 9 weights and the full model space thus comprises  $2^9=512$  models. The search space, accordingly, was reduced to 144 models.

We also tested a class of hyperbolic model for the explicit and dissociation task (see equation (4)). As opposed to the linear formulation, the discount of the incentive by the cued difficulty level is hyperbolic. There was a  $D$  term, denoting difficulty in the implicit and explicit tasks,  $De$  and  $Dc$  denote the experienced and cued difficulty in the dissociation task. For the explicit task,  $De$  and  $Dc$  have exactly the same values; the model can nonetheless be estimated unambiguously when the effects of  $De$  and  $Dc$  are not linear, like in the hyperbolic formulation. With the hyperbolic forms, there are dependencies between weights since a null numerator prevents the denominator from impacting the model. Thus we discarded models with a null numerator and a non-null weight at the denominator (shading in Table 5). 78 models were eventually included in the analysis after reducing the model space with the same restrictions as before.

#### 3.1.3.4.2 List of models

Models are characterized by the modulation of parameters  $A$ ,  $S_E$ ,  $S_R$  by the factors (Difficulty, experienced or cued; Incentive). The list of possible modulations should be read column-wise: '1' denotes that the modulation is allowed, '0' that the modulation is absent. Modulations for a given factor are combined with modulations in the other factors. This means that any row from the Difficulty column can be combined to any row from the Incentive column, which characterizes a given model. Red highlight means that the modulations are not included in the model space and bold blue highlight indicates that the modulations in the best model. Note that there is no blue bold highlight in the tables of linear models for the Explicit task and the Dissociation task because these linear models were worse than their hyperbolic counterpart (last table), that were retained for the analysis.

Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1
1	1	0	1	1	0
1	0	1	1	0	1
1	0	0	1	0	0
0	1	1	0	1	1
0	1	0	0	1	0
0	0	1	0	0	1
0	0	0	0	0	0

Table 2: Implicit task, model space.

Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1
1	1	0	1	1	0
1	0	1	1	0	1
1	0	0	1	0	0
0	1	1	0	1	1
0	1	0	0	1	0
0	0	1	0	0	1
0	0	0	0	0	0

Table 3 Explicit task, model space for linear models

Experienced Difficulty			Cued Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1	1	1	1
1	1	0	1	1	0	1	1	0
1	0	1	1	0	1	1	0	1
1	0	0	1	0	0	1	0	0
0	1	1	0	1	1	0	1	1
0	1	0	0	1	0	0	1	0
0	0	1	0	0	1	0	0	1
0	0	0	0	0	0	0	0	0

Table 4 Dissociation task, model space for linear models

Experienced Difficulty			Cued Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1	1	1	1
1	1	0	1	1	0	1	1	0
1	0	1	1	0	1	1	0	1
1	0	0	1	0	0	1	0	0
0	1	1	0	1	1	0	1	1
0	1	0	0	1	0	0	1	0
0	0	1	0	0	1	0	0	1
0	0	0	0	0	0	0	0	0

Table 5 Dissociation & explicit task, model space for hyperbolic models

### 3.1.3.4.3 Bayesian model selection

See page 108.

## 3.1.4 Limitation of introspection

### 3.1.4.1 Introduction

The question addressed now is whether we can access through introspection the cost-evidence variable. The experienced and expected difficulty levels had distinct impact on the behavior, which suggest that there is something special about introspection in our task. It could be more strategic to modulate rest durations according to the experienced difficulty level; the inability to do so may suggest troubles in the introspection. Interestingly, the sport literature reported that athletes with good interoception abilities perform better. For instance, people who estimate their own heart beat rate accurately perform better on ergometer cycling, which could be attributed to more precise tuning of the behavioral self-regulation of the work load (Herbert et al., 2007). Another finding is that when the anticipated task difficulty is different from the experienced difficulty in running, athletes update their estimation after several minutes (Tucker, 2009). Although this timescale is much longer than that of trials in our tasks, this suggests that the introspection could play a critical role to improve physical performance.

The cost-evidence model posits that cost increases during the effort at a given accumulation rate, so that the cost-evidence level increases as the product of this rate and the effort duration. The model fit showed that this accumulation rate increases with the effort difficulty, so that the cost-evidence level corresponds to the product of the difficulty and the duration of effort. This makes the prediction that equivalent levels of cost-evidence cannot be achieved by substituting linearly

difficulty levels and durations; instead, the product makes the equivalence curves of cost-evidence levels convex relatively to pairs of difficulty levels and durations. For instance, the product between two numbers is equal to 9 over several pairs of values, like  $1*9$ ,  $3*3$  and  $9*1$ , which corresponds to a convex curve. First, we re-analyzed the behavioral results the Implicit Task (see Figure 21) to check whether cost, as inferred from behavior, is convex. Second, we run a separate experiment and ask participants to rate cost-evidence levels associated to effort for which the duration and difficulty were controlled experimentally to check whether the ratings would exhibit this convexity property.

### **3.1.4.2 Results**

To estimate cost-evidence levels from the behavior, we could fit our cost-evidence accumulation model. However, given the above results, we already know that it would satisfy the convexity property. It would be more convincing to show this property from the data independently of the model. We reasoned that when people are free to decide the effort duration, higher cost-evidence levels should make the probability to stop the effort more likely.

We first estimated from our data in the Implicit Task the probability to stop at a given effort duration for each difficulty level. For each participant, this probability is the cumulated frequency of effort durations observed over the task for a given difficulty level. In other word, the probability to stop at effort duration  $T$  is equal to the number of efforts for which the duration is at most  $T$ , divided by the total number of effort. We then compared predictions of this probability to stop from several models of cost-evidence, using a sigmoid link function to convert cost-evidence levels into a probability (value between 0 and 1). Models of cost-evidence differed in whether they allow non-linearity of factors, main factor effects and interaction between factors. Bayesian model comparison favored the pure interactive model (model expected frequency: 0.72, exceedance probability  $x_p=0.998$ , see Figure 28B).

A different set of subjects participated in the Cost Rating Task in which the duration and difficulty of efforts were imposed and varied experimentally. For comparability with the previous task, we also manipulated the incentive level (see Figure 28C). On each trial, participants won the coin value serving as incentives. More exactly, they won a portion of this value determined by the proportion of time spent squeezing at the required target force level or higher. Participants were asked to be as accurate as possible, so that this ratio was almost 100% (mean over subjects: 98.7%, extreme subjects: 94.6% and 99.9%). The difference between the force produced and the force required did not varied significantly across conditions (incentives,  $p = 0.28$ ; effort duration  $p=0.99$ ; difficulty level  $p=0.59$ ; interactions between these factors all  $p>0.21$ ), suggesting that effort production was well controlled by the design. The cost rating was framed in term of exhaustion, which is more intuitive

for participants: ‘Are you exhausted now?’, with ratings between the extremes ‘not at all’ and ‘completely’. Cost ratings were not impacted by incentives ( $p=0.1$ ), and marginally by the initial position of the cursor on the rating scale ( $p=0.056$ ). Critically, the cost rating increased both with the effort duration ( $p=0.028$ ) and the effort difficulty ( $p=5 \cdot 10^{-6}$ ), without significant interaction between these factors ( $p=0.96$ ). Bayesian model selection confirmed that there was no evidence for interaction, the best model was simply additive (expected frequency: 0.68, exceedance probability  $x_p=0.997$ ) (see Figure 28E). The fit of this model is shown in Figure 28F and should be compared to the convex profile of Figure 28C.

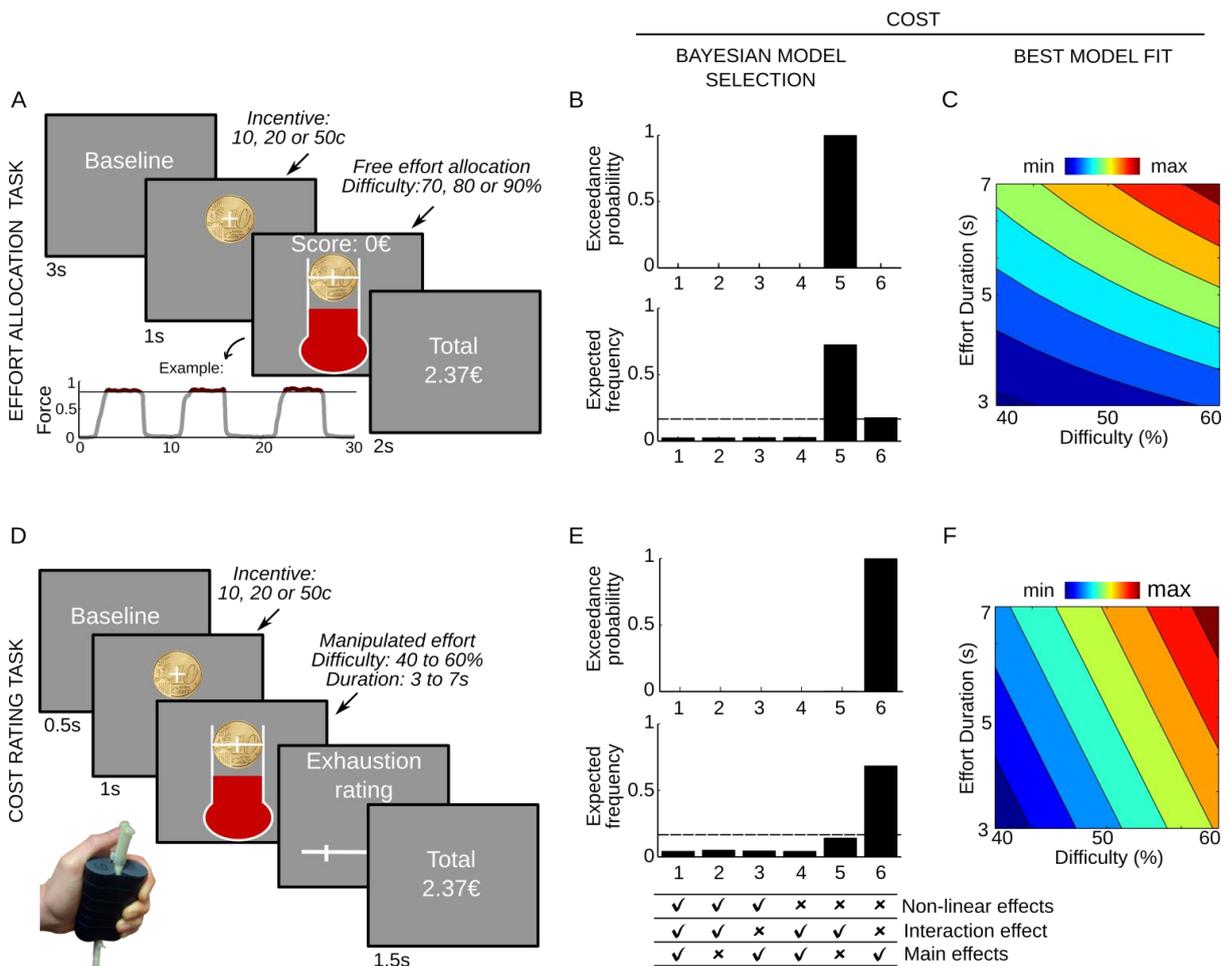


Figure 28 Introspection task and results.

A: Implicit effort allocation task. The illustrated screenshots were successively presented in every trial. When the thermometer image was displayed, participants could squeeze a handgrip to win as much money as possible. Subjects were provided with online feedback on force level and cumulative payoff. The payoff was only increased when force level was above the target bar, at a constant rate proportional to the monetary incentive. The incentive (10, 20 or 50 cents) and the difficulty (i.e. the force required to reach the target bar: 70, 80 or 90% of maximal force) were crossed over trials. The last screen indicated the money won in the trial and summed over all preceding trials.

B: For each participant, we estimated the cost level that triggers the decision to stop and fitted several models to predict cost from effort durations and difficulty levels and compared them. Models could include main effects, interaction and a

*non-linear modulation of factors, when marked with a tick, or not, when marked with a cross (bottom chart). C: The fit of the best model is shown. For this fit, we used the median parameter values over subjects and the same range of difficulty and effort duration levels as in (F) for visual comparison.*

*D: Cost Rating Task. On each trial, participants were asked to squeeze the hand grip up to the target level (horizontal bar, controlling for the difficulty level: 40% to 60% of maximal force) and as long as the thermometer was displayed (controlling for the duration: 4s to 7s). After each effort, participants rated the cost level associated to the effort with a visual horizontal analog scale. The last screen of each trial summarized the trial payoff and payoff cumulated over preceding trials (not shown here for clarity). Model comparison of cost ratings is shown in E. The fit of the best model is plotted with the median parameter values over subjects (F).*

To double-check the key difference between the experiments: linear vs. convex, we fitted a model with constant elasticity of substitution between effort duration and difficulty to characterize the curvature of costs (see methods). Costs in the Effort Allocation Task was convex, as reflected by curvature indices below one (median: 0.66, s.e.m.: 0.08, sign-test of the median against 1,  $p=4.7 \cdot 10^{-4}$ ), this was not the case in the Cost Rating Task (median: 1.18, s.e.m.: 0.09) and curvature indices significantly differed between tasks ( $p=2 \cdot 10^{-5}$ , Wilcoxon rank sum test for equal medians).

During the debriefing of the Cost Rating Task, participants unambiguously reported that variations of effort difficulty and effort durations were obvious. Then, they were asked whether one of these two factors had a greater impact on their ratings. 13 subjects reported that it was the duration, 3 that it was the difficulty, and 2 that there was an interaction. The comparison of the standardized effect size revealed a greater impact of the difficulty on ratings than the duration (paired t-test,  $p=0.016$ ). Among the 16 subjects who reported main effects, the comparison between guess and true effect was not correct for 12 of them, which is more than expected by chance (binomial test,  $p=0.028$ ).

### **3.1.4.3 Discussion**

To summarize the findings, cost-evidence levels that drive behavior and that are reported by subjects differ in their shape, suggesting that the true cost-evidence level is not what subjects report. Additionally, subjects have a systematic bias in what they think their report is impacted by. Overall, these findings suggest poor introspective capabilities on effort-related costs.

The absence of interaction between effort difficulty and duration is not due to these factors being ignored altogether, since the main effects were significant. Note that this conclusion is not based on accepting the null hypothesis (absence of interaction) in the classical statistical approach. We avoid this flaw by comparing the evidence of each hypothesis (or model), i.e. the probability of observing the data given a model that includes interaction or not. This absence of interaction should be contrasted with the findings that, in the effort allocation tasks, people adapt their effort duration to

this interaction, which is captured by the cost-evidence metric. Why is the reported cost different from the variable that underpins effort allocation?

A first explanation might be that when there is no need to optimize behavior, cost-evidence is not computed and therefore, not reported. Subjects, who were asked for a rating, actually report something else. To circumvent this issue, we could have asked ratings during the effort allocation task. However, this is methodologically difficult: the task should have been interrupted to ask the rating, which could alter the temporal structure of effort allocation problem. Another issue is that it is hard to know whether participants do what they tell or tell what they do: do subjects sample an internal variable that underpins their behavior (like cost-evidence) or do they adapt their behavior to match what they say? It could also be that they learn from what they do and try to make a report consistent with their behavior. Overall, it is likely that the effort allocation task and the introspection task would interfere.

Another explanation is that the cost-evidence variable is computed even when it is not used for behavior, and that this variable can be introspected, but that the way we framed our introspective question does not suit this variable. A last possibility is simply that the cost-evidence variable is computed, but cannot be introspected. These three explanations have in common that the exhaustion feeling that subjects reported in the cost-rating task cannot be used by the brain as a cost-evidence signal because it does not have the required properties.

The comparison of the implicit, explicit and dissociation tasks previously suggested that the true, experienced difficulty level was not used to optimize rest duration. Can this be related to the results of the cost-rating task?

From the effort allocation tasks alone, it seemed possible that the participants are not aware of the difficulty level, which prevents them from using this information to optimize rest duration. In the cost-evidence accumulation model, there is no contradiction in having both an effect of the difficulty on effort and no effect on rest: the comparison between the cost-evidence level and a threshold does not carry information about the difficulty level. This difficulty level can only be estimated from cost-evidence by considering its dynamics: steeper increases indicate higher difficulty levels. Therefore, it is possible that the decision to stop the effort is affected by the effort difficulty while the former remains blind to the latter.

Is the explanation that participants cannot estimate the difficulty level sensible given the results of the cost rating task? Apparently not since the reported exhaustion was impacted by the difficulty level. The difficulty range in the cost-rating task was 40 to 60% of the maximal force whereas it was

70 to 90% in the effort allocation task. Lower values were used because a cost rating task with a sufficient amount of trials for the analysis was simply not doable in a higher difficulty range: our pilot studies showed that it is way too exhausting. One could therefore argue that people can sense the different levels in the 40-60% range but not in the 70-90% range. I do not believe it is the case because a pilot study with lab mates (data not shown) revealed that when people are asked to discriminate between 70, 80 and 90%, their correct guess rate is around 2/3 and significantly higher than the chance level 1/3. However, the exhaustion ratings were also affected by the duration of the effort. It is possible that if participants rely on their perceived exhaustion and do not keep track precisely of the duration, their estimate of the difficulty is mistaken by the duration effect. In the case of the effort allocation task, it would be all the more confusing that, for a given incentive level, higher difficulty levels are associated to shorter effort durations.

Another possibility is that the effort difficulty can be estimated accurately, but that this information is not salient enough to be taken into account to optimize rest duration. This alternative explanation is compatible with the explicit effort allocation task in which explicit (hence more salient) difficulty levels were used in rest optimization. The reported exhaustion levels were impacted more by the effort difficulty than the effort duration, but the post-task report revealed the opposite feeling: this could support the idea that there is a lack of salience of the difficulty level.

#### **3.1.4.4 Method**

The methodological details on the Implicit Effort Allocation Task are presented in the method section page 116.

##### **3.1.4.4.1 Set up**

Same as page 105

##### **3.1.4.4.2 Pre processing**

The effort onset were identified on-line and used to update the screen displayed to the participants. The effort onset was determined as the first sample exceeding 20% of the participant maximal force.

##### **3.1.4.4.3 Maximal force estimates**

Same as page 105

##### **3.1.4.4.4 Behavioral task**

The Cost rating task was presented on a computer screen, and was programmed with Matlab using Psychtoolbox (<http://psychtoolbox.org>).

The task included 7 sessions for each hand that alternated between sessions. Each session comprised 21 trials. The design was factorial and crossed all condition levels: 3 incentives levels (10c, 20c, 50c), 7 effort duration levels (equally spaced from 3s to 7s) and 7 effort difficulty levels (equally spaced from 40% to 60% of the participant maximal force). Each condition was presented only once and corresponded to a trial. The order of condition was pseudo-randomized so that each session had exactly the same difficulty and incentive mean value, and slight differences in mean durations. Both values and sequences of the mean durations were random and varied between participants, as well as the sequences of incentive-difficulty pairs within sessions. This procedure reduces the predictability of conditions between trials. Every trial started with a baseline (1s), followed by the display of the incentive (1s) and then the display of a thermometer that served as a 'go' signal to start the effort. The level of fluid within the thermometer varied in proportion of the force exerted on-line and was scaled so that the target level presented at the top of the thermometer corresponded to the manipulated difficulty level (40% to 60% of the maximal force). The thermometer was displayed as long as the participant had to sustain the effort. This duration was varied experimentally and started when the participant reached the target force, and not when the go signal appeared. The rating screen followed the effort. 'Avez-vous épuisé vos ressources?' ('Are you exhausted now?') was written, and participants indicated their answer between 'Pas du tout' ('not at all') to 'Totalement' ('completely') with a cursor that they moved with left/right key press. The scale had 50 steps but no visible graduation. Rating and validation were self-paced. The last screen lasted 1.5s and summarized the payoff earned at the current trial and the payoff cumulated over the preceding trials. The amount earned during a trial corresponded to the coin values times the proportion of time spent above the target level after it was reached.

#### 3.1.4.4.5 Participants

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. All subjects were recruited via email within an academic database and gave informed consent prior to participating in the study. Participants were between 20 and 39 years old, no self-reported psychiatric or neurological history and no self-reported current psycho-active substance consumption. Participants were paid in cash at the end of the experiment. The payoff corresponded to the money won during the task (there was no fixed amount here). The amount earned during the task was eventually down-scaled (divided by 2.48) so that the payoff during the task corresponded to the monetary value of the coin image presented. Participants were told about this correction prior to the experiment. 'Exp.' refers to the author who collected the data (FM: Florent Meyniel)

Task	Exp .	Period	N included	N male	N excluded	Mean age +/- s.e.m.	Fixed (€)	Var. (€)	Var. (€) range
Cost Rating Task	FM	02/2013	18	7	0	22.2 +/- 0.5	0	29.8	29 – 30

### 3.1.4.4.6 Statistical analysis

Ratings of the Cost Rating Task were submitted to a multiple regression analysis to characterize the effect of different factors. The explanatory variables comprised the manipulated factors (incentive level, effort difficulty, effort duration) and covariates (a constant per session to capture the mean, a linear trend per session to capture drift over trials, and the initial position of the rating cursor). Two-way interaction terms were also included. Variables were z-scored over all trials, except trends that were z-scored within their sessions and padded with 0, and constant regressors. The significance of the parameter estimates was assessed using a random effect at the group level with a two-tailed t-test.

The implicit effort allocation task and the cost rating task were contrasted to assess whether cost levels (inferred from the behavior or reported) were underpinned by similar combinations of effort duration and effort difficulty, and more precisely, whether interaction or summation of these latter variables were better predictors of cost levels. Models of cost had the following general form:

$$C = \beta_0 + \beta_1 X_1^{\lambda_1} + \beta_2 X_2^{\lambda_2} + \beta_3 X_1^{\lambda_1} X_2^{\lambda_2}$$

where  $X_1$  and  $X_2$  are the difficulty and duration levels. Setting the  $\lambda$  terms to 1 made the model linear with respect to factors, setting  $\beta_3$  to zero made the model non-interactive and setting  $\beta_1$  and  $\beta_2$  to zero, but not  $\beta_3$  made the model purely interactive. All these combinations, but the pure constant model, were tested resulting in 6 models. Costs were reported in the Cost Rating Task, they were corrected for covariates (session mean, drift over session, initial position of the cursor) prior to estimation. Costs in the Effort Allocation (implicit) task were inferred from the behavior. The cumulated probability to stop the effort after a given exertion duration was derived directly from the distribution of effort duration, for each difficulty level. The costs were transformed into probability to stop by assuming a sigmoid, also termed probit function:

$$P = \frac{1}{1 + e^{-c}} .$$

This sigmoid was not parameterized, i.e. we did not include a scaling and offset parameter to C since it is redundant with the fit of C itself.

A constant elasticity of substitution (CES) model was also fitted to characterize the curvature of cost. The CES model is:

$$C = (\alpha X_1^\delta + (1 - \alpha) X_2^\delta)^{\frac{1}{\delta}}$$

in which  $\alpha$  ranges from 0 to 1 and characterizes the substitution ratio between  $X_1$  and  $X_2$ , i.e. how much both are equivalent to one another, and  $\delta$  is strictly positive and characterizes the curvature of this equivalence.  $X_1$  and  $X_2$  were rescaled to the same range of values. As there is no scaling factor between the right and left sides of the equality, ratings were rescaled so that C also has the same range of value as  $X_1$  and  $X_2$  in the Cost Rating Task. For the Implicit task, costs were transformed into probability to stop the effort with a sigmoid function. However, we let the offset and scaling in the sigmoid as free parameters so that the probability can be fitted by the CES model.

### 3.1.5 Behavioral results: general discussion

#### 3.1.5.1 Summary

What do we adapt our behavior to when allocating effort over time? We designed tasks to investigate the effort allocation process in the laboratory on short time scales, using hand grip effort. We proposed a computational model to account for this process and we both tested its main predictions and refined its description with model fit. We demonstrated that our behavior is adapted to two different kinds of processes.

The first process is an adaptation on the fly to effort-induced level of cost evidence. We manipulated this level by varying the effort and rest duration and showed that people exert shorter effort when cost-evidence is previously increased by longer effort and they exert longer effort when cost-evidence is previously decreased by longer rest. Our data revealed that cost-evidence imposes a higher bound on effort production but also that a lower bound that corresponds to the maximal improvement rest can bring. We also showed that when people are allowed to adapt rest duration, they lengthen rest to compensate the accumulated level of cost-evidence.

The second process is strategic adaptation. Incentives can regulate the impact of cost-evidence on the effort allocation behavior by pushing back the limit of cost-evidence for higher incentives, which lengthens effort duration, and by speeding up the cost-evidence dissipation at rest which quickens effort resumption. We also showed that in addition to this incentive effect, the effort difficulty

impacts the behavior through both processes and that these effects can be dissociated behaviorally and computationally: the difficulty level the motor system experiences increases cost-level on the fly during the effort whereas the difficulty level that is anticipated strategically increases cost-level dissipation during rest.

Last, we tried to assess whether this computational variable corresponds to the introspective feeling of exhaustion. We estimated the shape of cost-evidence from the decisions to stop the effort in an effort allocation task and from exhaustion ratings. Our model predicts that cost-evidence increases as the product of the effort difficulty and duration, which we found when cost-evidence was inferred from the behavior but not when it was based on the self-report.

### *3.1.5.2 Comments*

In the following discussion, I remind that accumulation models are also used in other domains and emphasize two key differences with respect to our approach. I then try to relate our model to its biological grounding referring to two literatures: the fatigue and sport medicine on the one hand and the effort-based decision-making on the other hand. With the computational insight, I suggest that these two views could correspond to the dissociation between the two kinds of processes put forward above and emphasize the present contribution relatively to other models. Last I acknowledge some limitations of this work, but also suggest some clinical applications.

Evidence-accumulation models have been used to account for perceptual (Gold and Shadlen, 2007; Heekeren et al., 2008; Wang, 2012) and value-based decision making (Basten et al., 2010; Krajbich et al., 2010). In both cases, the accumulation is a way to improve the estimate of a noisy input value (Brunton et al., 2013). In the effort allocation problem, cost is not stationary but dynamic in essence: it gradually increases during effort due to fatigue and it is cleared during rest. I suggested an interpretation of the cost-evidence accumulation model in which the (constant) benefit yielded by an additional unit time of effort is compared online to the (increasing) cost this effort entails. The difference, which is the net utility, decreases over time up to a bound that triggers the effort cessation when the cost does not worth the benefit anymore. From an economic point of view, this interpretation of the cost-evidence accumulation corresponds to a dynamic cost-benefit comparison. This dynamic comparison fits well the idea that the cost-benefit tradeoff is unfolded over time in the effort allocation problem. This interpretation of the accumulation is different from that used in (Basten et al., 2010). In this study, the cost and benefit were compared to estimate the net value of the option presented. The dynamic of the accumulation did not reflect the net value, which was constant, but the integration of this net value over time to decide whether the option should be accepted or rejection.

The second key difference with respect to other accumulation models is that we take into account not only the accumulation, but also the dissipation of the variable to explain effort and rest. One of our results is that the dissipation rate can be modulated strategically according to the incentive and difficulty levels, which suggests that the dissipation is not just a passive relaxation process. The idea that the brain prepares the effort during rest is not surprising. The signature of this preparation can be found in neuronal signals that build up progressively before initiating an effort such as the reduction of motor beta synchrony (15-30Hz) and the readiness potential in the motor cortex. The impairment in motor beta synchrony suppression correlates to the reduced motor ability of Parkinson's disease patients (Kühn et al., 2004; Brown, 2006) and its suppression in healthy subjects corresponds to a gradual commitment in a motor response associated to a decision (Donner et al., 2009; O'Connell et al., 2012). In short, this motor beta synchrony corresponds to a gating of the motor command that needs to be suppressed to allow motor initiation (Engel and Fries, 2010). On the other hand, the readiness potential corresponds to the preparation of the motor plan associated to the action (Lang, 2003; Shibasaki and Hallett, 2006). The dissipation of cost-evidence during rest would therefore not be passive, but active and possibly rooted in some bodily signals corresponding to readiness.

Though both the motor beta synchrony and the readiness potential signals support the idea that the rest contains active task-related processes, they say little about clearing a motor cost. The physiological grounding of cost-evidence could shed interesting light on this process. The fatigue signal should co-vary with proprioception, nociception and the monitoring of the body state (Craig, 2002). Interestingly, the brain areas associated to nociception and proprioception, the so-called pain matrix (comprising in particular, posterior insula, the secondary somatosensory cortex and the anterior cingulate cortex) (Friebel et al., 2011; Wager et al., 2013) are also involved in the placebo effect (Wager et al., 2004). The placebo effect is under the dependence of the opioid system (Benedetti, 2008) which is involved in pain relief (Staahl et al., 2009) and analgesia induced by opioids and placebo could share the same neural network (Petrovic et al., 2002). This suggests that the brain has an endogenous way to control pain with the opioid system. This could also be a means by which it reduces cost-evidence. Indeed, motor fatigue, among many other causes, is due to an inhibition by muscle afferent fibers type III and IV on the nervous system (Enoka and Stuart, 1992; Gandevia, 2001; Martin et al., 2008) and this inhibition can be relieved by opioids (Hilty et al., 2011b). However, opioids are not the only neuromodulator that could be at play at central levels to relieve from pain and fatigue. The catecholamines are also involved. The serotonin pathways in particular participate in the antalgic effect of common pain killers such as acetaminophen (a.k.a. paracetamol) (Smith, 2009) and genetic polymorphisms associated to serotonin re-uptake transporters account for

part of the between-subject variability in pain sensitivity (LaCroix-Fralish and Mogil, 2009; Lindstedt et al., 2011). This is in line with the fact that serotonin levels increase in the brain during effort (Gandevia, 2001) and participate in the sensation of fatigue (Boyas and Guével, 2011). The two neuromodulators outlined here, opioids and serotonin, reveal that the brain can regulate cost-related variables and that these systems could down-regulate cost-evidence depending on our motivations.

This short overview of the biological grounding of cost-evidence tracking in the brain supports the claim that they are two kinds of processes in the effort allocation that are complementary and dissociable. The first process is the online tracking of cost-related variables; the second is strategic regulation of this signal depending on goal pursued. Incentive motivation, i.e. the prospect of higher rewards (Berridge, 2004), would be involved in the second kind of process. The effort difficulty would participate to both: it impacts effort cost for biomechanical reasons (Ma et al., 2009), and it can impact strategic adaptation through the expected cost value. In particular, the expected difficulty discounted the expected reward, which is commonly found in effort-based decision, (Botvinick et al., 2009; Prévost et al., 2010). Interestingly, the two effects of the difficulty were dissociable behaviorally and computationally.

The resulting model accords well with the recent endeavor to think fatigue as the combination of effort-induced physical perturbations and regulatory mechanisms (Abbiss and Laursen, 2005; Noakes, 2011). What our contribution adds to this endeavor is that physical perturbations and regulatory mechanisms are both accounted for by the cost-evidence variable. We add to the previous view a functional distinction: the first level would adapt behavior on the fly to effort-induced perturbations, and the second level would adapt strategically to the first level. The strategic regulation consists in adjusting the parameters of the accumulation process. But critically, both levels are integrated into a single variable (cost-evidence), the dynamic of which guides behavior and implements utility-seeking: the maximization of benefit at the lowest physical cost. Such an optimization principle is at play in other fields such as motor control (Wolpert and Ghahramani, 2000; Todorov, 2006; Rigoux and Guigon, 2012) and decision-making between alternatives (Rangel et al., 2008). The cost-evidence model could contribute to make the link between all these literatures.

Note that the distinction between on the fly, reactive adaptation and strategic, planned adaptation was of particular interest to discuss the dissociated effects of experienced and anticipated effort difficulty. Is it common that costs are less explicit or easy to infer than benefits? At least, any combination of implicit or explicit costs and benefits can be met in our environment. For instance, at

the restaurant you know the cost in advance with certainty: the price on the menu, but you only know for sure the benefit when you experience having the food. The restaurant is the symmetric of the implicit effort allocation task in which benefit is anticipated and cost is experienced. The asymmetry between explicit benefit and implicit cost is pervasive in effort-based decisions. It may even be that such asymmetry is more frequent for effort-based decisions; however, this is not a claim, just an intuition. Moving a fridge is a representative example: it is clear why you have to move the fridge; it is less clear in advance how difficult it is.

The distinction between on the fly adaptation and strategic anticipation could be of clinical relevance to better understand apathy. In particular, based on our model, we could predict two different kinds of apathy: effort could be limited because the expected cost is perceived abnormally high or because the true effort-induced cost is abnormally high. This dissociation could map clinically onto two different diseases with a strong apathy component: depression disorders vs. chronic fatigue. However, apathy is not characterized only by the duration of effort or rest, but also (and maybe more classically) by the intensity of the effort produced (Schmidt et al., 2008). This is a clear limitation of our effort allocation tasks: the free parameters from the subject's perspective are the durations of effort and rest, not the intensity of the effort that is imposed. In our everyday lives, it is on the contrary often operant for the subject to vary the effort intensity, so as to increase the magnitude of the reward, or to increase the probability to obtain it, or to reduce the delay within which this reward is reached.

### *3.1.5.3 What we still do not know and how the neuroimaging and manipulation studies could help*

The behavioral data provided valuable tests and refinements on the cost-evidence accumulation model for effort allocation. However, there are still some unaddressed issues, which neuroimaging studies and manipulation studies have the potential to uncover.

The search for a correlate of the theoretical cost-evidence signal in the brain is motivated by several goals. First, so far this variable has the potential to account for the behavior, but it is a pure theoretical construct. Finding correlate in the neural activity could indicate that this variable is computed in the brain, which is necessary (but not sufficient) for this variable to be used to guide behavior. A strong claim of the model is that a single variable account for decisions related to both effort and rest. This is a strong constraint for the neural correlate. Another constraint on this correlate is that it should not simply wax and wane during effort and rest. The model fitted on behavioral data makes specific predictions on how the slopes and bounds of this signal should be impacted by the difficulty and incentive levels. Such a match between the behavioral predictions and

the neural correlate would be further convincing. Second, if the neural correlate is taken in the strong sense as a way to access the cost-evidence signal through neuroimaging, we could address issues that the behavior left unanswered. For instance, the model fit on behavioral data is blind to whether the change of amplitude is implemented by changes on the signal upper or lower bound. Third, the idea that rest is active, and not passive, was mentioned several times. To support this idea, we can try to show whether the beta band synchrony, a well-known motor preparation signal, could play a pivotal role between increased motivation and quicker effort resumption. Last, localizing the correlate of this signal in the brain could help to refine its biological meaning, in particular whether it is associated to proprioception and nociception.

Manipulation studies could test further the implementation of the cost-evidence accumulation process in the brain. Three key topics deserve to be outlined. First, the strategic aspect of effort allocation is driven, in part, by the use of incentives. Instead of manipulating the incentive, we could manipulate the process at the origin of the incentive effect with a pro-motivational hypnotic suggestion. Second, it seems that the cost-evidence signal is crucially related to the estimation of cost, which could involve proprioceptive and nociceptive signals. To test this idea, we can manipulate the antalgic properties of the body with pro-antalgic hypnotic suggestion and paracetamol. Last, this cost signal, following the rationale presented in the general introduction to this manuscript, should be related to fatigue. To challenge this idea, we could test the role of a neuromodulator involved in fatigue like serotonin, although I acknowledge that fatigue is not likely to involve only one neuromodulator and that the relationship between fatigue and serotonin is still debated.



## 3.2 How does the nervous system solve the effort allocation problem?

### 3.2.1 Neuroimaging data support the cost-evidence model of effort allocation

#### 3.2.1.1 Introduction

The aim of this neuroimaging study is to find whether any brain region could support a signal similar to the cost-evidence signal. The motivation for this search is that the cost-evidence signal is so far, just a theoretical construct. A neural correlate would support our model because if this variable underpins the behavior, it should be computed by the brain. Note that cost-evidence (and its putative correlate) is not trivially related to the behavior: the behavior is only a discrete reflection of the effort allocation process through effort onset and offset, whereas the cost-evidence signal should be a continuous variable varying between these events. Besides, the cost-evidence signal is not a signal simply waning and waxing, the model predicts specific modulation by the effort difficulty and the incentive. Therefore, the requirements to fulfill to be called a cost-evidence signal are selective. In addition, such a correlate might help to refine the model since some properties cannot be addressed by behavioral data alone. For instance, the behavioral model fit is blind to which bound of cost-evidence is impacted by the incentive. Despite distinct interpretation of having modulation of the lower bound (better preparation) or the upper bound (limits pushed back), the amplitude between these bounds is all that matters from the computational point of view.

The previous section on behavioral results showed that the incentive, the experienced difficulty and the expected difficulty affect the effort allocation. It would be interesting to know how these factors impacts the brain signal related to effort allocation. However, the implicit effort allocation task was the first developed and was the one we used for neuroimaging recordings. The explicit and dissociation task were follow-up studies. Therefore, despite the potential interest of how the experience and expected difficulty may be processed by the brain, the neuroimaging results focus on the incentive and the difficulty factors as manipulated in the implicit task. Nevertheless, not knowing beforehand the effort cost can be encountered in common situations, which raises the interesting issue of how the brain may track cost and benefit in this case.

For this neuroimaging study, we used fMRI and MEG because both methods have complementary spatial and temporal resolutions. We took benefit of this complementarity: first, we assumed what the temporal profile of cost evidence should be and searching it in the brain with fMRI, hence using a model-driven approach and second, we extracted the time course of MEG activity in the brain regions identified with fMRI to check, in an independent set of subjects, whether the cost-evidence signal actually meets the assumptions made, hence using a data-driven approach.

### 3.2.1.2 Results

#### 3.2.1.2.1 Behavioral results

The fMRI and MEG data were recorded while participants performed the implicit effort allocation task. The task and the behavioral results were presented in the behavioral section (see page 109), as well as the fit of the cost-evidence accumulation model (see page 119).

#### 3.2.1.2.2 fMRI Results

We used fMRI during task performance to assess 1) whether some brain activity is correlated with cost evidence accumulation during effort and dissipation during rest and 2) whether the amplitude of this putative brain signal is modulated by incentives.

We estimated a first general linear model (GLM1) including cost evidence as a parametric modulator of neural activity at every time point, with constant amplitude throughout all conditions (Figure 29A). This cost evidence signal was significantly expressed (surviving both voxel- and cluster-wise whole-brain correction) in the bilateral posterior insula (secondary somatosensory cortex SII) and the ventromedial thalamus (Figure 29B and Table 4 ). A follow up analysis revealed that the bilateral activation of the insula was not due to a summation of unilateral effects corresponding to each hand: separate contrasts for left and right hands yielded similar results with bilateral activations.

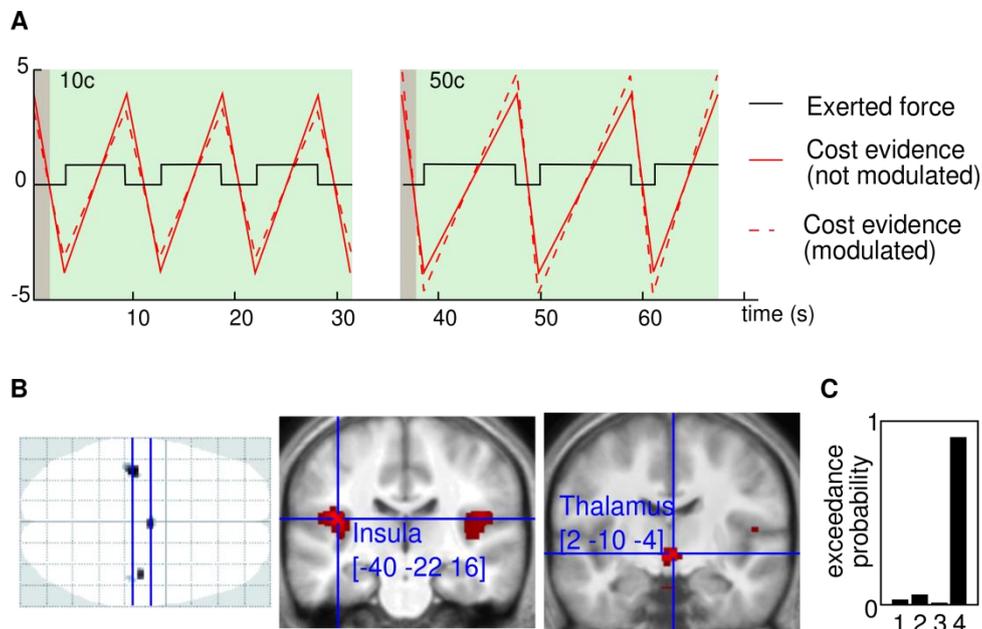


Figure 29 *Implicit task: fMRI result*

**A:** Example of two successive trials with the corresponding cost evidence modeled for fMRI data analysis. Green and grey shading indicates incentive display and effort exertion periods. The incentive was 10c in the left trial and 50c in the right trial. The exerted force is shown in black and cost evidence in red. Two alternative cost evidence regressors are illustrated: one with a constant amplitude (solid line) and one with both bounds modulated by the incentive (dashed line).

**B:** Neural correlates of cost evidence. The statistical parametric maps show brain regions where activity was significantly correlated with cost evidence with constant amplitude. Statistical threshold was set at  $p < 0.05$  with voxel-wise (axial projection plan on the left) or cluster-wise (coronal slices on the right) family-wise error correction for multiple comparisons over the entire brain. The coronal slices were taken along the planes indicated by the blue lines on the glass brain. The [x y z] peak coordinates refer to the Montreal Neurological Institute (MNI) space.

**C:** Modulation of cost evidence amplitude. The bar graph represents the result of a Bayesian model selection comparing the fit of different cost evidence regressors to the activity recorded in the significant clusters shown on slices (bilateral posterior insula and

These two brain regions are considered components of the so-called pain matrix and more generally of the proprioception network (Peyron et al., 2000; Friebel et al., 2011). We extracted them together to form a single region of interest (ROI) that we used in all the following analyses. Note that the same maps were obtained whether or not the exerted force level was included as an additional regressor in the GLM (compare Tables 4 and 6).

	Peak t	Peak unc. p	Peak FWE p	Cluster FWE p	Cluster size (voxels)	Peak coordinates
Left posterior insula	5.825	0.000	0.011	0.006	226	[-40 -22 16]
Ventromedial thalamus	5.693	0.000	0.013	0.018	69	[2 -10 -4]
Right posterior insula	5.469	0.000	0.018	0.006	254	[42 -16 10]
Hypothalamus	4.819	0.000	0.053	0.031	36	[0 0 -20]
Fusiform gyrus	4.774	0.000	0.057	0.018	68	[22 -52 -10]
Cerebellum	4.422	0.000	0.100	0.041	24	[4 -46 -40]
Fusiform gyrus	4.234	0.001	0.132	0.038	27	[18 -70 -6]
Fusiform gyrus	4.227	0.001	0.133	0.035	30	[-10 -68 -4]

Table 4 *fMRI list of significant clusters*

All clusters are listed that were observed using a voxel-wise threshold of  $p < 0.001$  uncorrected and a cluster-wise threshold of  $p < 0.05$  family-wise error (FWE) corrected. The [x y z] peak coordinates in millimeters refer to the Montreal Neurological Institute (MNI) space. The contrast tested is the cost-evidence signal.

	Peak t	Peak unc. p	Peak FEW p	Cluster FEW p	Cluster size (voxels)	Peak coordinates
Ventromedial thalamus	5.490	0.000	0.017	0.022	88	[2 -10 -4]
Right posterior insula	5.463	0.000	0.018	0.007	367	[42 -16 8]
Left posterior insula	5.410	0.000	0.019	0.008	286	[-42 -22 18]
Fusiform gyrus	4.661	0.000	0.064	0.005	616	[22 -52 -10]
Hypothalamus	4.578	0.000	0.073	0.033	59	[0 0 -20]

cerebellum	4.540	0.000	0.078	0.032	61	[4 -48 -40]
------------	-------	-------	-------	-------	----	-------------

*Tableau 5 fMRI significant clusters when including the motor output as a covariate*

*All clusters are listed that were observed using a voxel-wise threshold of  $p < 0.001$  uncorrected and a cluster-wise threshold of  $p < 0.05$  family-wise error (FWE) corrected. The [x y z] peak coordinates in millimeters refer to the Montreal Neurological Institute (MNI) space. The contrast tested is the cost-evidence signal.*

Three additional GLM were built to account for amplitude modulation by incentives with changes in the upper bound (GLM2), lower bound (GLM 3) or both bounds (GLM4). BMS analysis (Figure 29C) indicated that the activity extracted from the ROI was best explained (with  $x_p > 0.99$ ) by GLM including amplitude modulation (GLM2-4 versus GLM1), i.e. GLM that were not used to identify the ROI. Among the three possible modulations, changing both bounds (as illustrated in Figure 29A) was the most probable ( $x_p > 0.96$  for GLM4). Thus, fMRI data revealed that proprioceptive regions continuously signal cost evidence over effort and rest periods. Additionally, they showed that the two bounds triggering effort cessation and return are moved apart when incentives are increased, an effect that could not be inferred from behavioral data.

### 3.2.1.2.3 MEG results

We used MEG to confirm the conclusions drawn from the fMRI study with a reverse approach: instead of a model-driven approach showing that a theoretical signal fits brain activity, we followed a data-driven approach to show that brain activity fulfills theoretical predictions. More specifically we assessed 1) whether scalp activity arising from the ROI sources was ramping up and down during effort and rest periods and 2) whether incentive and difficulty effects on bounds and slopes conformed to the model optimized on behavioral and fMRI data.

MEG time-series were epoched into behaviorally-defined rest and effort periods, which were resampled to a same duration and averaged over conditions and subjects. The principal component analysis (PCA) performed on this grand average revealed that both effort and rest periods were dominated by ramping activity (Figure 30A). Indeed, the first components were linear variations that accounted for most of the signal variance (91% for effort and 72% for rest periods). The sources of these first components were reconstructed subject-wise with a minimum-norm procedure, either unconstrained or informed by setting priors on the fMRI-based ROI. BMS showed that the reconstruction using the anatomical priors had a much higher exceedance probability ( $x_p = 0.90$  for effort and  $x_p = 0.94$  for rest). Thus, the effort and rest ramping signals that dominated scalp activity were most likely to arise from the same sources: the ROI identified with fMRI. This is consistent with the idea that these regions (posterior insula and ventromedial thalamus) generate an accumulation signal throughout task performance.

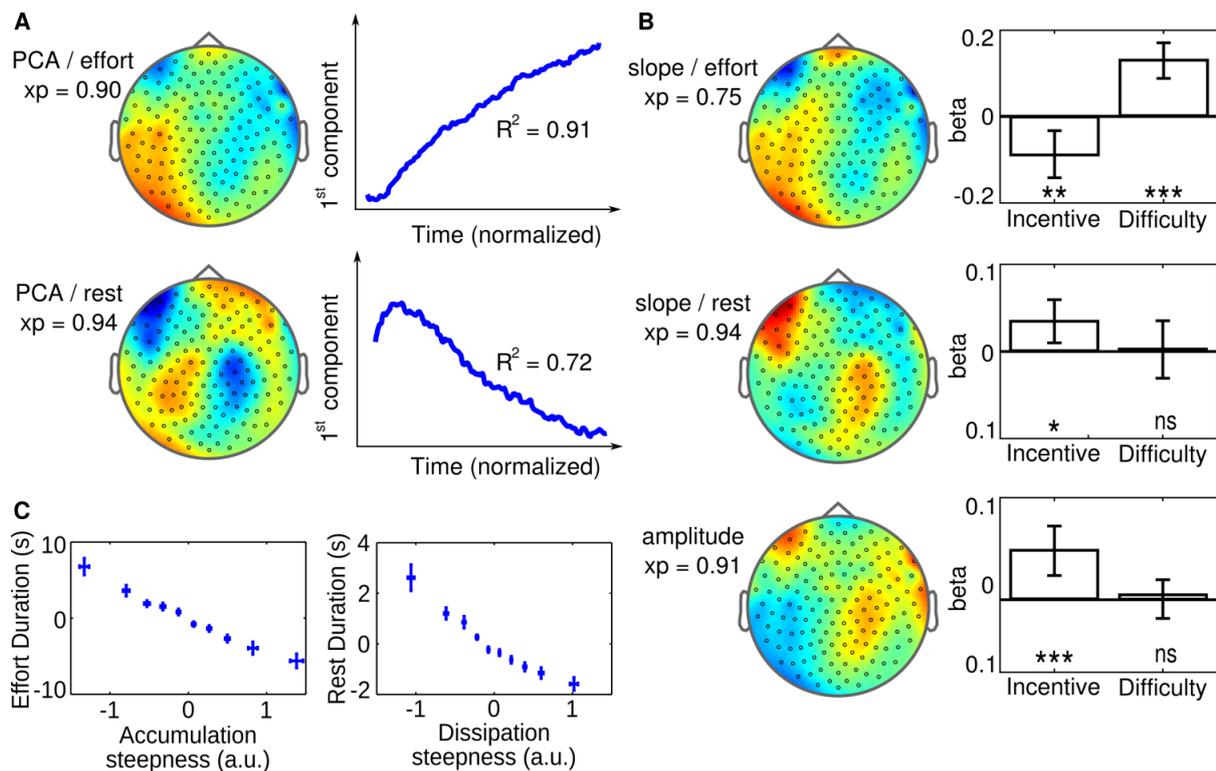


Figure 30 *Implicit task, MEG results*

**A:** Principal component analysis performed on average activity for effort and rest epochs separately. The scalp topography, time series and proportions of variance explained ( $R^2$  statistics) correspond to the first component.

**B:** Regression analysis performed on every single event to estimate accumulation slope (by fitting a line to rest and effort epochs separately) and amplitude (by fitting a V shape to two consecutive rest and effort epochs). For each analysis the scalp topography shows the slopes and amplitudes averaged over trials and subjects. The bar graph represents the coefficients (betas) obtained for the two experimental factors (incentive and difficulty) with a linear regression model fitted on the slopes and amplitudes reconstructed within the regions identified with fMRI. Error bars are inter-subject 5% confidence intervals. Significance of group-level t-tests: \*\*\*  $p < 0.0005$ , \*\*  $p < 0.005$ , \*  $p < 0.05$ .

For each scalp topography the  $x_p$  value denotes the exceedance probability of the source reconstruction model that included as priors the regions identified with fMRI.

**C:** Correlation between residual durations and accumulation or dissipation steepness, obtained by regressing out incentive and difficulty effects. Durations were defined from behavioral data and slopes from MEG data reconstructed in the source space. Effort duration was plotted against accumulation steepness (left) and rest duration against dissipation steepness (right). For illustration, data were binned into deciles in every subject. Dots represent inter-subject means  $\pm$  standard errors for the 10 bins.

Note that we use steepness instead of slope for the accumulation signal (power) reconstructed in the source space since it is not signed.

We then returned to scalp raw time-series (no resampling and averaging) and estimated accumulation and dissipation slopes using linear regression for every single effort and rest epoch on each channel (Figure 30B). In each subject slopes were averaged over epochs and conditions, for effort and rest separately. The sources of slope topography were reconstructed subject-wise, using either unconstrained or informed minimum-norm procedure as above. The BMS showed again that

including fMRI-based priors largely improved reconstruction plausibility ( $x_p=0.75$  for effort and  $x_p=0.94$  for rest).

Individual reconstruction matrices were then used to estimate accumulation and dissipation slopes in the source space for each epoch and subject. The variations of these slopes were fitted with a linear model including incentive and difficulty levels as regressors. Incentives affected both accumulation slope during effort and dissipation slope during rest ( $p<0.004$  and  $p<0.008$ , respectively), whereas difficulty only impacted accumulation, not dissipation slope ( $p<4 \cdot 10^{-6}$  and  $p>0.9$  respectively). Thus, the modulations identified from behavior (incentive increasing  $S_r$  and difficulty increasing  $S_e$ ) were confirmed and an additional modulation (incentive decreasing  $S_e$ ) was revealed. The effect of this additional modulation is to prolong effort periods for higher incentives, which was so far entirely imputed to larger amplitude between bounds.

To assess whether the amplitude of the accumulation signal identified in MEG activity was modulated, we fitted a single V shape to every contiguous rest-effort duplet. Regression coefficients were taken as amplitude estimates and were submitted to the same analysis as slopes. The source reconstruction for these amplitudes as well was much more plausible when including the fMRI-based ROI as priors ( $x_p=0.91$ ). The linear regression performed on amplitudes in the source space showed a significant increase with higher incentives but no significant modulation by difficulty ( $p<3 \cdot 10^{-4}$  and  $p>0.65$ , respectively). Thus, MEG data revealed that incentives in fact affected all three parameters of the accumulation model ( $S_e$ ,  $S_r$  and  $A$ ), while the difficulty effect remained relatively specific (only impacting  $S_e$ ).

Additionally, we tested whether incentive and difficulty effects would also be observed across subjects. Between-subject correlations between mean durations and accumulation parameters replicated all four effects that were found to explain the modulation of behavior by incentive and difficulty levels (Figure 31). Thus, subjects with greater effects on accumulation / dissipation slope or amplitude had greater effects on effort / rest duration.

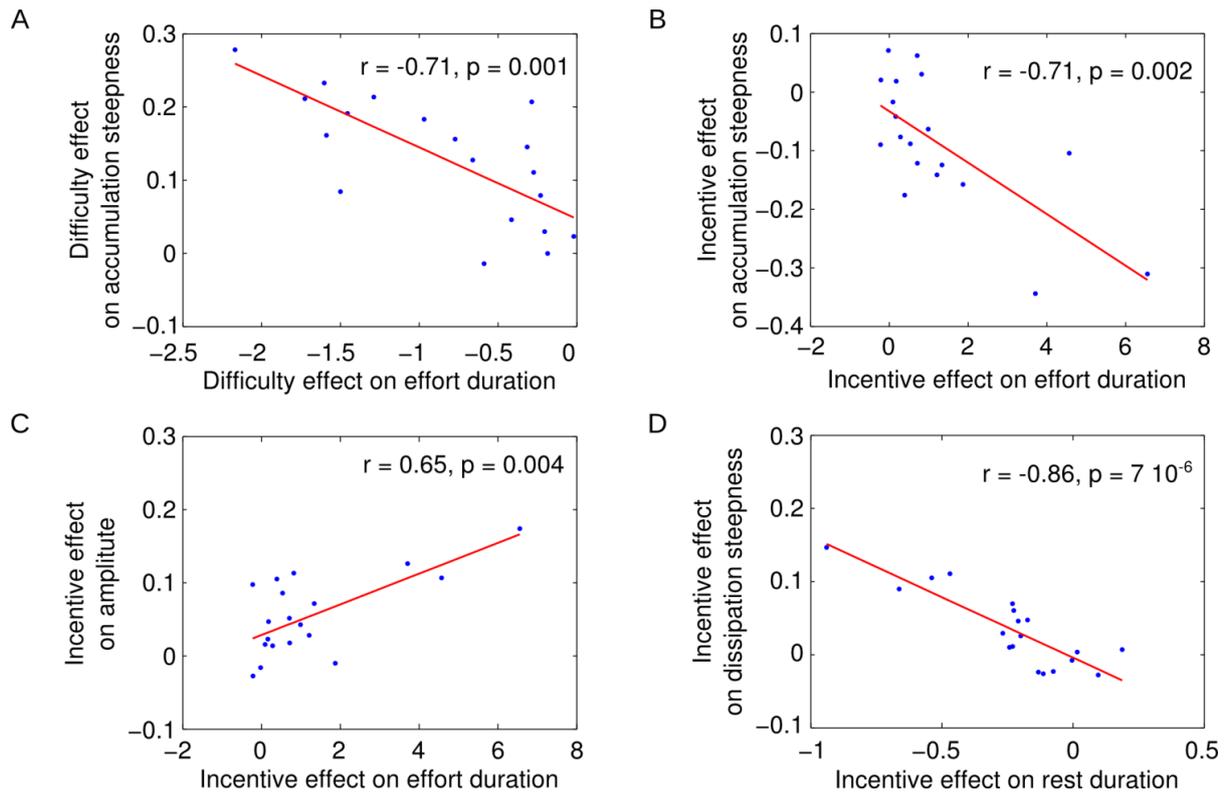


Figure 31 **Across-subject correlations between behavioral and computational effects of experimental factors (monetary incentive and task difficulty).**

Behavioral effects are: shortened effort duration with higher task difficulty, prolonged effort duration with higher monetary incentive, and shortened rest duration with higher monetary incentive. Correlations were searched with all four computational effects that were found to explain these behavioral effects in MEG recordings. Computational effect refers to modulation of one parameter (accumulation steepness, dissipation steepness or amplitude between bounds) of the signal reconstructed in our ROI. Note that we use the term ‘steepness’ and not ‘slope’ because this power signal has no sign. The four computational effects are steeper accumulation with higher difficulty (A), slower accumulation with higher incentives (B), larger amplitude with higher incentives (C) and steeper dissipation with higher incentives (D). R-values are Pearson rho correlation coefficients; p-values indicate the significance of robust-fit regression (which under-weights potential outliers).

The above analyses suggest that incentive and difficulty effects on effort and rest durations are underpinned by modulation of cost-evidence accumulation slopes and bounds. However, a strong prediction of the model is that decisions to engage and terminate effort are triggered by the cost-evidence signal reaching a predetermined threshold. This implies that effort and rest durations should be respectively correlated with accumulation and dissipation slopes, even when the correlation induced by our experimental manipulation is removed. To assess this prediction, we regressed the variance related to incentive and difficulty levels out of the reconstructed signal and duration, and tested the correlation between residual steepness and duration (Figure 30C). The correlation across trials was negative for every subject and for both effort and rest periods: the steeper the slope, the shorter the duration. Correlation coefficients were highly significant at the group level for both effort ( $p = 1.4 \cdot 10^{-8}$ ) and rest periods ( $p = 5.9 \cdot 10^{-11}$ ).

Taken together, fMRI and MEG findings demonstrated that cost evidence is indeed tracked in proprioceptive brain regions and that the impact of potential benefits on the accumulation process could be more complex than suggested solely on the basis of behavior.

### *3.2.1.3 Discussion*

We addressed the issue of whether the human brain encodes a hidden variable that could underpin effort allocation over time: cost-evidence. We found a brain signal that linearly accumulated cost evidence during effort production and dissipates at rest, in the posterior insula, secondary somatosensory cortex and ventro-medial thalamus. The observed decisions to stop and resume effort production corresponded to this cost evidence signal reaching upper and lower bounds. In addition, this brain signal contains all the modulations by the incentive at stake and the effort difficulty that could be inferred from the behavior, which strengthen the similarity between the computational variable and its brain counterpart. If this signal is taken as truly reflecting the cost-signal in the brain, the opportunity to access it through neuroimaging makes it possible to uncover effects that could not be inferred from the behavior alone, such as the modulation of both bound of the signal by the incentives. This suggests two distinct mechanisms by which effort production can be improved: shifting the upper limit corresponds to pushing back the limit of effort production when more motivated whereas shifting the lower limit corresponds to a better preparation, when more motivated, to the subsequent effort to come.

This complex neural signal with precise spatio-temporal characteristics was identified with both fMRI and MEG using standard analytical tools, which validates the usual assumptions about the relationships between electromagnetic and hemodynamic activity (Logothetis, 2008; Gutschalk et al., 2010b; Lee et al., 2010; Rosa et al., 2010; Vartiainen et al., 2011). Indeed, both depend on the neuronal activity, in particular the extra-cellular fields. The MEG (and EEG) signals correspond to extracellular currents that are structured in the cortical layer in such a way that they do not cancel out each other, this happens for instance with large pyramidal neurons (Hämäläinen et al., 1993; Baillet et al., 2001; Buzsáki et al., 2012). On the other hand, the BOLD signal is also well correlated to the local field potential (Logothetis et al., 2001; Logothetis and Wandell, 2004; Shmuel et al., 2006; Goense and Logothetis, 2008) rather than the spiking activity (Sirotnin and Das, 2009). Due to its high temporal resolution, MEG brought evidence that the electromagnetic activity emitted by our regions of interest indeed followed the neural dynamics that was modeled and convolved with hemodynamic response to fit fMRI data. And due to its high spatial resolution, fMRI confirmed that our theoretical cost evidence signal was indeed represented in proprioceptive brain regions.

Our interpretation was confirmed by the signal being encoded in regions pertaining to the pain matrix, such as the posterior insula and ventral thalamus (Peyron et al., 2000; Friebel et al., 2011; Brodersen et al., 2012; Wager et al., 2013). More precisely, the main activation foci were located in the operculum parietale area 1, which has been implicated in high order somatosensory processing, in connection with both the ventral thalamus and parietal network (Eickhoff et al., 2010). The classical pain matrix additionally includes mid-cingulate regions (Mohr et al., 2005; Beckmann et al., 2009), which were also activated in relation to cost evidence in our results but slightly below the statistical threshold for significance. Interestingly, direct electrical stimulation of the posterior insula was shown to induce painful sensations (Ostrowsky et al., 2002), which is actually very specific to this brain region when this effect is contrasted to stimulation of other sites of the brain (Mazzola et al., 2012). However, we cannot infer from brain localization alone that the cost evidence variable can be equated to a subjective pain sensation. Indeed, different functions have been assigned to this brain network and particularly to the posterior insula, such as the monitoring of bodily states (Craig, 2009b; Naqvi and Bechara, 2009; Jones et al., 2010).

It should be emphasized that the signal labeled here as cost evidence did not mirror the behavioral output. Its linear fluctuations, dipping when effort starts and peaking when effort ends, were decorrelated from the force produced, which followed a boxcar dynamics. Thus, contrary to the behavioral output, the cost evidence signal spanned the same range of values in effort and rest periods. We believe that the dissipation at rest arises from an active process rather than from a passive relaxation. This is because, at the beginning of trials when subjects had not yet squeezed the handgrip, the cost evidence signal was first brought down to the lower bound in anticipation of effort exertion. In addition, it should be notice that the cost-evidence is not defined only at behavioral transition, but continuously during effort and rest. A previous report could support the current finding: Hilty and colleagues found that task failure, i.e. inability to maintain effort was preceded by high activity in the posterior insula and medial thalamus, when compared to similar effort that was succeeded (Hilty et al., 2011a). In other words, the BOLD signal was higher in these regions when participants gave up compared to when they maintained: this is precisely what is predicted by the cost-evidence signal.

The anticipation of the effort reflected by cost-evidence dissipation during rest resembles the motor readiness potential that is known to precede limb movements by a few seconds. However, this readiness potential has been localized in motor cortical areas and not in deep proprioceptive regions (Deecke et al., 1969; Colebatch, 2007). Besides, the complex modulation of our cost evidence signal by monetary incentives suggests that it is not merely related to motor output.

Task difficulty probably impacted the behavior on the fly, as it was not explicitly mentioned to subjects. It was noticeable in the implicit effort allocation task that rest durations were not affected by the effort difficulty. Consistently, difficulty effects specifically manifested as an increased accumulation rate during effort, leaving unchanged the dissipation rate during rest. Incentive effects were two-fold: 1) there is clear evidence that they speeded dissipation of cost evidence during rest, and that they slowed accumulation during the effort (though this effect was less convincing), and 2) they moved the bounds within which cost evidence fluctuates. The first process could reflect the intervention of an opponent motivation signal that would be continuously subtracted to cost evidence throughout effort and rest periods. This signal might come from brain regions involved in reward processing, or in the trade-off between reward and effort, such as the ventral striatum, the anterior cingulate cortex or the ventromedial prefrontal cortex (Croxson et al., 2009; Kolling et al., 2012; Schmidt et al., 2012). The second process might implement the intuitive psychological phenomenon that, when motivated, we literally push back our limits, allowing our body to work closer from exhaustion. It could be explained by reward-related regions adjusting decision thresholds in regions that are downstream to the posterior insula in the chain leading to motor outputs. Yet these speculative mechanisms that might adjust accumulation parameters to expected benefits would require further investigation. In particular, it is interesting that incentive signals in the BOLD data were actually limited in this task: there was little incentive effect during the trial, either sustained or at effort or rest transitions. The only clear incentive effect was found in the striatum at the display of the coin image. It is therefore unknown how the incentive effect remains pervasive during the trial and affects the cost-evidence parameters and the behavior.

Note that incentive effects argue against the possibility that the signal might encode money and not cost accumulation, since the slope observed during effort tended to decrease, not increase, with higher incentives. These effects are also consistent with reports that placebo analgesia reduces insular and thalamic BOLD responses to pain stimulation (Wager et al., 2004), suggesting that the brain can adjust the sensitivity of these regions depending on expectations. The placebo analgesia could rely on the opioid system, although it is not the only potential mechanism: increased connectivity within the prefrontal cortex (between the rostral anterior cingulate cortex and the ventro lateral prefrontal cortex) was also detected specifically during placebo analgesia (Petrovic et al., 2010). Evidence in favor of the involvement of the opioid system in placebo analgesia includes the report that placebo-induced and opioid-induced analgesia share common networks in the brain (Petrovic et al., 2002); that placebo-analgesia is associated to endogenous binding of opioid receptors, as evidenced with positron emission tomography (Zubieta et al., 2005) and that opioid receptor antagonists, like naloxone, block the placebo analgesia (Bandettini et al., 1997). It is

therefore possible that the endogenous analgic mechanisms are at play in the effort allocation problem, in particular to allow distinct cost-evidence level depending on the incentives. The localization of the cost-evidence signal in brain regions processing the placebo and opioid related analgesia is not a demonstration, but gives credit to this idea.

### **3.2.1.4 Methods**

Details on the behavioral task and the participants are provided in the Methods section page 116. We used the implicit version of the effort allocation task.

#### **3.2.1.4.1 Imaging data acquisition**

##### *3.2.1.4.1.1 Experimental setting*

Prior to scanning, participants were given written instructions to the task, which were repeated step by step orally. Subsequently, they were escorted inside the MRI or MEG scanning room and invited to find a comfortable body position that they could keep throughout the experiment. The only change was passing the power grip from one hand to the other between sessions. We used homemade power grips composed of two plastic cylinders compressing an air tube when squeezed. The tube led to the control room, where it was connected to a transducer converting air pressure into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to the force exerted. The signal was fed to the stimuli presentation PC via a signal conditioner (CED 1401, Cambridge electronic design, UK) and read by a Matlab program (The MathWorks Inc., USA). Stimuli presentation was also programmed with Matlab using Cogent 2000 (Wellcome Department of Imaging Neuroscience, London, UK).

##### *3.2.1.4.1.2 fMRI data acquisition*

Subjects' head was constrained using foam and sand packs to limit movements. Functional echo planar images (EPI) were acquired with a T2\*-weighted contrast on a 3 Tesla scanner (Siemens Trio). Inter-leaved 2mm slices separated by a 1.5mm gap and oriented along a 30° tilted plane were acquired to cover the whole brain with a TR of 2s. The first 5 scans were discarded to allow for equilibration effects. All preprocessing steps were performed using SPM8. Structural T1-weighted images were also acquired, co-registered with the mean EPI, segmented and normalized to SPM standard MNI (Montreal Neurological Institute) T1 template. Normalized T1-images were averaged between subjects to localize group-level functional activations. EPI were spatially realigned and normalized (using the same transformation as for structural images), and spatially smoothed with a 8mm FWHM (full-width at half-maximum) Gaussian kernel.

#### 3.2.1.4.1.3 MEG data acquisition

A whole-head MEG system comprising 151 axial gradiometers (CTF Systems) was used to sample brain activity at 1250 Hz with online low-pass filter of 300 Hz. Head position was determined using marker coils at fiducial points (nasion, left and right ears). Ocular artifacts were marked manually and removed using the Gratton method with DataHandler (Cogimage, CRICM, Paris, France). Data were imported in Matlab and displayed using FieldTrip (Donders Institute, Nijmegen University, The Netherlands). MEG signal was low-pass filtered offline at 30 Hz. Effort onsets and offsets were detected manually based on the electromyogram (EMG). A template mesh (8196 tessels) and individual fiducials were used in SPM8 to compute a single shell head model and a lead field matrix per subject and session.

#### 3.2.1.4.2 Statistical analysis

##### 3.2.1.4.2.1 fMRI data analysis.

All general linear models (GLM) included realignment parameters as covariates of no interest to correct for movement artifacts. Regressors of interest were specified at the 125ms scale, and convolved with the canonical hemodynamic response function (HRF) and its first temporal derivative. The GLM included two categorical regressors: one modeling coin display onset with a delta function and one modeling the entire session with a boxcar function. There were also two parametric regressors: one modulating coin display by its value (10, 20 or 50 cents) and one modeling cost evidence variation over the entire session. Cost evidence was continuously modeled over effort and rest periods defined from the behavior, with linear increases and decreases between constant minimum and maximum, and then z-scored. Thus, the parametric cost evidence regressor was ramping up and down, between positive and negative values, during task trial and put to zero between trials (see Figure 29A).

Regression coefficients were estimated at the subject level using the standard restricted minimum likelihood (ReML) estimation implemented in SPM8. Individual linear contrasts of HRF regressors were then taken to a group-level random-effect analysis using one-sample t-tests. Statistical thresholds corresponding to correction for multiple comparisons over the entire brain were determined using the randomization (n=10 000 permutations) technique implemented in FSL (Centre for Functional MRI of the Brain, Oxford, UK). Two thresholds were used: a voxel-wise family-wise error (FWE) rate of  $p < 0.05$  and a cluster-wise FWE rate of  $p < 0.05$  (defined for voxels surviving  $p < 0.001$  uncorrected). The cost evidence regressor mapped onto three regions (bilateral posterior insula and ventromedial thalamus) that survived both voxel-wise and cluster-wise corrections. These

three clusters formed at  $p < 0.001$  uncorrected were grouped together to form a single region of interest which was used for all subsequent analyses.

Three other GLM were built that differed only on the cost evidence regressor, which now had an amplitude modulated by incentives in the same proportion as the model optimized on behavior. This modulation could in principle rely on the upper bound only (GLM2), or the lower bound only (GLM3) or be shared between both bounds (GLM4). Variational Bayes estimation procedure implemented in SPM8 was used to estimate for each subject the GLM log-evidences, which were summed over all voxels included in the ROI. Individual log-evidences were then submitted to a group-level random-effect BMS to identify the most probable cost evidence model given the ROI activity.

#### *3.2.1.4.2.2 MEG data analysis.*

For the principal component analysis (PCA), MEG data were epoched into rest and effort periods, resampled to 1250 points and averaged over conditions. A PCA was computed on the grand average (over subjects) to estimate the  $R^2$  statistic of each component for rest and effort periods, separately. PCA were also computed in each subject to reconstruct the sources of the first component with SPM8, using a minimum-norm algorithm that could include or not include the fMRI-based ROI as priors. The best reconstruction method (with or without priors) was determined using group-level BMS (Daunizeau et al., 2005).

For the slope and amplitude analyses, three regressions were performed on the scalp raw time series. A linear trend was first fitted separately on rest and effort epochs to estimate accumulation and dissipation slopes, respectively. Then a V-shape was fitted on contiguous rest-effort epochs to estimate the signal amplitude. Sources of the mean slopes and amplitudes were reconstructed subject-wise using the same procedure as for PCA. The estimated scalp-to-source reconstruction matrices were then used to invert each epoch. Activity in the source space was rectified (absolute value), log-normal transformed and averaged within the ROI. The resulting activity was analyzed using a linear model including three regressors: incentive (10, 20, 50) and difficulty (70, 80, 90) levels plus a constant. The significance of regression coefficients was estimated at the group level using a two-tailed t-test.

#### *3.2.1.4.3 Methodological clarification on MEG results*

The statistical significance of some MEG results might have been overestimated. Here, the bias uncovered is described and a correction is proposed. Fortunately, the conclusions remain unchanged. In other words, despite the bias, our data truly carry the effects that we initially reported.

As a reminder, the analysis dealt with a signal, recorded using MEG, which was waxing and waning while subjects alternated effort and rest periods. The aim was to characterize the modulation of the accumulation and dissipation slopes by the experimental factors, in a region of interest (posterior insula) that was identified beforehand using fMRI. To characterize the effect of experimental factors (incentive and difficulty levels) on the signal arising from the ROI (Figure 30B-C), we followed the pipeline described below.

- **Extracting the dependent variable.** We estimated the signal slope using linear regression, for every epoch and every sensor at the scalp level.
- **Estimating the underlying source.** We estimated with a minimum-norm procedure, including our ROI as prior, the source of the mean slope topography, separately for effort and rest epochs. This estimation provides an inversion matrix  $M$  that relates the scalp signal  $\mathbf{Y}_m$  to the source activity  $\mathbf{J}_m$ :  $\mathbf{J}_m = M * \mathbf{Y}_m$ . We used this inversion matrix  $M$  to estimate the source of slope topography in every epoch: for any given epoch  $k$ , the corresponding source distribution is  $\mathbf{J}_k = M * \mathbf{Y}_k$ .
- **Calculating regional source intensity.**  $\mathbf{J}$  (in bold) is a vector of source intensities, with each data point corresponding to a source location. The sign corresponds to the direction of the current, inward or outward with respect to the cortical surface. The regional source intensity was computed by averaging the rectified (absolute value) source intensities within this region. Removing the sign was the key step that produced the bias.
- **Testing experimental factors.** The source intensity in our ROI was then log-transformed to ensure Gaussian data, and regressed against a General Linear Model (GLM) that included incentive and difficulty levels as predicting variables.

If we consider now a stochastic auto-correlated noise  $s$ , such that for each time point  $t$ ,  $s(t+1) = s(t) +/- 1$  (with an equal chance to increase or decrease between two time points). Mathematically, this is a random walk on the integer numbers. The expected value of  $s(t)$  is null for any time point, because going up or down is equally likely. However, the expected distance from the starting point at time  $t$ , ( $|s(t)|$ , also known as the dispersion), scales with the square root of  $t$ . The absolute value of the signal slope follows a similar relationship; it is on average steeper for shorter duration.

Biological signals have an auto-correlation component (Schurger et al., 2012). Because the sign of the signal is lost in our pipeline of analysis, this component could be sufficient to produce the effect we report: the shorter the epoch duration, the steeper the slope. Thus, once we know that, for instance, higher incentives results in shorter rest duration, we can expect from the bias alone that the reconstructed slope in the ROI will be steeper.

The question we had to ask was whether our experimental factors produced any effect on the slope steepness, on top of the artifact due to auto-correlated noise in our data? We used a subtraction approach to answer this question. We removed from our dependent variables (slopes estimated at the sensor level) any effect that could be linearly explained by the duration itself, or the incentive and difficulty levels. We then ran our pipeline and obtained GLM parameters ( $\beta_{\text{noise}}$ ) that correspond to artifacts linked to auto-correlated noise. We then compared these parameters to those obtained in our initial analysis ( $\beta_{\text{full}}$ ). The difference between  $\beta_{\text{full}}$  and  $\beta_{\text{noise}}$  reflects the effect of interest, i.e. a corrected effect size, uncontaminated by the bias potentially introduced by auto-correlated noise. The table below reports the corrected parameters and associated Student statistics for the effects of the two factors on the two epochs.

	Factor	regression parameter	95% Student confidence interval	Group level p-value
EFFORT	Incentive	-0.035	[-0.074 0.004]	0.038
	Difficulty	0.037	[0.013 0.061]	0.002
REST	Incentive	0.009	[0.002 0.016]	0.008
	Difficulty	-0.004	[-0.018 0.011]	0.30

Compared to the initial results, the direction of the effects is the same, and their significance (unilateral t-test) is also preserved. We note that the incentive effect on effort duration is now weaker, but it was already concluded in the publication that this effect (which was not present in our modeling and fMRI analyses) was not necessary to explain the behavior. The corrected results provided here strengthen this conclusion.

### 3.2.2 Get back to work: A mechanism to translate incentives into speeded effort resumption

#### 3.2.2.1 Introduction

The aim of the present study was to better understand how the rest durations are regulated. Behavioral and neuroimaging data showed that rest duration is strategically modulated according to the monetary incentive in the implicit effort allocation task. This effect contradicts the view that rest would be a passive relaxation. Instead, it is more similar to a behavioral activation, motivated by a reward prospect. This incentive motivation process has been extensively investigated in neuroscience (Berridge, 2004; Schultz, 2006). Central to this incentive process is the issue of how the value of a world state can be translated into a motor code. How could this process be implemented in the brain

and could we find an electrophysiological marker of it? Two lines of research suggests that the motor loop between the basal ganglia and the motor cortex could implement this motivated regulation of speeded effort resumption, which could be tracked in the level of motor beta synchrony. In other words, from the electrophysiological point of view, the motor beta synchrony could have a pivotal role to translate higher incentive into speeded effort resumption. The rationale for this hypothesis is presented as the conjunction of these two lines of research.

As introduced in the general introduction of this dissertation, the basal ganglia system seems a likely candidate for mediating this incentive motivation process, as revealed by electrophysiology or functional MRI in healthy subjects (Schmidt et al., 2012; Tachibana and Hikosaka, 2012), and by the effects of focal lesions or degenerative disease in patients (Czernecki et al., 2002; Schmidt et al., 2008). The case of Parkinson's disease (PD) is particularly enlightening. PD is primarily due to degeneration of dopaminergic neurons, which results in a lack of dopamine projection in the striatum and a set of motor symptoms (akinesia, rigidity, tremor). However, this reduction of movement (or hypokinesia) could easily be reframed in terms of dysfunctional motivation, i.e., as a difficulty in activating motor plans leading to better states (Mazzoni et al., 2007; Baraduc et al., 2013). This idea accords well with an abundant literature that has implicated dopamine in incentive motivation (Berridge, 2004; Salamone and Correa, 2012).

Interestingly, one electrophysiological hallmark of PD is the high level of synchronous oscillations in the beta band (Schnitzler and Gross, 2005; Brown, 2006; Uhlhaas and Singer, 2006). Dopamine replacement therapy, as well as deep brain stimulation, simultaneously reduce beta synchrony and alleviate hypokinesia (Brown et al., 2001; Kühn et al., 2008). More precisely, delay in voluntary movement initiation was linked to insufficient reduction of beta synchrony in the sub-thalamic nucleus (Kühn et al., 2004; Williams et al., 2005). Furthermore, simultaneous deep and surface recordings showed that beta oscillatory activities were coherent across basal ganglia nuclei and motor cortical areas (Klostermann et al., 2007; Litvak et al., 2010). Taken together, these findings suggest that reducing motor beta synchronization (MBS) represent a neural mechanism through which expected rewards may facilitate action initiation.

In healthy people, there is indeed a progressive reduction of MBS (typically in the 13-30Hz range and centered on the precentral cortex), preceding movement initiation (Jasper and Penfield, 1949; Feige et al., 1996; van Wijk et al., 2012). It is thought to play a "gating role", meaning that high beta synchrony maintains the motor status quo, whereas low beta synchrony allows for a motor change (Engel and Fries, 2010; Jenkinson and Brown, 2011). Direct evidence of this idea is the finding that driving motor cortical activity in the beta rhythm slows motor performance (Pogosyan et al., 2009;

Joundi et al., 2012). Therefore, we reasoned that in incentive motivation paradigms, the reward magnitude should modulate MBS, which, in turn, should have an impact on movement initiation. To test this hypothesis, we analyzed MEG activity previously recorded in healthy participants while they were trying to minimize the duration of breaks during force production, in order to maximize their payoff.

### 3.2.2.2 Results

#### 3.2.2.2.1 Factors affecting rest behavior

We used the implicit effort allocation task that was presented earlier (see the behavioral section, page 109) with the fMRI and MEG study. The same MEG data were analyzed here. The behavioral results presented earlier were pooled over the fMRI and MEG groups. Although the results were similar in both groups, the results for the MEG group alone are presented here since this group is the focus of the present study. All subjects spontaneously alternated effort and rest periods during the course of trials, suggesting that incentive levels were high enough to induce effort production, and difficulty levels high enough to impose breaks.

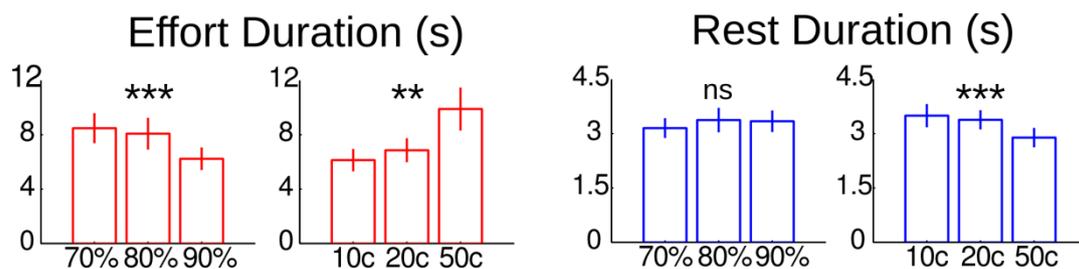


Figure 32 *Implicit task, MEG behavioral results.*

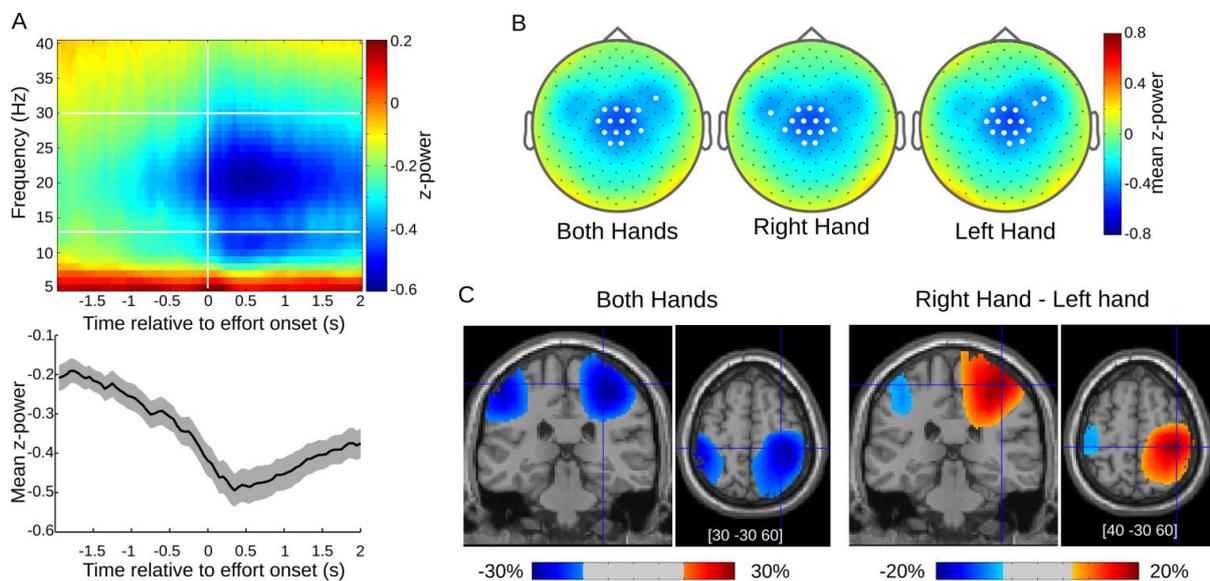
Average data sorted by incentive and difficulty levels. Bars are mean effort and rest epoch durations and error bars the inter-subject standard errors. Significance of repeated-measure ANOVA main effects: \*\*\*  $p < 0.0005$ , \*\*  $p < 0.005$ , \*  $p < 0.05$ .

Two repeated-measure ANOVA were performed to characterize the effect of the manipulated factors (incentive and difficulty levels) on the durations of rest and effort periods, separately (see Figure 32). For higher incentives, participants prolonged effort duration ( $F_{2,36}=11.1$ ,  $p=2.7 \cdot 10^{-3}$ ) and shortened rest duration ( $F_{2,36}=10.5$ ,  $p=3.2 \cdot 10^{-4}$ ). These two effects contributed to increase the payoff, as they augmented the total time spent squeezing the grip when more money is at stake. Higher difficulty shortened effort duration ( $F_{2,36}=14.0$ ,  $p=4.3 \cdot 10^{-5}$ ) but did not significantly affect rest duration ( $F_{2,36}=1.0$ ,  $p=0.35$ ). In the following, we focus on rest duration, which was only affected by incentive level. The aim of the following MEG data analysis was to test whether incentive effect on rest duration was mediated by MBS reduction.

### 3.2.2.2.2 Spatiotemporal characteristics of motor beta synchrony

To be qualified as MBS, our signal had to exhibit three critical features: 1) a dip around movement initiation, 2) in a specific frequency band (13-30 Hz), 3) from a source located over the central sulcus, lateralized with respect to the hand used.

After spectral decomposition of oscillatory activity, the power was normalized by a z-score transformation. This ‘z-power’ was calculated at each frequency by subtracting the mean and dividing it by the standard deviation of the trial baseline, which was defined as the 2-s window preceding incentive display. To verify the scalp topography of MBS at effort initiation, we averaged z-power over time (within 1s centered on effort onset) and frequency (within a range of 13 to 30 Hz). Sensors showing the lowest synchrony level were indeed localized around central brain surface (Figure 33A, right). In order to analyze variations of MBS level, we selected as sensors of interest the 10% sensors with lowest beta synchrony level. We based the selection on pooled data, since the sensors obtained for left-hand or right-hand sessions were almost identical. In other words, including the sensors that were specific to left-hand or right-hand sessions would not make a significant difference. The time course of z-power, averaged across sensors of interest, showed the classical progressive reduction in the beta band, over the 2s preceding effort initiation, dipping just after effort onset (Figure 33A, left).



**Figure 33 Characterization of motor beta synchrony**

*A: Beta synchrony level recorded from central sensors. The time-frequency map shows synchrony level around effort onset, averaged over central sensors (white dots on both hands topography). The mean time series over the beta band (13-30 Hz) is plotted underneath. Z-power means that the power is z-scored relative to the trial baseline.*

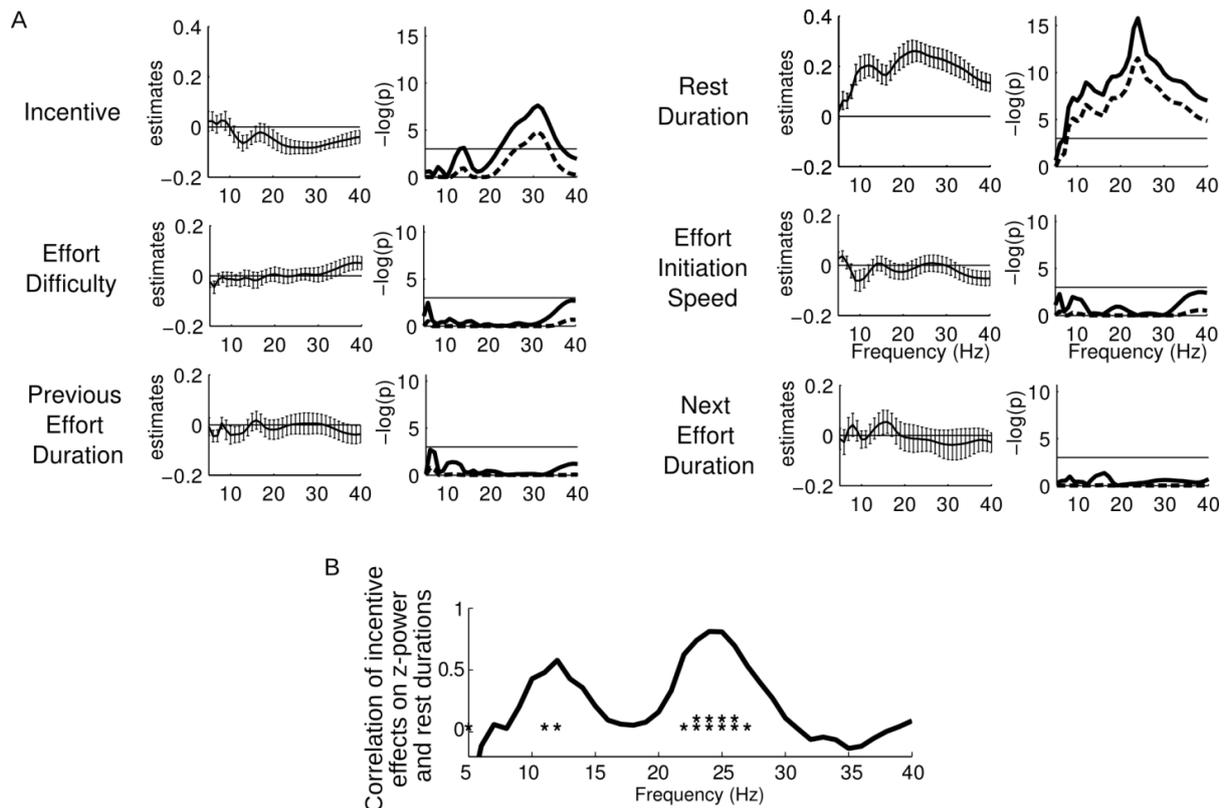
*B: Scalp topography of beta synchrony level. Maps show synchrony level averaged over beta-range frequencies (13 to 30 Hz) in a time window centered on effort onset (-0.5 to 0.5 s), for pooled sessions and for left and right hand sessions, separately. White dots correspond to the 10% sensors with lowest synchrony level.*

*C: Sources of beta synchrony reduction. Coronal and sagittal sections show regions exhibiting the most pronounced reduction of beta synchrony level, within the 2s preceding effort onset relative to baseline, irrespective of the hand used (left panels) and contrasted between hands (right panels).*

To confirm that beta synchrony reduction arose from motor regions, we reconstructed the sources of power in the 13-30 Hz range using a beamformer (see Methods). Contrasting rest to baseline revealed that two main sources underpinned beta power reduction in the left and right sensorimotor cortex (MNI peak coordinates: [30 -30 60] and [-50 -40 50]). When contrasting sessions using left and right hand, a significant asymmetry was observed in favor of the contralateral sensorimotor cortex (see Figure 33C).

#### 3.2.2.2.3 Factors modulating motor beta synchrony

The above results confirmed that our task elicits MBS reduction prior to effort, using classical fix-window analyses. Since we aimed at explaining rest duration in the following analyses, we now considered variable windows: the entire epochs between effort offset and onset. To identify which factors have an impact on MBS, we averaged, for each rest epoch, z-power over time points and sensors of interest, which were selected on the basis of independent criteria (see above). This mean z-power was fitted with a linear model including both manipulated and observed variables: incentive level, difficulty level, duration of the preceding effort, rest duration, initiation speed of the following effort, and duration of the following effort. The model also included factors of no interest that would capture fatigue or adaptation effects at different time scales: the ordinal position of the considered rest period within a trial, that of the trial within a session, and the session number. Regression coefficients were estimated independently for every frequency and their significance was calculated after family-wise error correction for multiple comparisons at the group level (see Methods). Z-power was specifically reduced in the beta band by two factors: higher incentive levels and shorter rest durations (Figure 34A). Among the motor execution parameters (previous and following effort duration, rest duration, effort initiation speed, and effort difficulty), MBS only varied with rest duration. The size of this effect was significantly higher than the others (paired t-test over frequencies and variables, all  $p < 0.05$ , except for 'next effort' at 14-16Hz and 'Difficulty' at 39-40Hz for which there was a weak trend  $p < 0.1$ ). Thus, the regression analysis supports the idea that MBS reduction may mediate incentive effects on a specific movement-related parameter: rest duration.



**Figure 34 Factors affecting motor beta synchrony**

*A: Within-subject effects on z-power (z-scored power change relative to baseline level). Six factors of interest were included in the multiple regression analysis. The coefficients (betas) obtained at each frequency are shown in the left sub-graphs (as inter-subject mean +/- s.e.m.). Their significance was estimated at the group level and plotted as p-value logarithms in the right sub-graphs (plain lines: uncorrected p-values, dashed lines: FWE-corrected p-values). The horizontal line corresponds to  $p=0.05$ .*

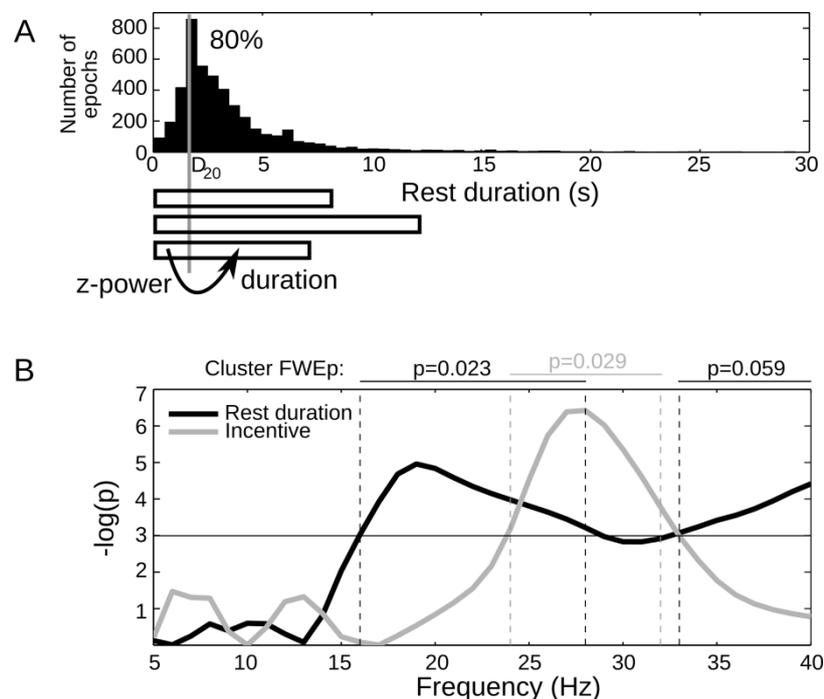
*B: Across-subject correlations between incentive effects on rest duration and z-power (the latter being estimated in the null space of rest duration effect). Spearman correlation coefficients are plotted for each frequency in and around the beta band (5-40 Hz). \* <math><0.05</math> uncorrected p-value, \*\* <math><0.05</math> FWE corrected p-value.*

This idea also predicts that subjects who exhibit strong incentive effects on MBS reduction should as well exhibit strong incentive effects on rest duration. Note that at the subject level, MBS level and rest duration are not independent: this may bias the between-subject correlation between incentive effects on MBS reduction and rest duration. We therefore orthogonalized both variables by regressing out of MBS the linear effect of rest duration prior to estimating the parametric effect of incentive levels. As predicted, the between-subject correlation was positive and significant in the beta range (Figure 34B), surviving family-wise error correction for testing multiple frequencies (see Methods).

#### 3.2.2.2.4 Comparison of causal models linking incentives, motor beta synchrony and behavior

In principle, the statistical dependencies between incentives, MBS, and behavior could result from rest duration mediating incentive effects on MBS. In other words, shorter duration would, as artifact,

reduce measures of MBS level. Even though we do not see what mechanism could support this scenario, we intended to rule this out formally. As consequences cannot precede their causes, we simply tested whether MBS reduction would anticipate rest shortening. Z-power was averaged within a limited time window starting with effort offset and independent from rest duration (see Figure 35A). This early z-power was then regressed against the subsequent rest duration. The end of the time window was set for each subject at the 20<sup>th</sup> percentile of rest durations, to ensure both a sufficient amount of z-power samples (the minimum being 24 samples at 20Hz, corresponding to 1.2s) and a sufficient amount of rest periods included (the 20% rest periods shorter than the time window being excluded). We also included incentive level in the regression model, since it is correlated with rest duration.



**Figure 35 Early motor beta synchrony predicts rest duration**

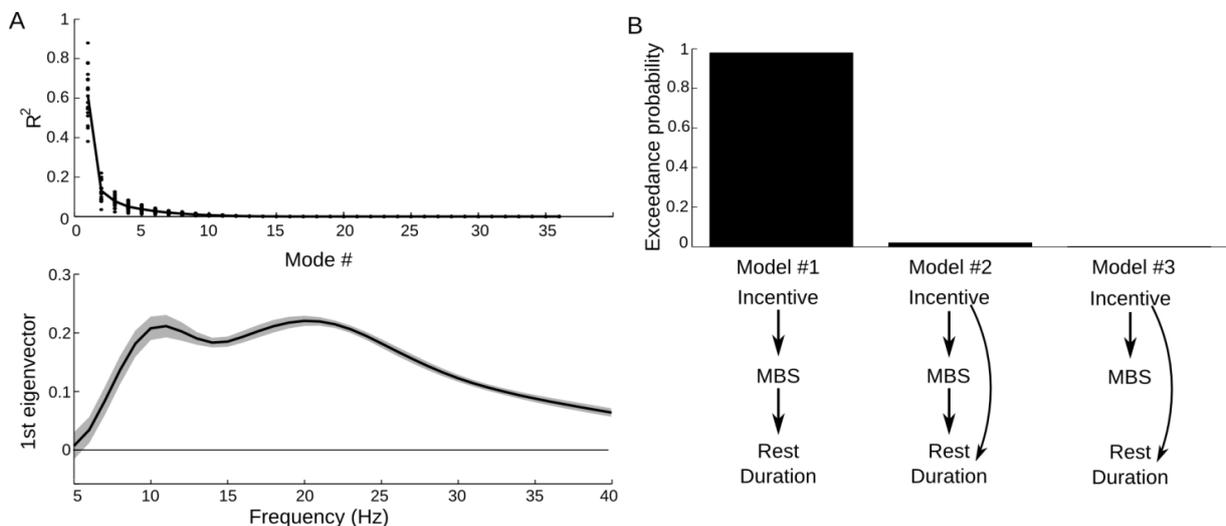
*A: Distribution of rest durations pooled over all participants. The underneath graph depicts the logic of the analysis: for each subject, the 20<sup>th</sup> percentile of rest duration ( $D_{20}$ ) was estimated and early z-power was defined as the mean z-power between rest onset and  $D_{20}$ . Then rest duration (together with incentive level) was regressed against early z-power, across all epochs that lasted longer than  $D_{20}$  (hence representing 80 % of the total).*

*B: Simultaneous regression of rest duration and incentive level against early z-power. Regression significance at the group level is shown as logarithm of uncorrected p-values, for each frequency in and around the beta band (5-40 Hz). Cluster-wise FWE-corrected p-values are also indicated, above the graph.*

The regression was done for each subject and frequency and then tested for significance at the group level (see Figure 35B). Several clusters of frequencies formed at a  $p < 0.05$  uncorrected threshold survived FWE correction, revealing that z-power in the beta range was correlated positively

with rest duration, and negatively with incentive level. Thus, the early MBS reduction was enhanced by incentive level and predicted the upcoming rest duration. Therefore, it seems unlikely that shorter rest duration may be the cause of MBS reduction.

However, the causal links between incentives, MBS, and behavior remain to be specified. Our hypothesis posits that incentives impact MBS reduction, which in turn controls effort initiation (model #1). A more complicated possibility is that incentives also impact rest duration, independently from their effect on MBS (model #2). A last alternative is that incentive level is a common cause of both MBS reduction and rest duration, which would induce a spurious correlation between MBS level and rest duration (model #3). To reduce dimensionality of z-power, we ran a singular value decomposition (SVD). Results showed that z-power could reasonably be reduced to its first mode, which captured most of the variance (Figure 36A). In all subjects, the first eigenvector mirrored closely the MBS pattern over frequencies at effort initiation. For display purpose (Figure 36B), we oriented this first eigenvector such that its mean value in the beta range (13-30Hz) was positive in each subject. Indeed, the orientations of eigenvectors over frequencies and observations (i.e. left and right singular vectors) are arbitrary and depend on each other. Thus, using the first mode of z-power, we had one vector of observations for all three variables (incentive level, MBS reduction, and rest duration).



**Figure 36 Evidence for motor beta synchrony reduction mediating incentive effects on rest duration**

A: Singular value decomposition of z-power variations over epochs, in the 5 to 40 Hz range. Top: Dots represent  $R^2$  statistics obtained for each mode and subject; solid line is the mean over subjects. Bottom: The curve indicates the first eigenvector (inter-subject mean  $\pm$  s.e.m), for each frequency in and around the beta band. Since the direction of eigenvectors is per se arbitrary, we flipped them subject-wise for their mean value over the beta range to be positive, which allows direct visual comparison of the group average with the other figures.

*B: Results of model comparison. The graphs illustrate the 3 models tested to account for statistical dependencies between incentive level, beta z-power and rest duration. Bars indicate model exceedance probability (i.e. the probability that the model is the most frequently implemented in the population).*

Models were specified by linear dependencies between these variables (depicted by arrows in Figure 36C). Given the data, we estimated these models and took their respective evidence to perform Bayesian model selection with a random-effect analysis at the group-level (see Methods). Model #1, with an expected frequency of 0.85, obtained a high exceedance probability ( $x_p = 0.99$ ), which is the confidence that this model is more frequently implemented than the two other models in the general population.

### **3.2.2.3 Discussion**

In the implicit effort allocation task, participants were not instructed when to start exerting physical effort. Instead, they were motivated by monetary incentives to spend more time working. They spontaneously adjusted effort allocation to these incentives, trading off benefits against costs. In particular, they shortened breaks when work paid more. We found evidence that such an effect of incentive motivation was underpinned by a reduction of motor beta synchronization (MBS), relative to trial baseline. Indeed, MBS reduction was correlated across trials with both incentive level and rest duration. In addition, subjects who exhibited stronger incentive effects on rest duration also exhibited stronger incentive effects on MBS reduction. Finally, direct Bayesian model comparison suggested that the most likely interpretation of statistical dependencies between our three variables of interest is that incentive effect on rest duration was mediated by the amplitude of MBS reduction. In the following, we first discuss the modulation of rest duration by MBS reduction, and then the modulation of MBS reduction by incentives.

Effort onset could be predicted by MBS reduction measured in a fixed time window at the beginning of rest periods. This observation discards the possibility that rest duration per se may bias MBS measurements. On the contrary, it suggests that MBS reduction favors the initiation of effort production. This is in line with the general idea that a high MBS level represents an 'idling rhythm' maintaining the motor status quo and that decreasing MBS allows for a motor change (Engel and Fries, 2010). More precisely, the corticospinal pathway might be less excitable during high MBS states, preventing any motor program from triggering movement initiation (Schoffelen et al., 2005). Interestingly, we observed that MBS reduction only impacted initiation time and no other effort-related parameter, such as speed or duration. This specificity echoes numerous reports that MBS reduction observed before the action onset is not linked to any movement parameter (van Wijk et al., 2012).

Growing evidence suggests that MBS reduction indeed plays on movement initiation. In PD patients, a higher MBS level in the sub thalamic nucleus was correlated across trials with longer initiation delay (Kühn et al., 2004), and with successful inhibition of the prepotent response in a Stroop task (Swann et al., 2009; Brittain et al., 2012). In healthy participants, faster finger tapping (with reduced intervals between movements) resulted in a lower MBS level (Toma et al., 2002), but prolonged movement duration (with constant intervals between movements) did not (Cassim et al., 2000), suggesting that MBS is specifically modulated by rest (not effort) duration. Interestingly, increasing response uncertainty (by augmenting the number of possible movements) also enhances both reaction time and MBS level (Tzagarakis et al., 2010). Accordingly, in a visual detection task, the progressive reduction of MBS correlated with the gradual commitment to a motor response, which was distinct from the confidence in the perceptual decision (Donner et al., 2009; O'Connell et al., 2012).

We should acknowledge that the MBS is not the only motor rhythms involved in the regulation of movement production. There is now growing evidence that there is a balance between pro-kinetic and anti-kinetic oscillations: the gamma (>40Hz) and the beta rhythms, respectively, with opposite effect of corticospinal excitability (Schoffelen et al., 2005). In Parkinson disease, therapies and improvement of the motor symptoms are correlated to opposite effect in the beta and gamma band (Brown et al., 2001). High gamma frequencies recorded in the sub-thalamic nucleus of implanted patients (200 Hz) were also reported to correlate with apathy score in Parkinson disease (Özkurt et al., 2011). We did not investigate gamma rhythm in our data because the task was physically strenuous, so that the data were heavily contaminated by muscular artifacts in higher frequency bands. We tried to minimize these artifacts during the recording with repeated instructions to use only the forearm to squeeze the grip while leaving the shoulder as steady as possible and by sitting the subjects comfortably to avoid contractions in the neck. However due to the nature of the task, muscular artifacts are unavoidable and prevented the investigation of higher frequency bands. The mu rhythm (around 10 Hz) is also reported to co-vary with the motor beta rhythm. This mu rhythms could be relate to somatosensory processing (Jones et al., 2009). We observed mu desynchronization in our data, and similar, though far less significant, effects of rest duration and incentive.

The specific link of pre-effort MBS to initiation time should be contrasted to the case of readiness potential or field (RP/RF), which also manifests as a slow ramping signal that precedes voluntary movement (Pedersen et al., 1998; Praamstra et al., 1999; Leuthold and Jentzsch, 2002; Shibasaki and Hallett, 2006). Many movement-related factors affect the RP/RF such as the force load, the effector used and the movement complexity (Lang, 2003). Thus, while MBS reduction may reflect motor gating in general (Engel and Fries, 2010), with a main source in the contralateral motor cortex

(Jurkiewicz et al., 2006; Donner et al., 2009; Tzagarakis et al., 2010), the RP/RF seems to reflect the preparation of a specific motor program (Shibasaki and Hallett, 2006), with main sources in the supplementary motor area in addition to the primary motor cortex (Ball et al., 1999; Cunnington et al., 2005). Such a clear-cut distinction between MBS-motor gating and RP/RF-motor preparation should however be tempered: some authors argue that the RP denotes the transition from intention to action (Lang, 2003), and others have recently proposed that the RP reflects the passive stochastic accumulation of a 'go' signal (Schurger et al., 2012).

Little is known about how these signals relate to reward processing, which has been overlooked by the EEG-MEG literature until recently. On the one hand, MBS reduction was characterized as a gating signal in the domain of motor control, the movement being directly instructed or related to a perceptual decision, without bearing any particular value for the subject. On the other hand, MBS was characterized in the domain of motor disorders as a pathological signal, which should be eliminated in order to alleviate symptoms such as hypokinesia in Parkinson's disease (PD). To our knowledge, the intuitive idea that MBS could represent a normal process adjusting motor behavior to subjective goals in healthy conditions has not been directly investigated. Here we provide evidence that MBS reduction may speed up effort initiation proportional to expected rewards in healthy subjects. Other neural mechanisms have been suggested for underlying such incentive motivation processes. For instance, reward representation may influence motor output through cortico-cortical connections, implementing a top-down regulation of behavior (Locke and Braver, 2008; Kounieher et al., 2009). Another possibility is that the interaction between reward and motor circuits occurs within the basal ganglia, with the ventral parts boosting the dorsal parts (Knutson et al., 2008; Schmidt et al., 2012; Tachibana and Hikosaka, 2012). An alternative suggestion is that dopamine release facilitates the expression of motor programs, either at the cortical or sub-cortical level (Berridge, 2004; Robbins, 2007; Salamone and Correa, 2012).

These possibilities are not mutually exclusive and could be articulated with the phenomenon of MBS reduction. It is known that degeneration of dopaminergic neurons in animal models of PD, as well as in human patients, results in abnormally high beta oscillations that can be reduced by dopamine replacement medications (Schnitzler and Gross, 2005; Uhlhaas and Singer, 2006; Hammond et al., 2007). This is possibly because dopamine release in the striatum helps filter cortical input and desynchronize basal ganglia output signals (McIntyre and Hahn, 2010). Thus, one could speculate that reward prospects represented in limbic circuits may amplify dopamine release, which in turn may facilitate movement initiation by lowering beta oscillations in the motor circuits. This view would invite reconsidering the status of bradykinesia or akinesia as motor symptoms. They would

instead represent dysfunction of motivational processes that are in between pure reward and motor representations. Several computational accounts of dopamine depletion support this interpretation, as they attributed delay and slowness to a shift in the movement cost/benefit ratio, rather than to sub-optimal control of the movement spatiotemporal trajectory (Mazzoni et al., 2007; Niv et al., 2007; Baraduc et al., 2013). Yet the link between dopamine release and MBS reduction in healthy conditions remains to be established.

Thus, we conclude that MBS reduction may represent a neural process translating expected reward into motor activation. There are however some limitations that should be acknowledged. First, our conclusion is based on statistical dependencies between variables, which suggest - but not prove - a causal pathway from incentive level to MBS reduction to effort initiation. Direct manipulation of MBS level, through dopaminergic medication or electrical stimulation, could provide more conclusive evidence for causality, by affecting subjects' sensitivity to incentives, and perhaps patients' apathetic symptoms. Second, the frequencies that were correlated to incentive level and rest duration were slightly different, even if they could all be labeled as 'beta'. It remains to be understood whether these differences in frequency are functionally significant for the incentive motivation process.

#### **3.2.2.4 Methods**

Details on the experimental setting and the behavioral task are reported in the behavioral section (see page 116).

##### **3.2.2.4.1 Subjects**

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. Subjects (n=19, 8 males, 24.9±0.7 years) were recruited within an academic database via email, and gave informed consent prior to participating in the study. They were right-handed, between 20 and 30 years old, free from magnetic artifacts, and with normal vision and no history of neurological or psychiatric disease. They believed that the money won while performing the task would be their remuneration for participating, but their payoff was eventually rounded up to a fixed amount (100€).

##### **3.2.2.4.2 MEG data acquisition**

A whole-head MEG system with 151 axial gradiometers (CTF Systems) was used to sample brain activity at 1250 Hz with an online low-pass filter of 300 Hz. Two bad sensors (MLT42 & MRT32) were excluded because of high noise levels. Head position was determined using marker coils at fiducial points (nasion, left and right ears) before each session. The first session served as a reference to control that head displacement in the 7 remaining sessions never exceeded 5 mm. The first two sessions were excluded from the analyses in two participants, due to excessive change in head

position. Electromyograms were recorded simultaneously with two pairs of disposable surface electrodes on each hand, which were placed to target the flexor digitorum superficialis (on the forearm) and the first dorsal interossei (between thumb and index).

The detection of the effort onset and onset was based on visual inspection of the EMG trace after 100 Hz high pass filter, with the help of the force level trace (low pass filtered at 20 Hz) to ensure a precise detection while avoiding false detection.

#### 3.2.2.4.3 MEG spectral decomposition

Data was imported into Matlab and analyzed using Fieldtrip toolbox (<http://fieldtrip.fcdonders.nl>, (Oostenveld et al., 2011)). For each session (lasting around 320s), the whole data set was decomposed into power overtime and frequency, using a product with a set of Morlet wavelets after fast-Fourier transform. A product in the frequency space is equivalent, but computationally faster, to a convolution in the time space. The Morlet wavelets trade temporal against spectral resolution ( $\sigma_t$  against  $\sigma_f$ ), such that  $\sigma_f\sigma_t=1/(2\pi)$ , and scale this tradeoff according to the frequency  $f$ , such that  $f/\sigma_f$  is constant. This ensures that finer temporal resolution is achieved for higher frequencies at the expense of a lower spectral resolution. We set the  $f/\sigma_f$  ratio to 7, which is standard for the frequency range investigated here. For each unit frequency between 5 and 40 Hz, the product between the wavelets and the data was computed for every 50-ms step. This is well below the original sampling rate (1250Hz) but sufficient considering the time scale of the process investigated (seconds). The length of each wavelet used for the computation was 3 times its temporal resolution ( $\sigma_t$ ). These power data were used for every analysis, except the source reconstruction.

#### 3.2.2.4.4 MEG source reconstruction

Templates of brain anatomy (single\_subj\_T1.nii) and meshes of cortical surface and head envelope (cortex\_5124.surf.gii) were taken from SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>). For every subject, sensor positions were co-registered in the MNI space using the mean fiducial positions over sessions as landmarks. The normalized lead field was estimated using a single shell head model in a 10x10x10mm grid of sources covering the entire brain. Time series were epoched into pairs of temporal windows, first with the 2s of rest preceding effort onset and second with the 2s of baseline preceding the corresponding trial. All epochs were analyzed together in the 13-30 Hz frequency domain, using multi-tappers to compute the cross-spectrum matrix between sensor pairs (in which the diagonal corresponds to the power spectrum at each sensor). A spatial filter was jointly computed for rest and baseline data, using DICS beamformer (Dynamic Imaging of Coherent Sources (Gross et al., 2001)), without regularization of the solution. Then, we projected the signal of each epoch separately through the filter to estimate the power in the source space, over the 13-30 Hz

range. A common filter allows contrasting rest and baseline power levels, hence computing the percentage of signal change. The group mean was interpolated onto the anatomical template for display.

#### 3.2.2.4.5 MEG statistical analysis

Multiple regressions of power level against the various factors of interest (Figure 34A) were estimated using Matlab statistical toolbox. The significance of all regression coefficients was estimated at the group level using a non-parametric procedure. To estimate uncorrected p-values, the null distribution of t-values was estimated by flipping the sign of regression coefficients over participants. To estimate family-wise error (FWE) corrected p-values, the null distribution of the maximal t-value was estimated in a similar fashion (Nichols and Hayasaka, 2003). In both cases, 200 000 distinct sign changes were used among the  $19^2 = 524\ 288$  possible changes.

For between-subject correlations (Figure 34B), uncorrected p-values were calculated from Spearman correlations, and FWE-corrected p-values were derived from Holm's step-down adaptive method, which strongly controls FWE without any assumption (Nichols and Hayasaka, 2003).

Cluster statistics in Figure 35B were calculated using a cluster-mass permutation scheme described in (Maris and Oostenveld, 2007). First, to determine uncorrected statistical threshold ( $T_{\text{thd}}$ ) corresponding to  $p=0.05$ , we used the same randomization procedure as described above for multiple regression analyses. These non-parametric thresholds (incentive:  $T_{\text{thd}}=2.092$ ; rest duration:  $T_{\text{thd}}=2.088$ ) were very close to their parametric counterpart ( $t_{0.975,18}=2.101$ ). Second, cluster-mass FWE-corrected p-values were estimated from the null distribution of maximal cluster mass (sum of cluster t-values) formed at  $T_{\text{thd}}$ . This distribution was approximated using 200 000 distinct randomizations (again among 524 288) of regression coefficient signs over participants. All p-values reported throughout this study correspond to bilateral tests.

#### 3.2.2.4.6 Bayesian model selection

See page 108.

### 3.2.3 Adaptation of the body activation system

#### 3.2.3.1 Introduction

The previous results showed that the behavior and the dynamics of brain signals that regulate behavior such as the cost-evidence signal or the motor beta synchrony were strongly modulated by the incentive. The fMRI data revealed a transient incentive signal in the striatum when the incentive level was revealed as a coin image (these data are not shown above) but there was no tonic signal reflecting the incentive level during the trial and there were neither phasic signals at behavioral

transitions between effort and rest. It is tempting to suggest from these observations that the incentive level is encoded early in the trial and then used to adjust the dynamics of cost-evidence and motor beta synchrony, which impacts the timing of effort and rest onset. But this view emphasizes the following contrast: on the one hand, the pervasiveness of incentive effect on several brain dynamics and on the other hand, the absence of a sustained physiological signature that would capture the global behavioral activation induced by incentives. Here, the aim is to find a marker of this global energizing effect of incentives. A good candidate is the autonomic signal that plays a key role in the bodily activation.

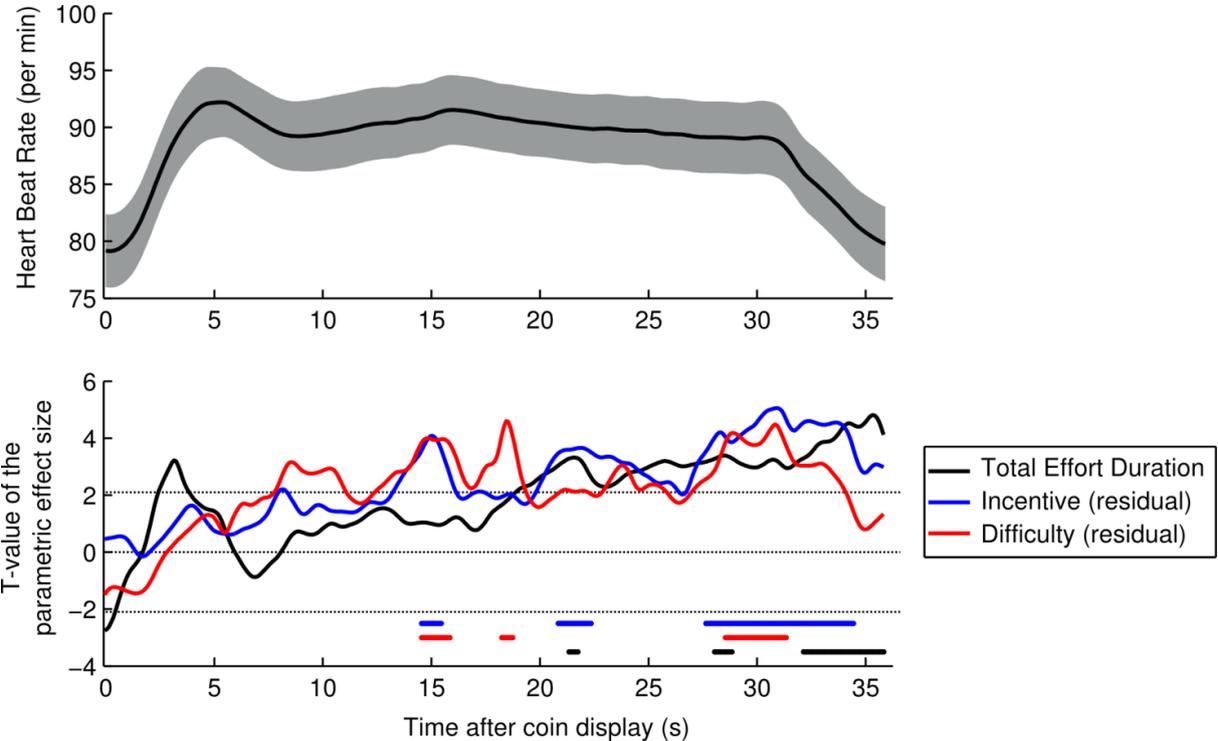
The autonomic system can be subdivided based on anatomical, physiological and functional evidence into the parasympathetic and the sympathetic systems (Iversen et al., 2000). In particular, these sub-systems rely on distinct neurotransmitters, respectively acetylcholine and noradrenaline. Both can be viewed as opponent systems, the sympathetic system up-regulating the behavioral activation whereas the parasympathetic system promoting housekeeping processes such as digestion and rest. The balance between these two systems controls, for instance, the heart beat rate (HBR) and the pupil diameter, such that reactive increases of HBR or the pupil diameter are under sympathetic control (Iversen et al., 2000). The skin conductance is more uniquely related to sympathetic control, and was extensively studied with the polygraph (Handler et al., 2010) and in cognitive science, for instance in decision making (Figner and Murphy, 2010). Overall, the skin conductance response, the HBR and the pupil diameter are often used as proxy for the activation level. For instance, the skin conductance response was increased for higher incentives in a task in which the reward prospect motivated hand grip effort production, when the incentive display was either supraliminal or subliminal (Pessiglione et al., 2007). Similar results were found for emotional arousal (Schmidt et al., 2009). However, it is important to note that this incentive modulation of the autonomic system is not necessarily a strict reflection of an increased behavioral output: they can be disconnected. In particular, in patients with focal striato-pallidal lesions, the autonomic modulation by the incentives was preserved although higher incentives were not translated into higher behavioral activation: force production was not impacted by the incentive level, as opposed to control subjects (Schmidt et al., 2008).

For the present study, we used the heart beat rate (HBR). The electrocardiogram was recorded during the MEG experiment introduced earlier. The HBR estimation is better suited to the task. Skin conductance is difficult to record in the effort allocation task because both hands are used alternatively. It is possible to record from the foot, however palm is preferred over foot for

methodological reasons (Handler et al., 2010). Pupil diameter is a discontinuous measure, interrupted by eye-blink which makes this analysis during the trial (30s) difficult.

**3.2.3.2 Results**

First, the heart beat rate (HBR) was estimated over the course of each trial (see Figure 37, top graph). The HBR increased immediately after the display of the incentive, reaching a peak after 5.4s. The median effort onset timing was on average 2.14s +/- 0.18 s.e.m. across subjects, the peak of HBR therefore followed effort initiation. After this peak, levels of HBR were maintained high until the end of the trial and then they abruptly dropped during the feedback and intertrial interval.



*Figure 37 Within-subject variations of heart beat rate*  
 The top graph shows the mean heart beat rate, with the shading corresponding to the s.e.m. across subjects. Time is aligned to the display of the incentive at t=0s, the period for effort allocation starts at t=1s and lasts 30s, the trial feedback is displayed at t=31s and is followed by the inter-trial empty screen at t=33s, and until t=36s.  
 The bottom graph shows variations around the mean heart beat rate explained by several factors in a linear model. The heart beat rate was regressed across trials against the cumulated effort duration rewarded in the trial (black) and the residuals of this effect were regressed against the incentive level (blue) and the difficulty level (red). T-values of the parameter estimates across participants are plotted; the dotted lines at +/- 2.1 correspond to a p-value of 0.05 uncorrected. The dots at the bottom indicate when the significance of parameters survives FWE correction < 0.05, to correct for multiple comparisons across time points within the 36s.

Given that the HBR time series closely matched the time course of the trial, it is likely that it was relevant for the behavior. To test this, HBR variations around the mean value were regressed across trials at every time point against the performance (total effort duration rewarded in the trial; see Figure 37, bottom graph, black line). The HBR variations were significantly related to the behavioral performance; this effect tended to build up over time so that the HBR was better explained by the performance at the end of the trial. A first peak of the effect of performance was found early in the trial. Post-hoc analysis revealed that this early effect was due to the fact that when participants performed better they tended to quicken effort initiation at the trial onset so that the first peak in the HBR time series is shifted to shorter latencies.

The behavioral analysis (see Figure 32) revealed that the performance, i.e. the total effort duration rewarded in the trial, increased with the incentive level and decreased with the difficulty. It is therefore likely that the HBR is correlated to the incentive and difficulty levels, since it is correlated to the performance that co-varies with these factors. Less trivial and more interesting is the question whether the HBR co-varies with the incentive and the difficulty levels on top of the performance effect. In other words: is the HBR impacted by the incentive and the difficulty levels when the performance is kept constant? To address this question, the incentive and difficulty effects were estimated after regressing out the total effort duration effect (see Figure 37, bottom graph, blue and red lines). The HBR was increased by higher difficulty levels, which could reflect that, when the effort duration is constant, higher difficulty levels make the exertion more demanding. The HBR was also increased by higher incentive levels. Note that when the effort duration rewarded is kept constant, trials with higher incentives are associated to higher payoffs. Therefore, for a given motor production, the HBR was increased when it was more motivating to exert effort, which could reflect an enhancement of the arousal and the activation system.

### **3.2.3.3 Discussion**

The results suggest that the HBR is a good marker of the behavioral activation in the sense that it reflects both higher behavioral outputs and higher incentive. First, HBR variations over time closely matched the task structure: HBR was increased during the trials as compared to the inter-trial intervals. Second, higher HBR correlated with better performance between trials in the task, i.e. longer cumulated effort duration. Third, when the performance effect was explained away, higher difficulty levels correlated to higher HBR, which corresponds to the intensity of the body activation. This increase of HBR with higher effort intensity was also observed for higher incentive levels.

Such modulation of the HBR by higher incentives (on top of the behavioral output) is classically seen as a marker of the motivation arousal and is physiologically accounted for by sympathetic activation

(Brehm and Self, 1989). This cardiovascular effect of motivation arousal does not depend on the physical nature of the effort produced. There are many reports of such effects in mental tasks such as copying. It is proposed that this activation is behaviorally relevant because the incentive effect is all the more pronounced that efforts are required: the incentive effect on HBR is amplified by the task difficulty. This effect could be mediated by beta-adrenergic activity (Richter and Gendolla, 2009). Such a view is comforted by the present results: the HBR was increased for higher incentive levels and higher behavioral outputs.

An interesting question is the direction of the causation: does the HBR activation lead to more physical efforts or does it result from higher efforts being produced? Since higher sympathetic activations also occur in mental task (Brehm and Self, 1989; Richter and Gendolla, 2009) and can be disconnected from the physical effort production (Schmidt et al., 2008), it is more likely that the increase in HBR precedes the increase in effort production. This causal relation could have been comforted with temporal precedence here; however the correlation between HBR and trial performance or incentive level, despite occurring early in the trial did not precede the latency of the first effort.

To support the claim that there is an incentive effect on top of the behavioral output, the effort duration produced was regressed out and the effort difficulty was included in a linear model. This raises two questions: if the incentive effect modulates the behavioral performance through the sympathetic system, how does the incentive effect on HBR survive after regressing out the behavioral performance? If the incentive effect truly precedes causally the behavioral effect, the effect might survive through partial de-correlation due to noise in the process. Another issue is that it cannot be ruled out that the incentive effect on HBR was not mediated by an incentive effect on behavior, because not all behavioral parameters were included in the analysis, such as the vigor of effort initiation (how steep is the force increase at the effort onset), the variability of the force produced, the pace of the effort production, etc. This is a limitation to the analysis.

Instead of testing the HBR as a marker of the behavioral activation, we could also have tested the HBR as a predictor of the behavior, in particular the timing of effort onset and offset. There is strong evidence that the level of perceived exhaustion is related to the HBR (Williamson, 2010). This relationship motivated Borg to align his exhaustion rating scale on the HBR (Borg, 1982, 1990). Assuming that the perceived exhaustion is related to cost-evidence, HBR could be a proxy of cost-evidence, waxing during effort and waning during rest. This is a perspective on the present analysis, its potential should however be tempered. First, the link between perceived exhaustion and the cost-

evidence level was not convincing (see the introspection study, page 128). Second, the temporal resolution of HBR may not be sufficient to match effort, and more critically, rest durations.

Finally, the idea that increased HBR for higher incentives was expected because motivational arousal enhances the sympathetic activation could be challenged. It was suggested previously that a potential mechanism for the brain to regulate cost-evidence could be mediated by the opioid system, like in the placebo effect. In this case, higher incentives would increase the opioid level to allow more effort production. This opioid-mediated placebo mechanism decreases the HBR (Benedetti, 2008): this reduction is blocked by naloxone, an opioid antagonist. Therefore in principle, it would have been possible that higher incentives reduce HBR, although such an effect could be overridden by the sympathetic effect.

### **3.2.3.4 Methods**

#### **3.2.3.4.1 Participants, Task & Payoff**

The electrocardiogram was recorded during the MEG study. See previous sections for details on the participants, the task and the behavioral effects.

#### **3.2.3.4.2 Data collection and pre-processing**

The electrocardiogram was recorded with disposable adhesive electrodes placed below the upper left clavicle and the bottom right of the abdomen, so that the heart is on a line between these two electrodes. The ground electrode was placed on the back of the left shoulder. The leads were connected to the magneto-encephalography system and acquired at the same frequency sampling (1250 Hz). The data were pre-processed off line to identify the R peak of the cardiac complex using a script of mine in Matlab using filtering tools from the signal processing toolbox (The MathWorks, Inc.). This script works as follows for each subject:

- 1) filtering: high-pass filter at 0.1 Hz and low pass filter at 60 Hz, using a 4-th order Butterworth filter (which ensures a flat response in the desired band and minimal changes elsewhere), that were applied twice in opposite direction to avoid phase shifts.
- 2) A first step identified R peaks based on a peak-above-threshold detection. Since the waveform could be locally distorted and the amplitude of the R peak was not constant, this detection was not perfect (misses and false alarms).
- 3) The R-peaks from the first steps served to epoch the signal around R-peaks. The epochs comprised 250 ms before and 400 ms after the R peak. Epochs entered a Principal Component Analysis, and the first component was used as the canonical template of the cardiac complex. This template was used in a template match procedure by convolving the

signal and the template, so that R peaks could be reliably detected as peaks of correlation between the template and the signal.

- 4) The heart beat rate was then computed at each R peak location by taking the R-R interval with the following peak.
- 5) For convenience, the heart beat rate was resampled with spline-interpolation at 10 Hz.

#### 3.2.3.4.3 Statistical analysis

Linear models were used for the analysis. Models comprised a constant to capture the mean and z-scored regressors to explain variations around the mean between trials. Note that the cumulated rewarded effort duration (called performance for the sake of brevity) was impacted by fatigue: values tended to be higher at the beginning of the experiment. So that the performance regressor can capture local variations in performance between trials and not a global fatigue trend, performance was 'de-trended', i.e. the mean of each session and linear trends over the trials of each session were removed prior to z-scoring the regressor.

The regression analysis was repeated on every time step of the data, which results in a time series of parametric effect estimates. To correct for multiple comparison over time, a family-wise error correction was estimated using a non-parametric procedure. Description of this procedure can be found in the method section page 170. Here,  $n=10\,000$  permutations were used to estimate the null distribution controlled for family-wise error and the bilateral threshold was determined to control the Type I error at the 0.05 level, i.e.  $FWE\ p < 0.05$ .

By design, the performance, the difficulty and the incentive regressors correspond to local variations between trials and are not confounded with fatigue effects. To be sure that the results presented cannot be attributed to fatigue effects captured differentially by the regressors, the same linear analyses were estimated after including a constant per session and linear trends over trials within each session as covariates. This did not alter the results; on the contrary, it tended to improve the significance levels.

### 3.3 Manipulation of the brain during effort allocation

The last part of this experimental work aimed at introducing perturbations in the brain during the effort allocation process to test our cost-evidence model. This section starts with a simple question: does nociception play any role in the effort allocation process? This addresses the nature of the cost-evidence signal encoded in the brain. We tested the relation to nociception in a first study with a pharmacological pain killer, and in a second study with a more original way of manipulating nociception and proprioception: the hypnotic suggestion. We extended the hypnotic manipulation of cost to a manipulation of benefit with a pro-motivational hypnotic suggestion in a third study. Unfortunately, these three studies yielded negative results. Building on negative results is flawed since there are too many reasons why it should not work. However, these studies were motivated by a principled approach and performed with care, thus I chose to include them in this dissertation with brevity. The last study aimed at testing the role of serotonin in the effort allocation process and yielded more interesting results.

#### 3.3.1 Manipulation of nociception with paracetamol

##### 3.3.1.1 Introduction

The cost-evidence model support that effort-related cost is the key variable to track on-line to guide the effort allocation based on a comparison with the potential benefit. The general introduction and the previous results suggested that effort-related costs could be signaled through proprioception and nociception, although non-proprioceptive and non-nociceptive estimation may also be at play. It is tempting to test experimentally the role of nociception in the effort allocation process. There is a large and ancient pharmacopeia that precisely aims at reducing nociception: pain killers are an opportunity to test whether nociception is involved in the effort allocation process.

We chose one of the most common pain killers: paracetamol (a.k.a acetaminophen). Paracetamol is not derived from morphine: it is thus a class I analgic. It has less anti-inflammatory effects compared to the so-called non-steroidal anti-inflammatory drugs, like aspirin: its effect is comparatively more central. The analgic effect of paracetamol is not pure but also antipyretic (as opposed to nefopam). Despite it has been used for more than a century, the mechanisms underpinning the analgic effect of paracetamol are still not perfectly understood; the reason is that there is a complex cascade of effects and many systems are impacted by paracetamol. Paracetamol is a potent inhibitor of the cyclooxygenase (COX), regulating prostaglandins and the eicosanoid system (Toms et al., 2008). This inflammatory process is however not believed to account for much of the analgic properties. More likely, the analgic effect is mediated by the serotonergic descending pathway, involving in

particular 5HT1A receptors at the spinal level, but other systems like the opioid system and the cannabinoid system may also be at play (Smith, 2009).

The use of paracetamol is therefore not motivated by the target of a precise system but rather because it is well tolerated: beside rare cases of allergy, there are no adverse effects in healthy people (Moore et al., 2000; Toms et al., 2008). In the implicit effort allocation task, paracetamol could alleviate the impact of physical cost, which should primarily lengthen the effort duration and likely shallow the effect of the difficulty on effort durations. Given that effects on rest duration may be mediated by expectations (as suggested in the behavioral section), they may not be affected by paracetamol, but if any effect, it should be in the direction of a faster recovery.

### 3.3.1.2 Results

The study with Paracetamol replicated the main results found in the fMRI & MEG group. When data were pooled irrespective of the treatment (placebo or paracetamol) or visit number (first or second visit), subjects increased effort duration for higher incentives ( $p=2.7 \cdot 10^{-4}$ ), reduced effort duration when more difficult ( $p=3.7 \cdot 10^{-10}$ ), reduced rest duration for higher incentives ( $p=0.037$ ) and did not change rest duration whatever the difficulty ( $p=0.83$ ).

The treatment was given before the estimation of the maximal force; we thus checked whether maximal force differed between treatments. Using paired comparisons over visits and participants, we found a non-significant increase of maximal force under Paracetamol (30.2 a.u. +/- 38.4 % CI,  $p=0.12$ ). The maximal force tended to decrease from the first to the second visit (-33.5 a.u. +/- 38.6 % CI,  $p=0.08$ ). The visit effect was counterbalanced relatively to the treatment over participants. Note that the treatment and visit interaction cannot be estimated in this cross-over design, unless the participant grouping is treated as a fixed effect (and hence, that observations between visits are not paired over subjects).

The payoff did not differ between treatments (Paracetamol compared to Placebo: -0.32€ +/- 1.01 % CI,  $p=0.51$ ). However, it increased over visits (second minus first visit: 1.61 € +/- 0.51 % CI,  $p=6 \cdot 10^{-6}$ ). To check whether the payoff difference between visits was due to the reduction of maximal force, it was regressed simultaneously against a constant (i.e. the effect of visit), the treatment order and the difference in maximal force. Lower maximal force produced during the second visit increased the payoff significantly ( $p=5.4 \cdot 10^{-4}$ ), but the effect of visit remained significant ( $p=8 \cdot 10^{-6}$ ).

As both effort and rest durations determine the payoff, the absence of payoff difference between treatments does not imply that there is no difference in the effort and rest durations, or their modulations by the experimental factors. All these variables were compared using paired differences

(see Figure 38). Participants did not change significantly their mean effort duration or their modulations by incentive and difficulty under Paracetamol (all  $p > 0.56$ ). Fatigue effects were not impacted either at any time scale: session, trial, effort position within a trial (all  $p > 0.82$ ). On the contrary, participant increased rest duration under Paracetamol ( $p = 0.05$ ). However, modulations of rest duration by the incentive ( $p = 0.52$ ), by the difficulty (0.45) or by any timescales ( $p > 0.23$ ) were not significant.

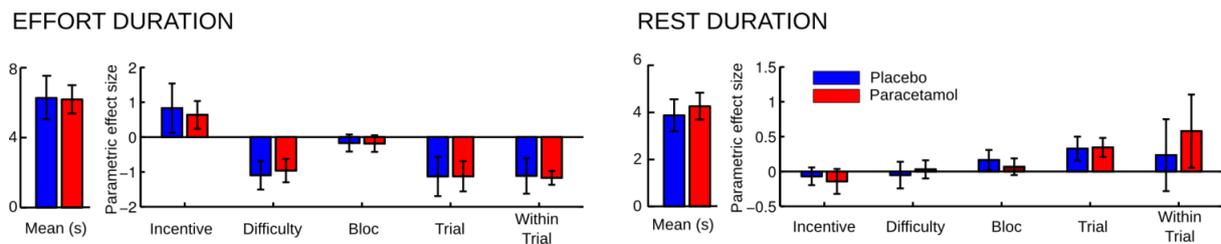


Figure 38 *Paracetamol vs. Placebo compared across visits: effort, rest durations and their modulations.*

Each color represents the data of a given treatment, that were varied for each subject across visits: The bars represent the effect size of regressors estimated from a linear model that include the mean and parametric modulations around the mean by the incentive level, the difficulty level, the position of rest or effort epochs between blocs (1 to 4) trials of a bloc (1 to 9) and within trials. Error-bars are 5% confidence interval across subjects. Means and effect sizes are plotted separately to provide a more convenient scaling.

On the contrary, variations between visits were more systematic. Participants increased their effort duration ( $p = 0.01$ ), they steepened the effect of the incentive of effort duration ( $p = 0.003$ ) and they smoothed the effect of the difficulty ( $p = 0.01$ ). Mean rest duration did not change ( $p = 0.63$ ), there was however a trend to increase the negative effect of the incentive ( $p = 0.08$ ) and increase positive modulation by the difficulty ( $p = 0.02$ ).

### 3.3.1.3 Discussion

To summarize, the predicted effects of paracetamol in the implicit effort allocation task was that the effect of physical cost could be alleviated. None of the predictions entailed by this hypothesis were met; the only marginally significant effect was actually the opposite: rest duration was lengthened under paracetamol.

If the increase of rest duration were to be taken seriously, how could it be explained? Drowsiness could increase rest duration, but there is no clinical evidence for such an adverse effect of paracetamol (Moore et al., 2000; Toms et al., 2008). Alternatively, participants under paracetamol could be less sensitive to the opportunity cost of rest, but to be consistent, this should have also produced a shallower incentive effect on rest duration, which was not the case. Does paracetamol have effects that are unrelated to physical pain? It was reported for social pain (Dewall et al., 2010),

however in this study, the effect was only seen after eleven day of protracted treatment. It was proposed that pain is characterized by at least two aspects: a sensory dimension, processed by the sensory cortices and the posterior insula, and the affective dimension, processed by the anterior insula and the dorsal anterior cingulate cortex (Eisenberger, 2012). The affective dimension would be a common signature to physical and social pain (Eisenberger et al., 2003) alleviated by pain killers.

The effects predicted after the hypothesis that cost-evidence depends on nociception were not observed. This could mean that the hypothesis is false, although it cannot be proved from a negative result. However, I should acknowledge a limitation in this study: we did not estimate whether the paracetamol treatment reduced pain sensation in our subjects with an independent test. It is possible that the treatment did not alter significantly pain sensation; in this case it is no clear that such a treatment provides a valid test for the role of nociception in the effort allocation task.

What follow-up studies or perspectives could be proposed after these negative results? I briefly discuss the use of opioids, which is a project initiated in the team, and the investigation of the placebo effect.

Opioid-based drugs are more potent pain killers than paracetamol (Staahl et al., 2009), and they have a more focused pharmacological effect. However, opioids have more adverse effects and must be used with caution and interpreted with care: the potential effects could dependent on side effects instead of the analgic effect. Note also that the analgic effect itself is ambiguous as it impacts both the sensory and affective aspect of pain (Zubieta et al., 2001, 2005). The relief of social pain by opioids could be related to this affective aspect (Eisenberger et al., 2003; Way et al., 2009; Eisenberger, 2012). This opioids could also impact hedonic experience: they modulated the feeling of losing and winning in a gambling task (Petrovic et al., 2008) and placebo was preferred to naloxone (an opioid blocker) in rats (Skoubis et al., 2005). These effects could be mediated by dopamine: opioids are generally inhibitory, but this inhibition produces a des-inhibition (inhibition of inhibition) in the ventral tegmental area, a dopaminergic nucleus projecting on the limbic system (Corbett et al., 2006). Overall, despite the rational would be to use opioids to modulate the sensory feedback, other opioid-related effects could impact the effort allocation process.

Another perspective would be to study the placebo effect. In principle, the placebo effect could be estimated in the current data by subtraction of the placebo visit to a control group for which nothing would be mentioned about pain and taking pain killer. The fMRI and MEG group could have served as a control group, however taking a placebo pill and being told about pain is not the only difference between the groups: payoffs differed (by design, 4 times higher in the imaging study) and

participants were physically more constrained in the imaging study. The placebo effect is interesting in the present context because it relies on the subject's expectations (Petrovic et al., 2010) and expectations could be at play in particular during rest in the implicit effort allocation task. The role of expectation and their relation to awareness could be dissociated in the placebo effect according to the recent finding that placebo (and nocebo) effect can be induced with subliminal cues (Jensen et al., 2012). The role of expectations is partially addressed in this dissertation with the work on hypnotic suggestion.

Finally, it was interesting in the paracetamol study, although unrelated to the aim, that behavior changed between the two visits: participants were more strategic during the second visit. In particular rest durations were more modulated by the difficulty level during the second visit. It is possible that participants better exploited or had a better access to this information during the second visit. Overall, it suggests that learning may also be at play in the effort allocation process.

### 3.3.1.4 Methods

#### 3.3.1.4.1 Participants

Participants were recruited within the campus of the Hopital Pitié-Salpêtrière among medical students. They gave informed consent prior to participation. 19 participants took part in the study, however, one was excluded because he came only to the first visit (the withdrawal was not related to any adverse effect, but to schedule issues), two other participants were not included in the analysis because of a mistake in the stimulation program (there is a scaling factor to translate the performance into the online and actual monetary payoff; this scaling factor was changed between the two visits). 16 were eventually included, 9 males, all were between 20 and 26 years-old, mean: 22.2, s.e.m. 0.41, their average weight was 65.2 kg, +/-2.4 s.e.m. The data were acquired by Lou Safra (mostly) and Florent Meyniel.

A preliminary interview ensured that participants received no chronic treatments (except anti-progestative and seasonal allergy treatments). In particular, participants were screened to avoid chronic disease or history of psychiatric, neurological, hepatic troubles or nephropathy. We made sure that no antalgic were administrated within 48 hours prior to the experiment, that participants had ever received a paracetamol treatment to avoid cases of allergy and that they had no known allergy to compounds of the pill.

#### 3.3.1.4.2 Task and Experimental procedure

The task used was the implicit effort allocation task. Participants run the experiment twice with a 48h interval of wash-out and at the same moment of the day. Participants took *per os* with a glass of

water, either paracetamol, two capsules of 500 mg each (Doliprane, Sanofi), or a Plabeco (the content of the capsules was replaced by wheat flour), using a double-blind, cross over design. The treatment type and visit number were counterbalanced over subjects. After a delay of 30 minutes during which subjects waited in a quiet room, the instructions of the task were delivered, the maximal force was estimated, and the task was performed.

#### 3.3.1.4.3 Payoff

The amount of money accumulated during the task was rounded up (6.9€ +/- 0.5 s.e.m.) and was added to a fix amount (20€) at each visit. The total was paid cash after the experiment of the second visit.

#### 3.3.1.4.4 Statistical analysis

The preprocessing to define effort and rest epochs, and the estimation of the effect size and their significance used the same procedure as in page 116. Comparison of variables (mean value, effect sizes) between treatments, or between visits used paired, bilateral t-test (treatment or visit number within-subject contrast).

### 3.3.2 Manipulation of expectations with motivational and analgic hypnotic suggestion

#### 3.3.2.1 Introduction

Strategies to kill pain involve both pharmacological and psychological manipulations. In the previous section, a pharmacological agent was used (paracetamol). Pharmacological studies are interesting because they enable to address biochemical mechanisms associated to cognitive processes. Psychological manipulation could have a complementary interest because the description is closer to the cognitive level. As in the paracetamol study, the aim is to test whether proprioception and nociception are involved in the effort allocation process. Hypnosis has proved to be a potent psychological pain killer, based on clinical evaluations (Wark, 2008), in particular when the use of pharmacological agent is difficult, like in labor and delivery (Landolt and Milling, 2011).

We run again the implicit effort allocation task, but with a prior manipulation of nociceptive and proprioceptive signals based on hypnotic suggestion. In general, the hypnotic suggestion works as follows: the participant is driven into a hypnotic state, then the hypnotist makes a suggestion like 'you shall stop smoking', 'you won't feel pain', and finally the participant is awakened. The suggestion is not a short injunction, but a little talk in which the hypnotist suggests strategies to achieve a specific goal. For instance to relief pain specifically in one arm, the hypnotist may suggest that the feeling is smoothened, as if this arm was conformably positioned in a harm and cottony place and the negative sensation ignored. Given the flexibility allowed by the procedure, we tested

another suggestion to boost the participant motivation to win as much money as possible by exerting a lot of effort. This latter suggestion is not used by clinicians, but was worth trying. The predictions were that the pro-antalgic suggestion would reduce the effect of cost (longer effort, less effect of the difficulty on the effort duration, and potentially a faster recovery at rest) whereas the pro-motivational suggestion would increase the effect of benefit (longer effort, shorter rest, enhanced effect of the incentive on effort and rest durations).

### **3.3.2.2 Results**

The study on the Hypnosis groups replicated the basic results of the fMRI & MEG study. Pooling data irrespective of the manipulation, subjects increased effort duration for higher incentives ( $p=0.03$ ), reduced effort duration when more difficult ( $p=0.001$ ), reduced rest duration for higher incentives ( $p=0.009$ ) and did not change rest duration whatever the difficulty ( $p=0.6$ ).

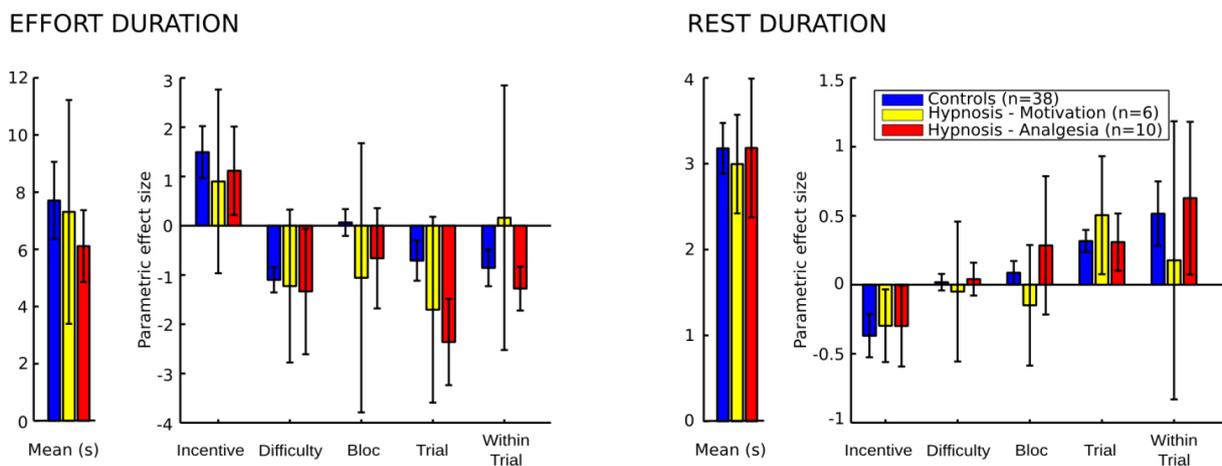
None of the suggestion had a significant impact on the payoff earned by the suggested hand compared to the control hand. In the antalgic group, payoff for the suggested hand was 8.9 CHF  $\pm$  2.6 (5% CI) and 9.4 CHF  $\pm$  3.1 (5% CI) for the control hand (paired t-test:  $p=0.74$ ). In the motivation group, payoff for the suggested hand was 8.4 CHF  $\pm$  4.8 (5% CI) and 7.8 CHF  $\pm$  3.1 (5% CI) for the control hand (paired t-test  $p=0.57$ ).

Mean effort and rest durations were also compared between groups. The direction of the change in the motivation group was as expected, but far from significance. Mean rest durations were shorter for the suggested vs. control hand (3.0 s  $\pm$  0.57 vs. 3.4  $\pm$  1.2, with 5% CI, paired t-test  $p=0.3$ ) and mean effort durations were higher for the suggested vs. control hand (7.3 s  $\pm$  3.9 vs. 6.3  $\pm$  2.9, with 5% CI, paired t-test  $p=0.2$ ). The changes were minimal and not significant in the analgesia group, for mean rest duration of suggested vs. control hand (2.9 s  $\pm$  1.0 vs. 2.9  $\pm$  1.1, with 5% CI, paired t-test  $p=0.9$ ) and for mean effort duration of suggested vs. control hand (6.6 s  $\pm$  2.5 vs. 7.0  $\pm$  2.5, with 5% CI, paired t-test  $p=0.7$ ). In the two groups, the paired comparison suggested vs. control hand of the incentive or difficulty effects never reach significance, for either effort or rest durations (all  $p>0.2$ ).

The absence of effect in the within-subject comparison approach could be due to a lack of effect of the suggestion. However, true effects of the suggestion could also be masked by the variability of the calibration between hands. Another possibility is that the effect of the suggestion is generalized to both hands instead of being hand-specific, which would kill the effect in the between-hand contrast.

To give a 'second chance' to the data, the results were re-analysed with a between subject approach, in which the analgesic and motivation groups are contrasted. For completeness, a control group was included in the comparison using the data from the MEG & fMRI group.

Data for the suggested hand in each group was compared between the hypnotic groups and between these groups and the control group. Figure 39 shows the results for mean effort and rest duration and the effect sizes of the modulation by the incentive, the difficulty, the bloc number, the trial number within a bloc and the effort or rest epoch position within a trial. The main comparisons of interest are for the mean and the effects of incentive and difficulty, which were all non-significant (all  $p > 0.24$ ). Some difference appear in the 'fatigue' effects (bloc, trial or epoch position) when compared to the control group, but the direct comparison between the analgesic and motivation group were never significant (all  $p > 0.1$ ).



**Figure 39 Suggested hand compared to a control group: effort, rest durations and modulations.**

Each color represents the data of a given group: both hands of a control group (blue), the hand targeted by a motivational (yellow) or analgesic (red) hypnotic suggestion. The bars represent the effect size of regressors estimated from a linear model that includes the mean and parametric modulations around the mean by the incentive level, the difficulty level, the position of rest or effort epochs between blocs (1 to 4,) trials of a bloc (1 to 9) and within trials. Error-bars are 5% confidence interval across subjects. Means and effect sizes are plotted separately to provide a more convenient scaling.

A similar between subject analysis that differed in the fact that data from both hand, and not only the hand targeted by the suggestion did not revealed any other significant differences for any comparison. The only interesting result of this analysis is that a consistent trend, despite always non-significant, emerged across comparisons: the suggested group have an impaired capacity to exert effort compared to the control group (longer rest and short effort, reduced effect of the incentive and higher effect of the difficulty, and overall higher effect of fatigue on both effort and rest durations).

### 3.3.2.3 Discussion

To summarize, the hypnotic suggestion had little effect: either when comparing the suggested hand to the control hand, or when comparing the suggested groups to a control group and when comparing the two suggested groups. Building on a negative finding should be avoided; however two comments can be made. First, the pro-motivational suggestion is not something our hypnotist had ever done: it is almost never used compared to the analgic suggestion. It is therefore possible that we were not able to induce any pro-motivational effect at all. Second, the practice of clinician is aimed at reducing kinds of pain that have little in common with what pain could be in our task. Indeed, exertion-induced pain is likely to be different from post-operative, cancer-induced, labor-induced, etc. pains. In addition, contrary to patients, in our task people can decide to stop pain almost instantaneously by stopping the exertion. It is therefore possible that nociception is involved in our task but that the clinical experience in killing pain is of little use in this context.

This difference in kinds of pain should be contrasted to what motivated this study. We reasoned that, in principle, there could be common mechanisms at play in the brain to deal with nociception in hypnotic analgesia and to deal with cost in the effort allocation task. The proposed mechanism for hypnotic analgesia is an attentional filtering. Indeed, highly hypnotizable subjects also have better attentional capabilities (Crawford, 1994). Interestingly, both the nociceptive sensation and the unpleasantness of pain are affected by hypnotic analgesia (Rainville et al., 1999; Faymonville et al., 2000), which suggests that hypnosis impacts both the sensory and affective aspect of pain. Related to the attentional hypothesis is the suggestion that the cognitive conflict elicited by pain would be alleviated under hypnosis, while the cognitive control is preserved. This could involve the anterior cingulate cortex, as evidenced by neuroimaging studies (Gruzelier, 1998; Egner et al., 2005; Raz et al., 2005). However, the effect of hypnosis might be more complex than just deciding to ignore pain. In another modality (induced paralysis), mechanisms associated to intended paralysis (acting as if paralyzed) and to hypnosis-induced paralysis were underpinned by different brain networks (Cojan et al., 2009).

As in the paracetamol study, the analgic suggestion did not yield the effects predicted after the hypothesis that cost-evidence depends on nociception. Again, it is possible that the hypothesis is wrong, although it is not safe to conclude from a negative finding. It is particularly not safe that our study lacks critical independent control tasks with the same subjects: first, that the hypnotic suggestion significantly modulated nociception or proprioception; second, that the suggestion does not impair motor actions. Indeed it would be a pity that a significant effect of the suggestion on nociception is counterbalanced by a motor impairment to result in a null effect on effort allocation.

### 3.3.2.4 Methods

This study was performed in the University of Geneva, LABNIC lab with the collaboration of Yann Cojan in the team of Patrik Vuilleumier.

#### 3.3.2.4.1 Participants

Participants were recruited within a data-base of the LABNIC and participated in the study in this lab. They gave informed consent prior to participation, following the local procedure. All were familiar with hypnosis; they had participated previously in sessions of hypnotic induction and were all highly susceptible to hypnosis. 16 subjects participated in the study, 13 with both Florent Meyniel & Yann Cojan using the experimental set-up used in the Implicit Task, and 3 other with Yann Cojan alone using a new experimental set-up developed at the LABIC lab to duplicate the original setup. This new set-up comprised a similar pneumatic handgrip that was connected to a BIOPAC system to convert pressure into voltage and condition the signal for the stimulation computer. Not all demographic data were collected. Sex is available for 13 participants, 8 were males; age is available for 9 participants, mean age 25.9 and s.e.m. 8.

#### 3.3.2.4.2 Task and Experimental procedure

The task was exactly the same of the Implicit Task (see page 116), but it was restricted to 6 (not 8) sessions for the sake of time. First, the task instructions were delivered and the maximal force was measured, then the hypnotic suggestion lasting 15 minutes was delivered and last the task was performed. Two types of suggestion was delivered: motivation (6 participants) or analgesia (10 participants), and targeted the right or left hand. The targeted hand was counterbalanced over participants within types of hypnotic suggestion. The experimenter was blind to the suggestion.

The hypnotic suggestion was delivered as follow. The participants were sitting comfortably, and the hypnotist induced the hypnotic state with an eye-roll procedure, which was familiar to participants. The hypnotist then suggested motivation or analgesia. The motivation suggestion emphasized the prospect of a gain and the eager to obtain it, with aspect of competition since most participants were either highly trained or at least much involved in sport. The hypnotist used during the suggestion the recall of a past experience of the participant, usually involving winning a sport contest. This experience was identified prior to the experiment through a specific interview between the participant and the hypnotist. In the analgesic suggestion, the hypnotist benefited from her practice as an experienced anesthetist at the local hospital and her routine medical use of hypnosis. The suggestion emphasized on gating the proprioceptive signal from the hand and fancying the hand as protected from pain. In both conditions, the suggestion targetted only one hand, the scope of the

suggestion was limited to the experiment. Finally, the participant was gradually returned to normal, awake condition.

#### 3.3.2.4.3 Payoff

Participants were paid in cash at the end of the experimental session. The amount won during the task was added to a fix payoff (10 CHF).

#### 3.3.2.4.4 Statistical analysis

The preprocessing to define effort and rest epochs, and the estimation of the effect size and their significance is the same as in page 116. Comparison of variables (mean value, effect sizes) within a group used paired, bilateral t-test (suggested and control hand within subject contrast) and between groups used un-matched sample bilateral t-tests.

### 3.3.3 Manipulation of central serotonin levels with Escitalopram

#### 3.3.3.1 Introduction

Serotonin may be of particular interest in the effort allocation task since it is associated to the processing of cost (Dayan, 2012), and as far as physical effort is concerned, to pain and fatigue.

The serotonin hypothesis of fatigue on physical performance posits that high levels of serotonin in the brain are responsible for supra-spinal fatigue during physical effort (Nybo and Secher, 2004). It is suggested that this mechanism is part of a humoral feedback loop: during intense exercise, muscular work induces an imbalance in the blood amino-acid concentration, in particular with increased tryptophan concentration. Tryptophan, a precursor of serotonin (5 hydroxytryptamine: 5HT), crosses the brain blood barrier and increases the serotonin brain levels (Gandevia, 2001). The serotonin hypothesis of fatigue is well supported by animal studies, but evidence is still lacking in humans: manipulation of serotonin levels was related to either increased, decreased fatigue or no effect on physical performance (Nybo and Secher, 2004). Inconsistency in the data could be accounted for by the diversity of 5HT receptors and their distinct post-synaptic effects. For instance, 5HT<sub>2A</sub> receptors increase the excitability of motoneurons (Jacobs and Azmitia, 1992), but 5HT<sub>1A</sub> receptors inhibit motoneurons, which could be a mechanism for physical fatigue at central levels (Cotel et al., 2013). Finally, the role of increased serotonin levels on fatigue, and sensation of fatigue, may have a direct effect, but indirect effects are also possible in particular through motivation and the cross regulation between neuromodulators, notably with dopamine (Boyas and Guével, 2011).

Serotonin is also involved in the regulation of pain. The small diameter fibers that convey nociception to the brain are modulated by serotonin, among many other things (Craig, 2002). 5HT<sub>1</sub> receptors participate in the postsynaptic inhibition of neurons in the dorsal horn of the spinal cord, which

conveys nociception (Jacobs and Azmitia, 1992). This serotonergic pain regulation mediates, through 5HT1A receptors, the analgic role of several treatments, such as paracetamol, although other systems are also involved (Smith, 2009). There is also evidence of serotonin-dependent pain regulation in the serotonin-related polymorphisms and the inter-individual differences in pain response, e.g. with variants of a gene controlling the expression level of the serotonin transporter: 5HTTLPR (LaCroix-Fralish and Mogil, 2009). People with the allele of 5HTTLPR that induces low expression of the transporter, hence higher concentrations of serotonin in the synaptic cleft, tolerate pain better: they have higher pain thresholds for heat pain and pressure pain (Lindstedt et al., 2011). This role of serotonin in pain regulation however finds only incomplete evidence in the clinical effect of antidepressants: tricyclic antidepressants and dual serotonin – noradrenaline reuptake inhibitors have a potent analgic effect, but selective serotonin reuptake inhibitors (SSRIs) only have a limited and inconsistent analgic effect across studies (Dharmshaktu et al., 2012).

We took advantage for the current study of a clinical trial by a pharmaceutical firm to assess the effect of Escitalopram on healthy males and females participants in the implicit effort allocation task. This study was designed as a randomized, double-blind placebo controlled study, where participants took either Escitalopram (therapeutic dose) or a placebo. Escitalopram is often qualified as the quintessential SSRI because of its selective effect on serotonin reuptake transporter (Stahl, 2008). Participants were submitted to the treatment (Escitalopram or placebo) during 8 weeks, which corresponds to the average duration of the acute SSRI treatment phase of depressive disorder. Within this duration, two thirds of the patients respond to the SSRI treatments. The effort allocation task was performed three times (at day 3, 14 and 56) during the treatment, which gave the opportunity to address the pharmaco-kinetic effect of the treatment. This is of particular interest because the antidepressant effect is delayed, and early phase of the treatment can be marked by side effects. The prediction of the effect of Escitalopram treatment, compared to placebo in the effort allocation behavior is actually difficult due to the complex pharmaco-kinetics of SSRI treatment and due to the sometimes opposed roles of 5HT receptors, for instance on fatigue and pain related processes.

### **3.3.3.2 Results**

Data were first analyzed irrespective of the treatment and the visit (for each subject, the average across visit entered the analysis) to check whether the results replicate the original findings of the fMRI & MEG group. Participants increased effort durations for higher incentives ( $p=6.4 \cdot 10^{-6}$ ) and decreased it when more difficult ( $p=8.9 \cdot 10^{-8}$ ). Rest durations were decreased by higher incentives ( $p=3.0 \cdot 10^{-4}$ ) and increased when more difficult ( $p=0.001$ ). Note that this later effect is not what was

found in the fMRI & MEG group, in the Paracetamol study and in the Hypnosis study, in which a non-significant effect was found. This could be a false positive, the likelihood of this explanation is however tempered by the good statistical power of this group (28 subjects for which effects are estimated from 3 replicates of the experiment). The effect was not due significantly to one group, or late compared to early visits since there was no Treatment effect ( $F_{1, 83}=1.4$ ,  $p=0.25$ ), no Visit effect ( $F_{2, 83}=0.85$ ,  $p=0.43$ ) and no interaction between the two ( $F_{2, 83}=0.84$ ,  $p=0.44$ ).

The maximal force produced by participants did not differ between groups (Placebo: 465 a.u. +/- 68 s.e.m.; Escitalopram: 368 a.u. +/- 47 s.e.m.;  $F_{1, 83}=2.2$ ,  $p=0.14$ ) and there was no effect of the visit ( $F_{2, 83}=8.8$ ,  $p=0.42$ ), but the interaction between the two was significant ( $F_{2, 83}=5.2$ ,  $p=0.008$ ), see Figure 40. The interaction was due to the fact that the difference was in the same direction over visits, but more significant during the first ( $p=0.03$ ) and the third visit ( $p=0.13$ ), compared to the second ( $p=0.98$ ).

The payoff was higher with Escitalopram than Placebo, however, not significantly (Escitalopram: £36.6 +/- 3.0 5% CI; Placebo: £31.8 +/- 3.3 5% CI,  $F_{1, 83}=1.92$ ,  $p=0.18$ ). There was no significant effect of the visit ( $F_{2, 83}=1.73$ ,  $p=0.19$ ), but a significant interaction between the visit and the treatment ( $F_{2, 83}=6.95$ ,  $p=0.002$ ), see Figure 40. The payoff was always higher in the Escitalopram group but did not showed any significant linear trend over visits ( $p=0.36$ ) whereas it increases over visit in the Placebo group ( $p=0.003$ ), such that the difference was higher during the first session ( $p=0.007$ ) and did not reach significance in the second and last visit ( $p>0.4$ ).

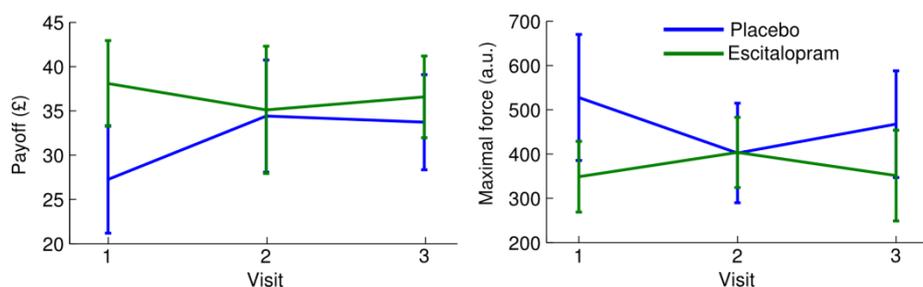


Figure 40 Escitalopram vs. Placebo: maximal force and payoff

Mean payoff and maximal force +/- 5% confidence interval across subjects.

The increased payoff in the Escitalopram group was driven by the following non-significant trends (see Figure 41): higher mean effort duration ( $F_{1, 83}=3.7$ ,  $p=0.067$ ), higher modulations of effort durations by the incentives ( $F_{1, 83}=1.96$ ,  $p=0.17$ ) and reduced mean rest durations ( $F_{1, 83}=1.47$ ,  $p=0.24$ ). For these three variables, there was no effect of the visit (all  $F_{2, 83}=0.57$ ;  $p=0.57$ ), and the interaction between the Visit and Treatment effect was marginal for effort mean duration ( $F_{2, 83}=2.5$ ,  $p=0.09$ ), and far from significance for incentive modulation of effort duration and mean rest duration (both  $F_{2, 83}$

$p > 0.93$ ,  $p = 0.4$ ). The visit-group interaction on the mean effort duration was due to a trend to the decrease the effect of the Escitalopram treatment over visits (linear trend:  $p = 0.058$ ) but not in the Placebo group (linear trend:  $p = 0.44$ ). The difficulty modulations of effort duration was not affected by the Treatment ( $F_{1, 83} = 0.22$ ,  $p = 0.6$ ) or the Visit ( $F_{2, 83} = 1.5$ ,  $p = 0.24$ ) or the interaction ( $F_{2, 83} = 0.57$ ,  $p = 0.57$ ). The difficulty modulation of rest duration was also not affected by the Treatment ( $F_{1, 83} = 1.4$ ,  $p = 0.25$ ), or the Visit ( $F_{2, 83} = 0.85$ ,  $p = 0.43$ ) or the interaction between the two ( $F_{2, 83} = 0.84$ ,  $p = 0.44$ ).

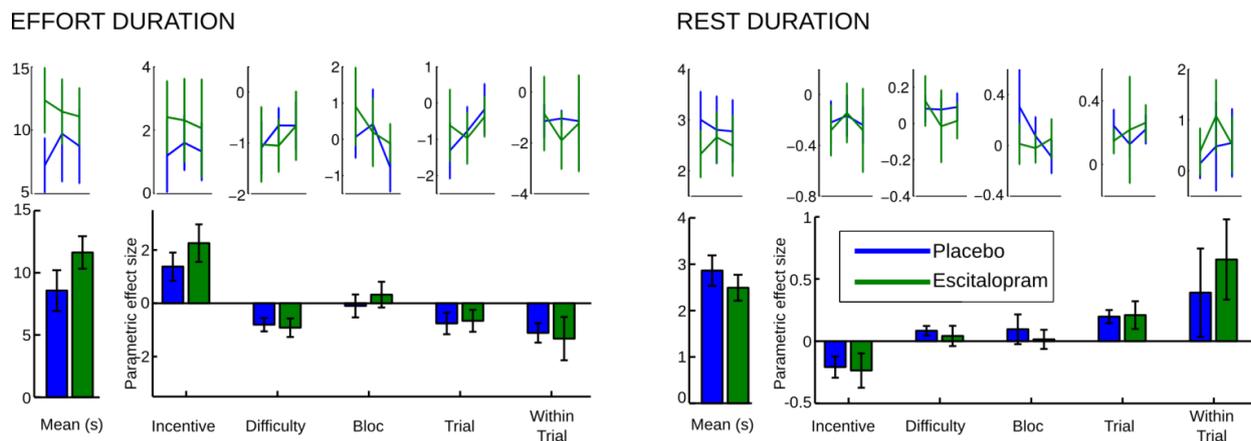


Figure 41: Escitalopram vs. Placebo groups compared across visits: effort, rest durations and their modulations.

Each color represents the data of a given treatment cohort. The bars in the bottom graphs represent the effect size of regressors estimated from a linear model that include the mean and parametric modulations around the mean by the incentive level, the difficulty level, the position of rest or effort epochs between blocs (1 to 4,) trials of a bloc (1 to 9) and within trials. The effect sizes were averaged subject-wise across visits. The upper graphs show the decomposition over the three visits, points from right to left correspond to Day 3, Day 14 and Day 56 of treatment. Error-bars are 5% confidence interval across subjects. Means and effect sizes are plotted separately to provide a more convenient scaling.

Regarding fatigue effects at different time scales: the session, the trials within a session and the effort or rest epoch position within a trial, there was not effect of the Treatment on effort fatigue timescales (all  $F_{1, 83} > 1.27$ ,  $p = 0.27$ ) nor interaction between the treatment and the visit (all  $F_{2, 83} > 2.15$ ,  $p = 0.13$ ), and there was no Treatment effect on the rest fatigue timescales (all  $F_{2, 83} > 1.87$ ,  $p = 0.18$ ) and a trend toward an interaction between the Treatment and Visit for the trial timescale ( $F_{2, 83} = 2.87$ ,  $p = 0.07$ ) but not for the session nor the rest position within a trial ( $F_{2, 83} > 1.5$ ,  $p = 0.23$ ).

### 3.3.3.3 Discussion

To summarize, participants under Escitalopram achieved a better payoff than the controls, through a lengthening of effort periods and a shortening of rest periods, and through a steeper effect of the incentive on effort duration. There was no clear pharmacokinetics effect, but a tendency to temper the Escitalopram effect over time. Many results are only trends toward significance, but overall, these results suggest that the Escitalopram treatment could have an energizing effect on the

behavior. This remains to be confirmed at the end of the study (only one third on the cohort was analyzed here).

The maximal force produced, that is used to calibrate the overall difficulty of the task, tended to be lower in the Escitalopram group. Could this explain the behavioral differences between the Escitalopram and the Placebo group? The maximal force produced differed little in the Escitalopram group between visits with, but if anything, a slight increase during the second visit. This pattern does not correspond to the variations over visits in the effects of interest. Furthermore, if the reported effect were due to the task being easier in the Escitalopram group than in the Placebo group, it would have been consistent to have smaller difficulty effects and fatigue effects in the Escitalopram group, which is not supported by the data: the trend is actually in the opposite direction.

It is intriguing that there is a significant effect of the difficulty level on rest duration in this study and not the others. Several differences between studies should be outlined. First, the hand grips used for this study are different from that used in the previous studies, with a width reduced by half. This alters the sensation of squeezing. This could change the proprioception in the task and make the detection of difficulty levels easier or more salient. We know that when they are explicit, cost levels affect rest durations. Second, different experimenters were involved in these tasks. Between-experimenter effect could be estimated on the full sample to test this explanation: 5 research assistants performed the tests. Last, participants in this study also participated in many other tests in this clinical trial, most of which did not involve incentive and performance-dependent payoff. These participants may have had different motivations, in particular an enhanced eagerness to optimize the payoff compared to participants in the other studies.

If the softening of the Escitalopram effects over visits were confirmed in the complete cohort, could it be related to the pharmaco-kinetics of antidepressant treatments? Under SSRI, there is an early increase of the neurotransmitter concentration and a delayed reduction of the receptor density. These changes lead progressively to normal levels of these neurobiological parameters that are abnormal in depressed patients. The mood improvement might be mediated by the delayed component and the adverse effects by the early component (Stahl, 2008). The early adverse effects and the delayed improvement is a big issue for the clinic because they limit the adherence to antidepressant treatments (Demyttenaere et al., 2001). This is all the more a problem that patients attribute these effects to the treatment and claim that it induces strange feelings that are not due to their depression. These adverse effects included an altered emotional processing: patients complain that the treatment make their emotions like artificial (Price et al., 2009). A potential explanation is that the treatment alleviates the negative bias that characterizes depression disorder before it

improves the mood, thus resulting in an emotional mismatch. An alternative explanation is that this effect is induced by the treatment irrespective of the disease. Indeed, the altered emotional processing under SSRI (Harmer et al., 2009) was also evidenced in healthy subjects with positive emotional bias under citalopram (Harmer et al., 2004).

There are not many studies using antidepressant in healthy subjects. Most of them addressed emotional processing (Harmer et al., 2004, 2009; Price et al., 2009). In fact, a meta-analysis on SSRI studies in healthy subjects revealed that consistent effects were not about anxiety and mood, but about emotional processing (Serretti et al., 2010). This positive emotional bias under SSRI is evidenced by improved memory recall for positive faces selectively and reduced startle responses to fearful faces (Harmer et al., 2004). By contrast, effects related to motivation and incentive processes are virtually unexplored. A noticeable exception is the study of SSRI in the probabilistic reversal learning task. In this task, people should distinguish errors due to the probabilistic predictive value of the cue from errors due to reversal in the contingencies (hence the name of the task). Healthy humans responded to acute citalopram treatment with higher sensitivity to losses (more lose-shift): they were more likely to change their strategies after negative feedbacks (Chamberlain et al., 2006). With a similar task, rats showed consistent results but a complex pharmo-kinetics: higher sensitivity to losses after acute citalopram and higher sensitivity to rewards after a chronic treatment (Bari et al., 2010). The increased behavioral energizing in the present results could be in line with the positive bias and the effect in the reversal learning task: there would be a general effect bias for positive feedback, biasing toward more money in the present task.

A more motor-oriented interpretation is also plausible. Indeed, it was stressed above that it was difficult to predict effect of SSRI in the effort allocation task was difficult because of the inconsistent post-synaptic effects of serotonin, e.g. between 5HT1A and 5HT2A. A recent PET study showed that remitters of depression treated by Escitalopram had higher binding at 5HT1A receptors in the raphe nucleus and all over the brain, compared to non-remitters, which suggests that 5HT1A receptors could play a key pharmacological role in Escitalopram treatments. It is tempting to speculate from previous reports that 5HT1A, by alleviating pain (Jacobs and Azmitia, 1992) could participate in the better performance of subjects in the Escitalopram group. However, this would be in contradiction with the pro-fatigue role of 5HT1A receptors that inhibit motoneurons in turtles (Cotel et al., 2013). Better motor performance was also reported under SSRI (paroxetine) in simple motor tasks such as tapping (Loubinoux et al., 2005). Although this behavioral results was comforted in other studies, the neural correlate of this effect was inconsistent across fMRI studies (Anderson et al., 2008).

### 3.3.3.4 Methods

#### 3.3.3.4.1 Participants

Healthy participants were recruited with public advertisement. Normal healthiness was insured by psychiatric, clinical and laboratory examinations, performed at selection visit. Participants with positive urine drug screening or medication in the course of the study (except paracetamol and oral contraceptives) were excluded. Pregnant women were not included. Most participants were students from the campus of the Universities of Oxford and Manchester. All participants gave informed consent prior participation in the protocol, in accordance with the British ethic comity that approved the study. The Placebo cohort had 15 participants, mean age 23.3 +/- 1.0 s.e.m., 6 males and 9 females; the Escitalopram cohort had 13 participants, mean age 24.7 +/- 1.2 s.e.m., 5 males and 8 females. The age difference is not significant ( $p=0.4$ ). Data were acquired by research assistants on the campus of Oxford and Manchester. Research assistants were trained by Florent Meyniel to deliver the test.

#### 3.3.3.4.2 Task and procedure

The study presented here is part of a wider phase I clinical trial, national (in the UK), multicentric, randomized, double-blind, with 4 parallel groups. The results presented correspond to the first third of the participants recruited in the study; it is a partial analysis in the course of the data acquisition.

The duration of the treatment was as follows:

- Selection period: 1 to 6 weeks (without treatment) between the selection visit and the inclusion visit. The treatment started at the inclusion visit, which served as a reference for the treatment duration and thus correspond to Day #0 and Week #0 (labeled D0 and W0).
- Double-blind treatment period: 8 weeks of treatment (Escitalopram or placebo) with visits at inclusion D0 (W0), D3 (W0), D7 (W1), D14 (W2), D35 (W5), D55 (W8) and D56 (W8).
- Double-blind tapering period (to reduce the Escitalopram dose): 1 week from D56 (W8) to D63 (W9) (Escitalopram or placebo).
- Follow-up period: 5-7 days after the end of the tapering period or study withdrawal (without treatment).

Participants were submitted to several tests and questionnaires. The implicit effort allocation task was performed at D3, D14 and D56.

The treatment dose was administrated as follow:

- Escitalopram – 10mg/20mg – oral route – 1 capsule in the evening at 8.00 p.m. (10mg the first week, 20 mg for the following 7 weeks and 10 mg for the tapering period).
- Placebo – oral route – 1 capsule in the evening at 8.00 p.m. for 8 weeks (and for the tapering period).

#### 3.3.3.4.3 Payoff

Participants were told they would receive a payment by bank transfer after the completion of the study and that a small part of this amount would be variable and depend on the performance in the task presented here. They were also told that for financial reasons, a fixed amount would be divided between the participants based on their relative performance, such that they could not predict their payoff but still would be motivated by the virtual money won in the game. Actually they all received the same amount in the end.

#### 3.3.3.4.4 Statistical analysis

The preprocessing to define effort and rest epochs, the estimation of the effect sizes and their significance followed the same procedure as in page 116. The effect size and the mean effort and rest duration then entered an ANOVA to estimate the effect of treatments and visits and the interaction. The ANOVA was performed with a mixed-effect in Matlab: subjects were treated as a random variable nested in the treatment variable, which was thus a between-subject effect; the visit was included as a within-subject effect. The subject, treatment and visit variables were categorical. Note that given that the visit was treated as a categorical variable, any effect (linear, quadratic, etc.) could be captured by the ANOVA. To follow-up the results of the ANOVA, the linear trends over visits were estimated in a separated analysis by regressing the variable of interest against the intercept and the linear effect of the visit number (1, 2 and 3; this values were z-scored). The significance of the linear effects of visit was estimated at the group level by testing the regression coefficients with bilateral t-test.

## 4 General discussion



## 4.1 Summary

We developed a paradigm to operationalize in the laboratory a temporal effort allocation problem at the time scale of several seconds. Although effort allocation is certainly at play in other paradigms, the tasks we developed are, to my knowledge, the first to investigate precisely this issue and to allow, given their short time scale, the manipulation of experimental factors within tasks, such as the incentive level and the effort difficulty.

The effort allocation behavior was not random but varied consistently with the factors manipulated, revealing that this process is solved under specific constraints. The first constraint evidenced is that effort allocation is adapted on the fly to the level of cost evidence, i.e. the cost associated to effort production. When cost evidence was increased by controlling the effort duration experimentally, the subsequent effort duration was impeded in proportion of this increase, suggesting that cost evidence is bounded. Conversely, when cost evidence was decreased by controlling the rest duration experimentally, the subsequent effort duration was facilitated in proportion of this reduction; this recovery was however bounded for high rest durations. The second constraint evidenced was that people seek to optimize utility, i.e. they adjust the bound of cost evidence and the recovery so as to maximize the benefit yielded through effort while minimizing physical cost. Higher incentives lengthened effort production and quicken effort resumption. Higher experienced difficulty levels speeded cost evidence accumulation, probably due to the physiology of effort production and higher expected difficulty levels strategically tempered the recovery rate. These two effects were dissociable.

The cost-evidence accumulation model we propose is a computational recipe to solve the effort allocation problem. This model builds on several advantages. First, this model is by essence dynamical, so that factors impacting this process can be translated into effects on the timing of decisions to stop and resume the effort. Second, the model is informed by the effort physiology: it assumes that effort costs increase during effort and are recovered during rest, and it combined this biological constraint to an optimization principle through the assumption that people maximize their utility, i.e. benefit discounted by cost. The resulting model accounts for both the behavioral adaptation on the fly to the cost evidence level, and the strategic modulations of this policy based on the effort cost and benefit. Therefore, this model makes the link between physiology, which imposes constraints that are not arbitrary but biologically relevant, and optimization principles that proved fruitful in many fields of research, in particular decision making. This model enables to relate the discrete observations that effort is stopped and resumed to the continuous process of effort allocation, so that this quantitative account can guide the search of a neural correlate of cost-

evidence. This hidden computational variable does not correspond to the perceived exhaustion. Indeed, the model predicts that the probability to stop the effort increases with the product of the effort duration and the effort difficulty. Behaviorally, the timing of effort cessation followed this interaction; however, the exhaustion rating calibrated on effort of controlled duration and difficulty reflected the sum, not the interaction, of these effort parameters.

The cost-evidence levels were correlated to the neural activity of proprioceptive regions (posterior insular cortex, secondary somatosensory cortex and ventro-medial thalamus) in a whole-brain functional magnetic resonance imaging (fMRI) search. A brain correlate, precisely in the proprioceptive regions, supports that this theoretical hidden variable may underpin the effort allocation process. Using a data-driven approach, we used magnetoencephalography (MEG) to confirm that the signal in these regions corresponded to time series of cost-evidence. This correlate corresponded not only to waning and waxing components of cost evidence but also to the modulations by the difficulty and benefit expected from the behavioral model fit. Neuroimaging data also uncovered properties of cost-evidence that could not be inferred from the behavior, e.g. that both bounds of cost evidence are impacted by the benefit, suggesting at the psychological level that higher benefits push the limit back during effort to allow more cost when more motivated, through modulations of the upper bound; and also improve the recovery during rest so that the system is better prepared to the following effort, through modulations of the lower bound. The neural correlate of cost-evidence is also impacted by factors not manipulated experimentally: observed variations of effort and rest durations within each experimental condition corresponded to variations in the steepness of the accumulation and dissipation of cost-evidence, respectively. This stresses the pervasiveness of the cost-evidence process and its ability to account for the behavior.

The investigation of the nervous system was not limited to finding correlates of cost-evidence, but also to relate the effort allocation process to the implementation of motor production. In particular, we showed that the faster resumption of effort for higher incentives could be mediated by the synchronization level in the beta band of motor activity (13 – 30 Hz). This synchronization level is known to be related to the facilitation of motor changes in the basal ganglia – cortical loop, and to be dependent on dopamine. Second, we showed that higher incentives are correlated to higher motivational arousal spanning the trial duration, as reflected in the heart beat rate activity. This could correspond to a tonic signal to energize behavior according to the benefit.

We tried to show that cost-evidence receives inputs from proprioception and nociception. To manipulate nociception, we used a pharmacological pain killer (paracetamol), and a psychological pain killer (hypnotic suggestion); however these two studies did not yield conclusive results. As the

prospect of a benefit affects effort allocation, we tried to manipulate this effect with a pro-motivational hypnotic suggestion, but this was not conclusive either. We also tested the role of serotonin in healthy subjects using a protracted treatment of a selective serotonin reuptake inhibitor: Escitalopram. Serotonin is involved in many processes in the brain, in particular the processing of cost, fatigue and pain. Escitalopram, compared to placebo, had an energizing effect: effort durations were increased, rest durations were reduced and the effect of incentive was enhanced during effort. This effect was higher when the treatment was started (day 3, compared to day 14 and 56). However these effects were marginally significant and are only preliminary data from on-going investigation on one third of the sample size.

## 4.2 Interpretations

The findings were discussed in the results sections after each piece of experimental work to improve the clarity and the logic of the dissertation. For the sake of brevity, I try not to repeat these above discussions but rather to complement them with transversal topics here. This is the opportunity to take a step back and rethink the relation between the accumulation model proposed here and the accumulation models often used in the literature on decision making. Second, I discuss what could be a more complete cost-evidence model with the relation to valuation processes, decision and motor implementation. Finally, I discuss the very meaning of cost-evidence.

### 4.2.1 The dynamic of effort allocation is supported by an accumulation model

The dynamics of cost is the key feature that enables the accumulation model to predict the timing of effort cessation and resumption. This dynamical property is used in other domain to account for the timing of perceptual decisions (Shadlen et al., 2006; Gold and Shadlen, 2007; Heekeren et al., 2008; Drugowitsch and Pouget, 2012) and value-based decisions (Basten et al., 2010; Krajbich and Rangel, 2011; De Martino et al., 2013) given the properties of stimuli. The meaning of the accumulation in these studies is however much different from that of the current work. Indeed, in these studies, the accumulation is a process to average out the noise of the input: the accumulation of evidence corresponds to the improvement of the estimation of the input. This accumulation is all the steeper that the signal to noise ratio is high. Setting a threshold a priori on the evidence accumulated to make a decision is an heuristic to implement (under specific conditions) the sequential probability ratio test (Cain and Shea-Brown, 2012; Churchland and Ditterich, 2012; Drugowitsch and Pouget, 2012). In other words, this accumulation is an integration of the input value, which improves the decision based on a noisy stimulus (Brunton et al., 2013). I suggest that it is not the case in the present model and that the primary meaning of the accumulation is not an integration of a noisy constant, but reflects the dynamics of a variable: the instantaneous physical cost. In this

interpretation, it is nevertheless possible that some form of integration is also involved: as any other signal, inputs to cost-evidence may be noisy and some form of integration may be used to filter these signals.

I now try to challenge my own proposal. In the present cost-evidence model, the instantaneous physiological cost is tracked over time, which results in dynamics with waxing and waning components over effort and rest. These dynamics are bounded, so that setting the upper bound according to the benefit during effort is advantageous relatively to utility maximization. A first alternative interpretation of the accumulation is that this dynamic reflects the integration of a constant net value, characterizing the effort utility. This approach was used in a different task. Stimuli were presented, each corresponding to a compound of potential monetary gain and loss; participants had to accept or reject these stimuli (Basten et al., 2010). The authors proposed that participants integrate the net value over time with a drift diffusion model to commit their choice. I doubt that the integration of a constant utility value, like in a standard drift diffusion model, would make sense in the effort allocation. There are at least three reasons. First, if constant net value was accumulated, the decision to stop the effort should be made always at the same decisional level, which was not the case since the upper bound was moved according to the benefit. A discrepant view would however acknowledge that decisional bounds of drift diffusion models can be adapted to the task parameters (Drugowitsch et al., 2012), so that this first argument may not be firm evidence to reject that cost-evidence is the integration of a net value. The second argument is that such integration would be severely sub-optimal to guide effort allocation. Indeed, when deciding whether to choose or reject an option, it makes sense to weigh the pros and the cons, possibly by integrating the net value. For instance we could choose to engage in an effort or not based on this mechanism. But what does it mean to make this comparison while we are already engaged in the effort? In this case, deciding to stop the effort would correspond to the decision that the effort is not worth doing, and hence that the exerted effort was actually a waste of energy. It is possible that it really takes several seconds to estimate the net value of effort. But in this case, people would stop resuming effort latter in a given trial or in the task when they get that effort is never worth doing, especially in the explicit version of the task. This would contradict the observed behavior. Third, we know from effort physiology that the instantaneous cost varies over time so that the effort utility is not constant.

Another possibility is that the brain uses a workaround to solve the utility maximization and to avoid excessive damages. For instance, the brain could integrate the difficulty level of effort that is manipulated experimentally and held constant in each trial: this integration over a constant value would increase over time and would increase more steeply for higher difficulty levels. To stop the

effort, this integration would be compared to the incentive level. Formally, this is exactly the dynamic of the cost-evidence accumulation supported by behavioral and neuroimaging data, so that computationally, this interpretation is equivalent to the one I propose. However, I do not see what kind of physiological tonic signal could encode the effort difficulty. The motor output is constant during the effort, but the difficulty level is not reflected in the visual display in the implicit task. The motor drive to produce this constant output is not constant: it increases over time to compensate muscular fatigue (Ma et al., 2009). The only physiological parameter that may remain constant is the muscle tension; this parameter is conveyed by Ia and Ib afferent fibers; however, this signal is modulated by other afferent fibers and by the efferent command. In addition, the neural correlate of cost evidence in the posterior insula strongly supports a connection to the physiological cost which increases over time.

Among these two alternative explanations, the first one makes little sense computationally and functionally; the second cannot be disentangled computationally from the interpretation I support, however, I may be less plausible functionally (i.e., from what we know from effort physiology and the neural correlate of cost evidence). Another reason to support my interpretation is that the meaning of cost-evidence is the same during effort and rest: it is the instantaneous effort-cost. If the increase in cost-evidence is interpreted as the integral over the difficulty level it is hard to explain why it should decrease during rest, unless one adds something to the model, like a leak in the integration. Instead, it seems more likely that cost-evidence accumulation truly reflects a change in the instantaneous effort cost and not the integration of a constant value over time.

Another key difference compared to other accumulation models is that cost-evidence is not only accumulated during effort up to a bound that triggers the effort cessation, it is also dissipated during rest down to a bound that triggers the effort resumption: the same variable subserves both decisions. It is easy to link the functional signal from the proprioceptive regions to the economic perspective. The signal from these regions increases during effort because the experienced cost increases, so that when bounded relatively to the incentive level, this mechanism to trigger effort cessation is economically sensible. It is less intuitive for rest to map the dynamics of the physiological signal onto the economic perspective: the opportunity cost of not doing the effort. We showed that the dissipation rate of cost-evidence is impacted by the incentive level and the expected difficulty level, as would be expected for the opportunity cost. It is unlikely that the proprioceptive regions encode the opportunity cost as such: the counterfactual nature of opportunity cost is at odds with the sensory nature of proprioception. However, it is possible that the physiological signal parallels the economic perspective, but following distinct mechanisms. First, it is not surprising that the

proprioceptive / nociceptive signal evolves during rest: muscles are still signaling their state, which is progressively less impacted by the preceding effort: there is a recovery. It may be more surprising that this signal is also impacted by the expected utility. What could be the mechanisms? We showed that there were lower levels of motor beta synchrony during rest. Beta (15 – 30 Hz) synchrony in the motor cortex is often anti-correlated to the gamma synchrony (> 40 Hz); both could have anti and pro-kinetic properties respectively, with opposite effects on cortico-spinal excitability (Schoffelen et al., 2005). Incentive-dependent changes at the spinal level could also modulate the ascending pathways, hence cost-evidence. It would be plausible since the descending and ascending pathways are tightly related. Another mechanism by which cost-evidence could be modulated during rest could involve the opioid system. Intrathecal injection of opioid agonist can relieve inhibition from nociceptive fibers and alleviate fatigue during physical exertion (Hilty et al., 2011b). It is possible that a placebo-like mechanism and opioids are at play to track cost evidence, especially given the anatomic correlate found for this variable (Meyniel et al., 2013). If opioids, or any other placebo-related mechanism were indeed involved, this could explain how cost-evidence is modulated by the utility during rest. Our data also showed an increased motivational arousal affecting the sympathetic system when higher incentives were at stake, this could also impact the body senses. Finally, it is also possible that the modulation of cost-evidence during rest by the utility is not mediated by the sensory aspect of cost-evidence itself, but through regulatory mechanisms in the brain. There are numerous reports across modalities that expectations modulate perception: motivation could bias this perception (Serences, 2008; den Ouden et al., 2010; Tallon-Baudry et al., 2011; Kok et al., 2012; de Lange et al., 2013). However, these are speculations and the issue remains quite open.

Last, the accumulation made it possible to have a continuous variable, waxing and waning over effort and rest period. Thus, we were able to track continuously cost-evidence and the related decisional process. This approach contrasts with the more widespread event-related analysis.

#### **4.2.2 What could be the complete model**

The cost-evidence model of effort allocation leaves several issues open and therefore appeal for additional experiments. In the following, I simply describe some of these issues.

The cost-evidence signal is impacted by the incentive and the effort difficulty, with the subtlety that the impact of the experience and the expected difficulty is different. Where do these input signals come from? Value signals were most convincingly observed in the fMRI signals when the incentive was revealed, in particular in the ventral striatum. We comparatively did not find correlate of the incentives later in the trial, either as a tonic signal or as a phasic signal during effort or rest or

behavioral transition. It is therefore not clear how the value information is transferred to the posterior insula / secondary somatosensory cortices. Actually, there is no need for this information to be transferred to these regions: it is possible that the incentive would affect directly the proprioceptive / nociceptive signal, so that it is reflected in these regions and that the bounds to this signal would not be set in the proprioceptive regions, but by a decisional system elsewhere in the brain reading these cost-evidence values and setting the threshold to make a decision. In both cases however, such other systems still need to be identified. Alternatively, the proprioceptive areas may combine a proprioceptive signal and a value signal to modulate cost-evidence and set bounds to send a stop and a go signal. In this case, where the value signal comes from, through which mechanisms it affects the proprioceptive regions and how the stop and go signals are implemented are open questions. Another open question is how closely related the incentive effect on the bounds and the slopes are and whether they correspond to common or distinct mechanisms to translate the incentive value into a behavioral response.

Similarly, it is not clear where the 'difficulty' signal comes from. The behavioral results suggest that the effect on the accumulation and dissipation slopes can be dissociated, which implies that at least two different routes can signal difficulty to the proprioceptive region. It is likely that the experienced cost is embedded in the ascending signal and that the expected cost is regulated by top-down control, for instance from the anterior cingulate cortex or the striatum that encoded the expected difficulty (Croxson et al., 2009; Prévost et al., 2010; Salamone and Correa, 2012; Kurniawan et al., 2013).

From the simplest perspective, this work showed that cost-evidence is a variable that may guide decisions to allocate the effort over time and that this variable is encoded in proprioceptive areas of the brain. First, we have no firm evidence for a causal role of this signal in effort allocation, it could be an epiphenomenon. Second, if it is not an epiphenomenon, it is not clear whether cost-evidence is the decisional variable, or a variable that guides decision. In both cases, how is the decision transformed into the actual motor response? The sensory cortex is tightly integrated functionally in movement production (Ghez and Krakauer, 2000b; Shadmehr and Krakauer, 2008). Sensory-motor transformation was studied in particular to understand its role in the control of motor parameters; it is also related to decisions to perform the action or not. The parietal-premotor cortical system could make the link between cost-evidence and the motor system, since it is involved for instance in the intention to move (Desmurget and Sirigu, 2009), and many paradigms involving accumulation of evidence assign a pivotal role to the parietal cortex between evidence tracking and motor output

(Gold and Shadlen, 2001; Donner et al., 2009; Louie and Glimcher, 2010; O'Connell et al., 2012; Wyart et al., 2012).

Another open issue is how cost-evidence is affected by learning. It was interesting in the paracetamol study that participants improved their performance in the second visit by optimizing their behavioral parameters: mean effort and rest duration and the sensitivity to difficulty and incentive. It would be interesting to know for instance whether people learn how to set the bound on cost-evidence and whether they adapt these bounds to their physical state, for instance as they become more fatigued along the task. It is particularly interesting for the lower bound of cost-evidence. In the current interpretation, this lower bound corresponds to how participants prepare for the following effort. When they resume effort too quickly, cost levels are still high and hence reduce the effort utility; on the contrary waiting longer improves the effort utility, but also reduces the time allotted to effort: how do people set the bound? Is it learned or updated?

Last, it is also unresolved how much automated or deliberative the effort allocation is. During debriefing, people reported some control on their decision to stop, sustain and resume the effort, but few of them reported explicit strategies; on the contrary, their report was rather vague. We interpreted the effort allocation process as comprising two levels: a level adapting the behavior on the fly to the cost evidence level, and a second level, superimposed on the first to implement strategic control on the mechanism parameters. To what extent do these two levels correspond to distinct processes? Could we find the dissociation, i.e. people with preserved on the fly adaptation, but impaired strategic modulations?

#### **4.2.3 What is cost-evidence?**

In this work, cost evidence is characterized functionally as the instantaneous effort-related cost, so that higher values make the effort less likely to be sustained and less likely to be initiated. This functional definition served to build a computational model, but left open the nature of this variable. Behavioral data showed that the increase in cost-evidence depended on the effort difficulty, which is consistent for a cost. On the contrary, reduction in cost-evidence during rest was not impacted by the true difficulty, but by the expected difficulty. This could make sense because during rest, the effort difficulty is rather a virtual parameter, so that the effort-related cost is not informed by the ongoing exertion and may rely more on expectations. The neural correlate of cost evidence was encoded in proprioceptive regions of the brain. Overall, these results would suggest that cost evidence is in part a sensory or, in a broader sense, a bodily signal, but it also reflects the predicted effort cost so that it depends on psychological factors and the expected effort utility.

This hybrid nature of cost-evidence maximizes the chance of not contradicting any literature. To what extent the sense of effort should be related to cost-evidence is not clear. In particular, our introspection study revealed a mismatch between the reported exhaustion level in a task where the effort was imposed and the exhaustion level inferred from the decision to stop when people could freely allocate their effort. There are limitations to this study, but it casts doubt on a simple relation between cost-evidence and the perceived exhaustion.

There is an interesting debate in the literature on perceived exhaustion about whether this sensation is related to the ascending signal: the sensory feedback, comprising proprioception, nociception, interoception; or the descending signal: the motor drive (Marcora, 2009; Smirmaul, 2012). If this debate was caricaturized into top-down vs. bottom-up origin of the sense of fatigue, it seems that besides the indefinite relation between cost-evidence and sense of fatigue, cost-evidence is neutral: it is affected by both top-down and bottom-up signals.

However, I must acknowledge that the origin of the cost signal still needs to be uncovered. Our attempt to demonstrate a relation to nociception was not conclusive. Proprioception and interoception are also plausible candidates. It is also possible that an efferent copy from the motor command contributes to the cost-signal. The modulations by strategic factors, in particular when they are not related to the effort produced as in the dissociation task, cast doubt that the cost-evidence signal is purely informed by sensations.

Last, was it a bright idea to give 'cost evidence' such a name? The 'cost' part of the name is rather intuitive and consistent given the results provided. It is less clear for 'evidence'. One drawback is that it may introduce confusion with the drift diffusion model literature and the estimation of a noisy input. However, evidence is a common word which is not used only in this specific literature. It has the advantage of drawing a connection with the sensory domain, to which cost-evidence is likely to be related. 'Evidence' indicates that this data is not a given, but an estimate: the brain constructs this value. And as in the perceptual domain, this estimate can be affected by priors and expectations through top-down regulation, which may also be the case for cost-evidence.

### **4.3 Limitations and unaddressed (related) topics**

I now make a short list of topics that were not addressed in the present work and that I am interested in. This is also the opportunity to stress the limits of the paradigms used. Note that not all limits are limitations: some limits usefully delineate the topic covered and characterize differences relatively to other topics. Other limits are more clearly limitations in the sense that improvements are needed.

I start with the limitations. The paradigms used implemented an effort allocation in which the intensity of the effort is not a free parameter but an imposed factor. In real life, not all actions have similar difficulty levels, but in many cases we can modulate the effort intensity. To take an example with physical effort: if you have to move your belongings to a new place, you may choose to lift two boxes at a time instead of one box at a time, this increases the effort intensity but also speeds up the relocation. Varying the effort intensity may be operant to reach the pursued goal (Pessiglione et al., 2007), and it may also be a way of allocating the effort over time (Tucker, 2009). A second limitation is that we ignored fatigue at a time scale longer than a trial, or at least it was treated as a co-variate in the analysis. The rationale was that we designed a paradigm to investigate the effort allocation within trials, so as to vary experimental factors between trials and quantify the effect of factors on effort allocation. However, fatigue at time scales longer than a given effort or rest epoch was significant and consistent across subjects. How do these effects impact the cost-evidence model and the behavior? A reason why we did not investigate these effects is the combinatorial complexity. The best cost-evidence model identified in the implicit effort allocation task has 5 free parameters, each of which can be affected by fatigue at several time scales (within trial, between trial, between sessions): the amount of models to test is thus 32768 (each parameter is potentially modulated by three time scales, resulting in 8 combinations, which makes  $8^5$  when crossed over 5 parameters). A third limitation is the indefinite understanding of what happens during the first rest of a trial, i.e. immediately after trial initiation. Participants did not resume effort immediately at trial onset, but prolonged rest; this rest was modulated by the incentive (and not by the difficulty) as the other rest epochs of the trial so that it was included in the analysis. Including this first rest in the fit of fMRI data also improved the results, suggesting that cost evidence is truly decreasing during this rest, as during the other rest epochs. This is consistent with the idea that there is a preparation to the effort. Nevertheless, to what extent this first rest is special or not, i.e. how the cost-evidence level is initialized would deserve further investigations, although the small amount of replicates impedes the statistical power of such analyses.

I now stress some limits of the paradigm. The first is actually more an open question than a claimed limit: we used handgrip effort, would it be much different with other types of effort? A related issue is how much the time scale of the effort allocation investigated matters: fatigue was readily observed within seconds, making effort allocation within trials necessary. Would it be fundamentally different for longer time scales? It is likely that the constraints on short and long time scales of effort allocation are different. In particular, humans have a circadian cycle, with periodic need for sleep that is a fundamental property of our brain (Krueger et al., 2008). It is also possible that in long run actions, there are different kinds of motivation along the time course of the effort. In particular, the

invested effort or time may alter our motivations, in particular by down-weighting future costs when much has already been invested, which is called the sunk cost effect (Arkes and Blumer, 1985). In the present study, we used money as an incentive, but it is not clear whether the benefit pursued by the subject was constant within the trial, even though the incentive value was kept the same. Last, it should be noted that we used physical effort, not mental effort. It was suggested that the encoding of benefit and cost of mental effort could use brain systems similar to physical effort (Boksem and Tops, 2008). How mental effort is allocated nonetheless remains an open issue. It is not clear whether mental effort could have similar dynamics of fatigue and recovery as the one used here over few seconds.

#### 4.4 Perspectives

I now conclude with perspectives on this work, more precisely things that I would do if I had more time. First, I am still convinced that nociception or proprioception may play a significant role in these effort allocation tasks. Nociception is easier to manipulate because many tools are available in the clinic. I would use drugs to target selectively the opioid system: morphine and naloxone, which are respectively agonist and antagonist of opioids receptors. The predicted effect of these treatments is an opposite pattern on the behavior, respectively with cost alleviation and cost aggravation. This project is about to be launched in the laboratory.

Second, I would like to manipulate the afferent signal directly to check whether it is an input to the cost-evidence signal. Transcutaneous nerve stimulation is used to relieve from pain, both on acute and chronic forms (Carroll et al., 2001; Rushton, 2002; Walsh et al., 2009). An electric periodic stimulation is applied on the afferent and disrupts the signal. The suggested mechanism is that the repeated stimulation induces an inhibition at the spinal level. The specificity of the effect and the direct manipulation of the peripheral signal are compelling, however a major drawback of the technic is that double blind manipulations are difficult to implement in practice because the stimulation is adapted to each participant to find the frequency that produces an effect.

Third, I am interested in whether the effort allocation could be implemented not only by switching effort on and off across time, but also by switching the effector used. In the effort allocation tasks we developed, this would correspond to allow the subject to freely use each hand alternatively. This is another way of solving the effort allocating problem. In real life situations, it is similar to using different means of pursuing the same goal. Does the cost-evidence model make specific predictions in this case? The same mechanism could serve at guiding decisions through the monitoring of effort cost compared to the expected benefit. What is not clear is how the costs from two different effectors may be combined, whether it is a simple summation or an interaction. The interactive

combination would suggest that both costs are not independent. This would happen for instance if there were common regulatory mechanisms, e.g. at the spinal level, or resource depletion at the body scale (e.g. oxygen, thermoregulation, although it may be limited with handgrip contraction, as opposed to more arduous effort such as cycling). This shared mechanism could also be in the brain: the brain would overweigh cost from a given effector if there is already high cost from another effector.

Fourth, the effort allocation tasks and the cost-evidence model could be used to characterize different kinds of apathy. Indeed, there could be a clinical interest of the task called 'dissociation effort allocation task' in this dissertation, because it behaviorally quantifies the weight of experienced cost and expected cost on, respectively, sustaining the effort and initiating the effort. Impairment of either process can result in less effort produced, which is a form of apathy. However, each process points to very different origins. One kind of apathy would correspond to avoiding effort because they are expected to be too demanding. The other kind of apathy would correspond to shortening effort because it is experienced as too difficult.

Finally, it seems that there is something special about expectation in the effort allocation task. The expected benefit impacts both effort and rest duration and the expected difficulty only impacts rest duration. The role of expectation is put forward in the placebo effect. It could be interesting to estimate the impact of the placebo effect in the effort allocation task. This would correspond to comparing two groups of subjects, the only difference would be that in one group, participants are given a (placebo) pill and told that it will reduce the experience of pain and fatigue ensued by effort. It would also be interesting to extend this expectation that costs are alleviated to the expectation that we are stronger. Instead of the antalgic property, we could tell the participant that the pill will make them stronger and full of energy.

## 5 References

- Abbiss CR, Laursen PB (2005) Models to explain fatigue during prolonged endurance cycling. *Sports Med Auckl Nz* 35:865–898.
- Amann M, Dempsey JA (2008) Locomotor muscle fatigue modifies central motor drive in healthy humans and imposes a limitation to exercise performance. *J Physiol* 586:161–173.
- Amemori K, Graybiel AM (2012) Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat Neurosci* 15:776–785.
- Anderson IM, McKie S, Elliott R, Williams SR, Deakin JFW (2008) Assessing human 5-HT function in vivo with pharmacMRI. *Neuropharmacology* 55:1029–1037.
- Arkes HR, Blumer C (1985) The psychology of sunk cost. *Organ Behav Hum Decis Process* 35:124–140.
- Baillet S, Mosher J, Leahy R (2001) Electromagnetic brain mapping. *Ieee Signal Process Mag* 18:14–30.
- Ball T, Schreiber A, Feige B, Wagner M, Lücking CH, Kristeva-Feige R (1999) The Role of Higher-Order Motor Areas in Voluntary Movement as Revealed by High-Resolution EEG and fMRI. *NeuroImage* 10:682–694.
- Bandettini PA, Kwong KK, Davis TL, Tootell RB, Wong EC, Fox PT, Belliveau JW, Weisskoff RM, Rosen BR (1997) Characterization of cerebral blood oxygenation and flow changes during prolonged brain activation. *Hum Brain Mapp* 5:93–109.
- Baraduc P, Thobois S, Gan J, Broussolle E, Desmurget M (2013) A common optimization principle for motor execution in healthy subjects and parkinsonian patients. *J Neurosci Off J Soc Neurosci* 33:665–677.
- Bari A, Theobald DE, Caprioli D, Mar AC, Aidoo-Micah A, Dalley JW, Robbins TW (2010) Serotonin Modulates Sensitivity to Reward and Negative Feedback in a Probabilistic Reversal Learning Task in Rats. *Neuropsychopharmacology* 35:1290–1301.
- Barros LF (2013) Metabolic signaling by lactate in the brain. *Trends Neurosci*.
- Barry BK, Enoka RM (2007) The neurobiology of muscle fatigue: 15 years later. *Integr Comp Biol* 47:465–473.
- Bartra O, McGuire JT, Kable JW (2013) The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage* 76:412–427.
- Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci U S A* 107:21767–21772.
- Bautista LM, Tinbergen J, Kacelnik A (2001) To walk or to fly? How birds choose among foraging modes. *Proc Natl Acad Sci U S A* 98:1089–1094.
- Becker GM, Degroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behav Sci* 9:226–232.

- Becker GS (1965) A Theory of the Allocation of Time. *Econ J* 75:493.
- Beckmann M, Johansen-Berg H, Rushworth MFS (2009) Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci Off J Soc Neurosci* 29:1175–1190.
- Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS (2007) Learning the value of information in an uncertain world. *Nat Neurosci* 10:1214–1221.
- Beierholm U, Guitart-Masip M, Economides M, Chowdhury R, Düzel E, Dolan R, Dayan P (2013) Dopamine Modulates Reward-Related Vigor. *Neuropsychopharmacol Off Publ Am Coll Neuropsychopharmacol*.
- Benedetti F (2008) Mechanisms of placebo and placebo-related effects across diseases and treatments. *Annu Rev Pharmacol Toxicol* 48:33–60.
- Berchicci M, Menotti F, Macaluso A, Di Russo F (2013) The neurophysiology of central and peripheral fatigue during sub-maximal lower limb isometric contractions. *Front Hum Neurosci* 7:135.
- Berridge KC (2004) Motivation concepts in behavioral neuroscience. *Physiol Behav* 81:179–209.
- Boksem MAS, Meijman TF, Lorist MM (2006) Mental fatigue, motivation and action monitoring. *Biol Psychol* 72:123–132.
- Boksem MAS, Tops M (2008) Mental fatigue: Costs and benefits. *Brain Res Rev* 59:125–139.
- Borg G (1990) Psychophysical scaling with applications in physical work and the perception of exertion. *Scand J Work Environ Health* 16 Suppl 1:55–58.
- Borg GA (1982) Psychophysical bases of perceived exertion. *Med Sci Sports Exerc* 14:377–381.
- Botvinick MM, Huffstetler S, McGuire JT (2009) Effort discounting in human nucleus accumbens. *Cogn Affect Behav Neurosci* 9:16–27.
- Bouret S, Ravel S, Richmond BJ (2012) Complementary neural correlates of motivation in dopaminergic and noradrenergic neurons of monkeys. *Front Behav Neurosci* 6:40.
- Bouret S, Richmond BJ (2010) Ventromedial and orbital prefrontal neurons differentially encode internally and externally driven motivational values in monkeys. *J Neurosci Off J Soc Neurosci* 30:8591–8601.
- Boyas S, Guével A (2011) Neuromuscular fatigue in healthy muscle: underlying factors and adaptation mechanisms. *Ann Phys Rehabil Med* 54:88–108.
- Brehm JW, Self EA (1989) The intensity of motivation. *Annu Rev Psychol* 40:109–131.
- Brittain J-S, Watkins KE, Joundi RA, Ray NJ, Holland P, Green AL, Aziz TZ, Jenkinson N (2012) A role for the subthalamic nucleus in response inhibition during conflict. *J Neurosci Off J Soc Neurosci* 32:13396–13401.
- Brodersen KH, Wiech K, Lomakina EI, Lin C-S, Buhmann JM, Bingel U, Ploner M, Stephan KE, Tracey I (2012) Decoding the perception of pain from fMRI using multivariate pattern analysis. *Neuroimage* 63:1162–1170.

- Bromberg-Martin ES, Hikosaka O, Nakamura K (2010) Coding of Task Reward Value in the Dorsal Raphe Nucleus. *J Neurosci* 30:6262–6272.
- Brooks AM, Berns GS (2013) Aversive stimuli and loss in the mesocorticolimbic dopamine system. *Trends Cogn Sci*.
- Brown P (2006) Bad oscillations in Parkinson's disease. *J Neural Transm Suppl*:27–30.
- Brown P, Oliviero A, Mazzone P, Insola A, Tonali P, Di Lazzaro V (2001) Dopamine dependency of oscillations between subthalamic nucleus and pallidum in Parkinson's disease. *J Neurosci Off J Soc Neurosci* 21:1033–1038.
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340:95–98.
- Bunzeck N, Guitart-Masip M, Dolan RJ, Düzel E (2011) Contextual novelty modulates the neural dynamics of reward anticipation. *J Neurosci Off J Soc Neurosci* 31:12816–12822.
- Burke CJ, Brünger C, Kahnt T, Park SQ, Tobler PN (2013) Neural integration of risk and effort costs by the frontal pole: only upon request. *J Neurosci Off J Soc Neurosci* 33:1706–1713a.
- Bush G, Luu P, Posner MI (2000) Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn Sci* 4:215–222.
- Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* 13:407–420.
- Cahuc P, Zylberberg A (2004) Labor supply. In: *Labor economics*. MIT Press.
- Cain N, Shea-Brown E (2012) Computational models of decision making: integration, stability, and noise. *Curr Opin Neurobiol* 22:1047–1053.
- Camerer CF, Hogarth RM (1999) The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework. *J Risk Uncertain* 19:7–42.
- Cardinal RN, Daw N, Robbins TW, Everitt BJ (2002) Local analysis of behaviour in the adjusting-delay task for assessing choice of delayed reinforcement. *Neural Networks Off J Int Neural Netw Soc* 15:617–634.
- Cardinal RN, Pennicott DR, Sugathapala CL, Robbins TW, Everitt BJ (2001) Impulsive Choice Induced in Rats by Lesions of the Nucleus Accumbens Core. *Science* 292:2499–2501.
- Carroll D, Moore RA, McQuay HJ, Fairman F, Tramèr M, Leijon G (2001) Transcutaneous electrical nerve stimulation (TENS) for chronic pain. *Cochrane Database Syst Rev Online*:CD003222.
- Carter RM, Meyer JR, Huettel SA (2010) Functional neuroimaging of intertemporal choice models: A review. *J Neurosci Psychol Econ* 3:27–45.
- Cassim F, Szurhaj W, Sediri H, Devos D, Bourriez J, Poirot I, Derambure P, Defebvre L, Guieu J (2000) Brief and sustained movements: differences in event-related (de)synchronization (ERD/ERS) patterns. *Clin Neurophysiol Off J Int Fed Clin Neurophysiol* 111:2032–2039.

- Chamberlain SR, Muller U, Blackwell AD, Clark L, Robbins TW, Sahakian BJ (2006) Neurochemical Modulation of Response Inhibition and Probabilistic Learning in Humans. *Science* 311:861–863.
- Charnov EL (1976) Optimal foraging, the marginal value theorem. *Theor Popul Biol* 9:129–136.
- Charron S, Koechlin E (2010) Divided Representation of Concurrent Goals in the Human Frontal Lobes. *Science* 328:360–363.
- Churchland AK, Ditterich J (2012) New advances in understanding decisions among multiple alternatives. *Curr Opin Neurobiol* 22:920–926.
- Cléry-Melin M-L, Schmidt L, Lafargue G, Baup N, Fossati P, Pessiglione M (2011) Why don't you try harder? An investigation of effort production in major depression. *Plos One* 6:e23178.
- Cohen D (1972) Magnetoencephalography: Detection of the Brain's Electrical Activity with a Superconducting Magnetometer. *Science* 175:664–666.
- Cojan Y, Waber L, Schwartz S, Rossier L, Forster A, Vuilleumier P (2009) The brain under self-control: modulation of inhibitory and monitoring cortical networks during hypnotic paralysis. *Neuron* 62:862–875.
- Colebatch JG (2007) Bereitschaftspotential and movement-related potentials: origin, significance, and application in disorders of human movement. *Mov Disord Off J Mov Disord Soc* 22:601–610.
- Corbett AD, Henderson G, McKnight AT, Paterson SJ (2006) 75 years of opioid research: the exciting but vain quest for the Holy Grail. *Br J Pharmacol* 147 Suppl 1:S153–162.
- Cotel F, Exley R, Cragg SJ, Perrier J-F (2013) Serotonin spillover onto the axon initial segment of motoneurons induces central fatigue by inhibiting action potential initiation. *Proc Natl Acad Sci U S A* 110:4774–4779.
- Cousins MS, Atherton A, Turner L, Salamone JD (1996) Nucleus accumbens dopamine depletions alter relative response allocation in a T-maze cost/benefit task. *Behav Brain Res* 74:189–197.
- Craig AD (2002) How do you feel? Interoception: the sense of the physiological condition of the body. *Nat Rev Neurosci* 3:655–666.
- Craig AD (Bud) (2003) Interoception: the sense of the physiological condition of the body. *Curr Opin Neurobiol* 13:500–505.
- Craig AD (Bud) (2009a) How do you feel — now? The anterior insula and human awareness. *Nat Rev Neurosci* 10:59–70.
- Craig ADB (2009b) How do you feel--now? The anterior insula and human awareness. *Nat Rev Neurosci* 10:59–70.
- Cramer SC, Weisskoff RM, Schaechter JD, Nelles G, Foley M, Finklestein SP, Rosen BR (2002) Motor cortex activation is related to force of squeezing. *Hum Brain Mapp* 16:197–205.
- Crawford HJ (1994) Brain dynamics and hypnosis: attentional and disattentional processes. *Int J Clin Exp Hypn* 42:204–232.

- Croxson PL, Walton ME, O'Reilly JX, Behrens TEJ, Rushworth MFS (2009) Effort-Based Cost-Benefit Valuation and the Human Brain. *J Neurosci* 29:4531–4541.
- Cunnington R, Windischberger C, Deecke L, Moser E (2003) The preparation and readiness for voluntary movement: a high-field event-related fMRI study of the Bereitschafts-BOLD response. *NeuroImage* 20:404–412.
- Cunnington R, Windischberger C, Moser E (2005) Premovement activity of the pre-supplementary motor area and the readiness for action: Studies of time-resolved event-related functional MRI. *Hum Mov Sci* 24:644–656.
- Czernecki V, Pillon B, Houeto JL, Pochon JB, Levy R, Dubois B (2002) Motivation, reward, and Parkinson's disease: influence of dopatherapy. *Neuropsychologia* 40:2257–2267.
- D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD Responses Reflecting Dopaminergic Signals in the Human Ventral Tegmental Area. *Science* 319:1264–1267.
- Dai T, Liu J, Sahgal V, Brown R, Yue G (2001) Relationship between muscle output and functional MRI-measured brain activation. *Exp Brain Res* 140:290–300.
- Dalsgaard MK (2005) Fuelling cerebral activity in exercising man. *J Cereb Blood Flow Metab* 26:731–750.
- Dalsgaard MK, Ogoh S, Dawson EA, Yoshiga CC, Quistorff B, Secher NH (2004) Cerebral carbohydrate cost of physical exertion in humans. *Am J Physiol Regul Integr Comp Physiol* 287:R534–540.
- Daunizeau J, Friston KJ, Kiebel SJ (2009) Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Phys Nonlinear Phenom* 238:2089–2118.
- Daw ND, Kakade S, Dayan P (2002) Opponent interactions between serotonin and dopamine. *Neural Networks Off J Int Neural Netw Soc* 15:603–616.
- Dayan P (2012) Twenty-five lessons from computational neuromodulation. *Neuron* 76:240–256.
- Dayan P, Abbott LF (2005) *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, 1st ed. The MIT Press.
- De Lange FP, Rahnev DA, Donner TH, Lau H (2013) Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *J Neurosci Off J Soc Neurosci* 33:1400–1410.
- De Martino B, Fleming SM, Garrett N, Dolan RJ (2013) Confidence in value-based choice. *Nat Neurosci* 16:105–110.
- Deecke L, Scheid P, Kornhuber HH (1969) Distribution of readiness potential, pre-motion positivity, and motor potential of the human cerebral cortex preceding voluntary finger movements. *Exp Brain Res Exp Hirnforsch Expérimentation Cérébrale* 7:158–168.
- DeLong M (2000) The basal ganglia. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 853–867. Mc Graw Hill.
- DeLuca J, Genova HM, Hillary FG, Wylie G (2008) Neural correlates of cognitive fatigue in multiple sclerosis using functional MRI. *J Neurol Sci* 270:28–39.

- Demyttenaere K, Mesters P, Boulanger B, Dewe W, Delsemme M-H, Gregoire J, Van Ganse E (2001) Adherence to treatment regimen in depressed patients treated with amitriptyline or fluoxetine. *J Affect Disord* 65:243–252.
- Den Ouden HEM, Daunizeau J, Roiser J, Friston KJ, Stephan KE (2010) Striatal prediction error modulates cortical coupling. *J Neurosci Off J Soc Neurosci* 30:3210–3219.
- Denk F, Walton ME, Jennings KA, Sharp T, Rushworth MFS, Bannerman DM (2005) Differential involvement of serotonin and dopamine systems in cost-benefit decisions about delay or effort. *Psychopharmacology (Berl)* 179:587–596.
- Desmurget M, Sirigu A (2009) A parietal-premotor network for movement intention and motor awareness. *Trends Cogn Sci* 13:411–419.
- Desmurget M, Sirigu A (2012) Conscious motor intention emerges in the inferior parietal lobule. *Curr Opin Neurobiol* 22:1004–1011.
- Dewall CN, Macdonald G, Webster GD, Masten CL, Baumeister RF, Powell C, Combs D, Schurtz DR, Stillman TF, Tice DM, Eisenberger NI (2010) Acetaminophen reduces social pain: behavioral and neural evidence. *Psychol Sci* 21:931–937.
- Dharmshaktu P, Tayal V, Kalra BS (2012) Efficacy of antidepressants as analgesics: a review. *J Clin Pharmacol* 52:6–17.
- Doñamayor N, Marco-Pallarés J, Heldmann M, Schoenfeld MA, Münte TF (2011) Temporal dynamics of reward processing revealed by magnetoencephalography. *Hum Brain Mapp* 32:2228–2240.
- Donner TH, Siegel M, Fries P, Engel AK (2009) Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Curr Biol* 19:1581–1585.
- Doya K (2008) Modulators of decision making. *Nat Neurosci* 11:410–416.
- Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A (2012) The cost of accumulating evidence in perceptual decision making. *J Neurosci Off J Soc Neurosci* 32:3612–3628.
- Drugowitsch J, Pouget A (2012) Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Curr Opin Neurobiol* 22:963–969.
- Duff E, Xiong J, Wang B, Cunnington R, Fox P, Egan G (2007) Complex spatio-temporal dynamics of fMRI BOLD: A study of motor learning. *Neuroimage* 34:156–168.
- Egner T, Jamieson G, Gruzelier J (2005) Hypnosis decouples cognitive control from conflict monitoring processes of the frontal lobe. *Neuroimage* 27:969–978.
- Ehrsson HH, Fagergren A, Jonsson T, Westling G, Johansson RS, Forssberg H (2000) Cortical Activity in Precision- Versus Power-Grip Tasks: An fMRI Study. *J Neurophysiol* 83:528–536.
- Eickhoff SB, Jbabdi S, Caspers S, Laird AR, Fox PT, Zilles K, Behrens TEJ (2010) Anatomical and Functional Connectivity of Cytoarchitectonic Areas within the Human Parietal Operculum. *J Neurosci* 30:6409–6421.

- Eisenberger NI (2012) The pain of social disconnection: examining the shared neural underpinnings of physical and social pain. *Nat Rev Neurosci* 13:421–434.
- Eisenberger NI, Lieberman MD, Williams KD (2003) Does Rejection Hurt? An fMRI Study of Social Exclusion. *Science* 302:290–292.
- Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? *Curr Opin Neurobiol* 20:156–165.
- Engel AK, Maye A, Kurthen M, König P (2013) Where's the action? The pragmatic turn in cognitive science. *Trends Cogn Sci* 17:202–209.
- Enoka RM, Stuart DG (1992) Neurobiology of muscle fatigue. *J Appl Physiol Bethesda Md* 1985 72:1631–1648.
- Erdler M, Beisteiner R, Mayer D, Kaindl T, Edward V, Windischberger C, Lindinger G, Deecke L (2000) Supplementary Motor Area Activation Preceding Voluntary Movement Is Detectable with a Whole-Scalp Magnetoencephalography System. *NeuroImage* 11:697–707.
- Euston DR, Gruber AJ, McNaughton BL (2012) The role of medial prefrontal cortex in memory and decision making. *Neuron* 76:1057–1070.
- Faymonville ME, Laureys S, Degueldre C, DelFiore G, Luxen A, Franck G, Lamy M, Maquet P (2000) Neural mechanisms of antinociceptive effects of hypnosis. *Anesthesiology* 92:1257–1267.
- Fehr E, Rangel A (2011) Neuroeconomic Foundations of Economic Choice Recent Advances. *J Econ Perspect* 25:3–30.
- Feige B, Kristeva-Feige R, Rossi S, Pizzella V, Rossini P-M (1996) Neuromagnetic study of movement-related changes in rhythmic brain activity. *Brain Res* 734:252–260.
- Fellows LK (2011) Orbitofrontal contributions to value-based decision making: evidence from humans with frontal lobe damage. *Ann N Y Acad Sci* 1239:51–58.
- Figner B, Murphy RO (2010) Using skin conductance in judgment and decision making research. In: A handbook of process tracing methods for decision research. New: Psychology Press.
- Floresco SB, Ghods-Sharifi S (2007) Amygdala-Prefrontal Cortical Circuitry Regulates Effort-Based Decision Making. *Cereb Cortex* 17:251–260.
- Fox MD, Snyder AZ, Barch DM, Gusnard DA, Raichle ME (2005) Transient BOLD responses at block transitions. *Neuroimage* 28:956–966.
- Fox P, Raichle M, Mintun M, Dence C (1988) Nonoxidative glucose consumption during focal physiologic neural activity. *Science* 241:462–464.
- Frank MJ, Doll BB, Oas-Terpstra J, Moreno F (2009) Prefrontal and striatal dopaminergic genes predict individual differences in exploration and exploitation. *Nat Neurosci* 12:1062–1068.
- Frank MJ, Moustafa AA, Haughey HM, Curran T, Hutchison KE (2007) Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proc Natl Acad Sci U S A* 104:16311–16316.

- Frank MJ, Seeberger LC, O'Reilly RC (2004) By Carrot or by Stick: Cognitive Reinforcement Learning in Parkinsonism. *Science* 306:1940–1943.
- Friebel U, Eickhoff SB, Lotze M (2011) Coordinate-based meta-analysis of experimentally induced and chronic persistent neuropathic pain. *Neuroimage* 58:1070–1080.
- Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007) Variational free energy and the Laplace approximation. *NeuroImage* 34:220–234.
- Frith CD, Blakemore, Wolpert DM (2000) Abnormalities in the awareness and control of action. *Philos Trans R Soc Lond B Biol Sci* 355:1771–1788.
- Gan JO, Walton ME, Phillips PEM (2010) Dissociable cost and benefit encoding of future rewards by mesolimbic dopamine. *Nat Neurosci* 13:25–27.
- Gandevia SC (2001) Spinal and Supraspinal Factors in Human Muscle Fatigue. *Physiol Rev* 81:1725–1789.
- Gardner EP, Martin JH, Jessel TM (2000) The bodily senses. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 713–736. Mc Graw Hill.
- Ghez C, Krakauer J (2000a) The organization of Movement. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 653–673. Mc Graw Hill.
- Ghez C, Krakauer J (2000b) Voluntary movement. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 653–673. Mc Graw Hill.
- Ghods-Sharifi S, Floresco SB (2010) Differential effects on effort discounting induced by inactivations of the nucleus accumbens core or shell. *Behav Neurosci* 124:179–191.
- Glascher J, Hampton AN, O'Doherty JP (2009) Determining a Role for Ventromedial Prefrontal Cortex in Encoding Action-Based Value Signals During Reward-Related Decision Making. *Cereb Cortex* 19:483–495.
- Glimcher PW, Camerer C, Fehr E, Poldrack RA (2009) Introduction: A brief History of Neuroeconomics. In: *Neuroeconomics Decision making and the brain*, Academic Press. (Glimcher PW, Camerer CF, Fehr E, Poldrack RA, eds). Elsevier.
- Glimcher PW, Rustichini A (2004) Neuroeconomics: The Consilience of Brain and Decision. *Science* 306:447–452.
- Gluth S, Rieskamp J, Büchel C (2013) Classic EEG motor potentials track the emergence of value-based decisions. *Neuroimage*.
- Goense JBM, Logothetis NK (2008) Neurophysiology of the BOLD fMRI signal in awake monkeys. *Curr Biol* 18:631–640.
- Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394.
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16.

- Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. *Annu Rev Neurosci* 30:535–574.
- Gross J, Kujala J, Hamalainen M, Timmermann L, Schnitzler A, Salmelin R (2001) Dynamic imaging of coherent sources: Studying neural interactions in the human brain. *Proc Natl Acad Sci U S A* 98:694–699.
- Gruzelier J (1998) A working model of the neurophysiology of hypnosis: a review of evidence. *Contemp Hypn* 15:3–21.
- Guigon E, Baraduc P, Desmurget M (2007) Coding of movement- and force-related information in primate primary motor cortex: a computational approach. *Eur J Neurosci* 26:250–260.
- Guitart-Masip M, Chowdhury R, Sharot T, Dayan P, Duzel E, Dolan RJ (2012) Action controls dopaminergic enhancement of reward representations. *Proc Natl Acad Sci U S A* 109:7511–7516.
- Gutschalk A, Hämäläinen MS, Melcher JR (2010a) BOLD Responses in Human Auditory Cortex Are More Closely Related to Transient MEG Responses Than to Sustained Ones. *J Neurophysiol* 103:2015–2026.
- Gutschalk A, Hämäläinen MS, Melcher JR (2010b) BOLD responses in human auditory cortex are more closely related to transient MEG responses than to sustained ones. *J Neurophysiol* 103:2015–2026.
- Hämäläinen M, Hari R, Ilmoniemi RJ, Knuutila J, Lounasmaa OV (1993) Magnetoencephalography: theory, instrumentation, and applications to noninvasive studies of the working human brain. *Rev Mod Phys* 65:413.
- Hammond C, Bergman H, Brown P (2007) Pathological synchronization in Parkinson's disease: networks, models and treatments. *Trends Neurosci* 30:357–364.
- Handler M, Nelson R, Krapohl D, Honts C (2010) An EDA Primer for Polygraph Examiners. Polygraph Available at: [http://scholarworks.boisestate.edu/psych\\_facpubs/48](http://scholarworks.boisestate.edu/psych_facpubs/48).
- Harenski CL, Thornton DM, Harenski KA, Decety J, Kiehl KA (2012) Increased frontotemporal activation during pain observation in sexual sadism: preliminary findings. *Arch Gen Psychiatry* 69:283–292.
- Harmer CJ, Goodwin GM, Cowen PJ (2009) Why do antidepressants take so long to work? A cognitive neuropsychological model of antidepressant drug action. *Br J Psychiatry* 195:102–108.
- Harmer CJ, Shelley NC, Cowen PJ, Goodwin GM (2004) Increased Positive Versus Negative Affective Perception and Memory in Healthy Volunteers Following Selective Serotonin and Norepinephrine Reuptake Inhibition. *Am J Psychiatry* 161:1256–1263.
- Harms MP, Guinan JJ Jr, Sigalovsky IS, Melcher JR (2005) Short-term sound temporal envelope characteristics determine multisecond time patterns of activity in human auditory cortex as shown by fMRI. *J Neurophysiol* 93:210–222.
- Hauber W, Sommer S (2009) Prefrontostriatal Circuitry Regulates Effort-Related Decision Making. *Cereb Cortex* 19:2240–2247.
- Hayden BY, Platt ML (2007) Animal Cognition: Great Apes Wait for Grapes. *Curr Biol* 17:R922–R923.

- Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci* 9:467–479.
- Herbert BM, Ulbrich P, Schandry R (2007) Interoceptive sensitivity and physical effort: Implications for the self-control of physical load in everyday life. *Psychophysiology* 44:194–202.
- Hillman KL, Bilkey DK (2010) Neurons in the Rat Anterior Cingulate Cortex Dynamically Encode Cost-Benefit in a Spatial Decision-Making Task. *J Neurosci* 30:7705–7713.
- Hilty L, Jäncke L, Luechinger R, Boutellier U, Lutz K (2011a) Limitation of physical performance in a muscle fatiguing handgrip exercise is mediated by thalamo-insular activity. *Hum Brain Mapp* 32:2151–2160.
- Hilty L, Lutz K, Maurer K, Rodenkirch T, Spengler CM, Boutellier U, Jäncke L, Amann M (2011b) Spinal opioid receptor-sensitive muscle afferents contribute to the fatigue-induced increase in intracortical inhibition in healthy humans. *Exp Physiol* 96:505–517.
- Hummel F, Kirsammer R, Gerloff C (2003) Ipsilateral cortical activation during finger sequences of increasing complexity: representation of movement difficulty or memory load? *Clin Neurophysiol Off J Int Fed Clin Neurophysiol* 114:605–613.
- Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MF, Behrens TE (2012) Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* 15:470–S3.
- Ide K, Horn A, Secher NH (1999) Cerebral metabolic response to submaximal exercise. *J Appl Physiol* 87:1604–1608.
- Iversen S, Iversen L, Saper CB (2000) The autonomic nervous system and the hypothalamus. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 853–867. Mc Graw Hill.
- Jacobs BL, Azmitia EC (1992) Structure and function of the brain serotonin system. *Physiol Rev* 72:165–229.
- James W (1907) The energies on men. *Science* 25:321–332.
- Jasper H, Penfield W (1949) Electrocorticograms in man: Effect of voluntary movement upon the electrical activity of the precentral gyrus. *Arch Für Psychiatr Nervenkrankh* 183:163–174.
- Jenkinson N, Brown P (2011) New insights into the relationship between dopamine, beta oscillations and motor function. *Trends Neurosci* 34:611–618.
- Jensen KB, Kaptchuk TJ, Kirsch I, Raicek J, Lindstrom KM, Berna C, Gollub RL, Ingvar M, Kong J (2012) Nonconscious activation of placebo and nocebo pain responses. *Proc Natl Acad Sci U S A* 109:15959–15964.
- Jones CL, Ward J, Critchley HD (2010) The neuropsychological impact of insular cortex lesions. *J Neurol Neurosurg Psychiatry* 81:611–618.
- Jones SR, Pritchett DL, Sikora MA, Stufflebeam SM, Hämäläinen M, Moore CI (2009) Quantitative analysis and biophysically realistic neural modeling of the MEG mu rhythm: rhythmogenesis and modulation of sensory-evoked responses. *J Neurophysiol* 102:3554–3572.

- Joundi RA, Jenkinson N, Brittain J-S, Aziz TZ, Brown P (2012) Driving oscillatory activity in the human cortex enhances motor performance. *Curr Biol* 22:403–407.
- Jurkiewicz MT, Gaetz WC, Bostan AC, Cheyne D (2006) Post-movement beta rebound is generated in motor cortex: Evidence from neuromagnetic recordings. *NeuroImage* 32:1281–1289.
- Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47:263.
- Kahnt T, Heinzle J, Park SQ, Haynes J-D (2010) The neural code of reward anticipation in human orbitofrontal cortex. *Proc Natl Acad Sci U S A* 107:6010–6015.
- Kahnt T, Heinzle J, Park SQ, Haynes J-D (2011) Decoding the formation of reward predictions across learning. *J Neurosci Off J Soc Neurosci* 31:14624–14630.
- Kalenscher T, Pennartz CMA (2008) Is a bird in the hand worth two in the future? The neuroeconomics of intertemporal decision-making. *Prog Neurobiol* 84:284–315.
- Kalivas PW, Nakamura M (1999) Neural systems for behavioral activation and reward. *Curr Opin Neurobiol* 9:223–227.
- Keisker B, Hepp-Reymond M-C, Blickenstorfer A, Kollias SS (2010) Differential representation of dynamic and static power grip force in the sensorimotor network. *Eur J Neurosci* 31:1483–1491.
- Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324:759–764.
- Klein-Flügge MC, Nobbs D, Pitcher JB, Bestmann S (2013) Variability of human corticospinal excitability tracks the state of action preparation. *J Neurosci Off J Soc Neurosci* 33:5564–5572.
- Klostermann F, Nikulin VV, Kühn AA, Marzinzik F, Wahl M, Pogosyan A, Kupsch A, Schneider G-H, Brown P, Curio G (2007) Task-related differential dynamics of EEG alpha- and beta-band synchronization in cortico-basal motor structures. *Eur J Neurosci* 25:1604–1615.
- Knutson B, Taylor J, Kaufman M, Peterson R, Glover G (2005) Distributed neural representation of expected value. *J Neurosci Off J Soc Neurosci* 25:4806–4812.
- Knutson B, Wimmer GE, Kuhnen CM, Winkielman P (2008) Nucleus accumbens activation mediates the influence of reward cues on financial risk taking. *Neuroreport* 19:509–513.
- Koechlin E, Hyafil A (2007) Anterior Prefrontal Function and the Limits of Human Decision-Making. *Science* 318:594–598.
- Koechlin E, Ody C, Kouneiher F (2003) The Architecture of Cognitive Control in the Human Prefrontal Cortex. *Science* 302:1181–1185.
- Kok P, Jehee JFM, de Lange FP (2012) Less is more: expectation sharpens representations in the primary visual cortex. *Neuron* 75:265–270.
- Kolling N, Behrens TEJ, Mars RB, Rushworth MFS (2012) Neural Mechanisms of Foraging. *Science* 336:95–98.

- Kornhuber H, Deecke L (1965) Changes in the brain potential in voluntary movements and passive movements in man: readiness potential and reafferent potentials. *Pflügers Arch Für Gesamte Physiol Menschen Tiere* 284:1–17.
- Kouneiher F, Charron S, Koechlin E (2009) Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci* 12:939–945.
- Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13:1292–1298.
- Krajbich I, Rangel A (2011) Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc Natl Acad Sci* 108:13852–13857.
- Kringelbach ML (2005) The human orbitofrontal cortex: linking reward to hedonic experience. *Nat Rev Neurosci* 6:691–702.
- Krueger JM, Rector DM, Roy S, Van Dongen HPA, Belenky G, Panksepp J (2008) Sleep as a fundamental property of neuronal assemblies. *Nat Rev Neurosci* 9:910–919.
- Kühn AA, Kempf F, Brücke C, Gaynor Doyle L, Martinez-Torres I, Pogosyan A, Trottenberg T, Kupsch A, Schneider G-H, Hariz MI, Vandenberghe W, Nuttin B, Brown P (2008) High-frequency stimulation of the subthalamic nucleus suppresses oscillatory beta activity in patients with Parkinson's disease in parallel with improvement in motor performance. *J Neurosci Off J Soc Neurosci* 28:6165–6173.
- Kühn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider G-H, Yarrow K, Brown P (2004) Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. *Brain J Neurol* 127:735–746.
- Kühn S, Gallinat J (2012) The neural correlates of subjective pleasantness. *Neuroimage* 61:289–294.
- Kuhtz-Buschbeck JP, Gilster R, Wolff S, Ulmer S, Siebner H, Jansen O (2008) Brain activity is similar during precision and power gripping with light force: An fMRI study. *NeuroImage* 40:1469–1481.
- Kurniawan IT, Guitart-Masip M, Dayan P, Dolan RJ (2013) Effort and valuation in the brain: the effects of anticipation and execution. *J Neurosci Off J Soc Neurosci* 33:6160–6169.
- Kurniawan IT, Guitart-Masip M, Dolan RJ (2011) Dopamine and effort-based decision making. *Front Neurosci* 5:81.
- Kurniawan IT, Seymour B, Talmi D, Yoshida W, Chater N, Dolan RJ (2010) Choosing to make an effort: the role of striatum in signaling physical effort of a chosen action. *J Neurophysiol* 104:313–321.
- LaCroix-Fralish ML, Mogil JS (2009) Progress in Genetic Studies of Pain and Analgesia. *Annu Rev Pharmacol Toxicol* 49:97–121.
- Landolt AS, Milling LS (2011) The efficacy of hypnosis as an intervention for labor and delivery pain: a comprehensive methodological review. *Clin Psychol Rev* 31:1022–1031.
- Lang W (2003) Surface recordings of the Bereitschaftspotential in normals. In: *The Bereitschaftspotential: movement-related cortical potentials*, pp 19–34. New York: Kluwer

Academic Publishers. Available at: <http://www.loc.gov/catdir/toc/fy036/2002040791.html> [Accessed February 19, 2010].

- Lea SEG, Webley P (2006) Money as tool, money as drug: the biological psychology of a strong incentive. *Behav Brain Sci* 29:161–176; discussion 176–209.
- Lebreton M, Jorge S, Michel V, Thirion B, Pessiglione M (2009) An Automatic Valuation System in the Human Brain: Evidence from Functional Neuroimaging. *Neuron* 64:431–439.
- Lee JH, Durand R, Gradinaru V, Zhang F, Goshen I, Kim D-S, Fenno LE, Ramakrishnan C, Deisseroth K (2010) Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature* 465:788–792.
- Leuthold H, Jentzsch I (2002) Distinguishing neural sources of movement preparation and execution: An electrophysiological analysis. *Biol Psychol* 60:173–198.
- Levy DJ, Glimcher PW (2011) Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *J Neurosci Off J Soc Neurosci* 31:14693–14707.
- Levy DJ, Glimcher PW (2012) The root of all value: a neural common currency for choice. *Curr Opin Neurobiol* 22:1027–1038.
- Libet B, Gleason CA, Wright EW, Pearl DK (1983) Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain J Neurol* 106 (Pt 3):623–642.
- Liljeholm M, O’Doherty JP (2012) Contributions of the striatum to learning, motivation, and performance: an associative account. *Trends Cogn Sci* 16:467–475.
- Lindstedt F, Berrebi J, Greayer E, Lonsdorf TB, Schalling M, Ingvar M, Kosek E (2011) Conditioned Pain Modulation Is Associated with Common Polymorphisms in the Serotonin Transporter Gene. *Plos One* 6:e18252.
- Litvak V, Eusebio A, Jha A, Oostenveld R, Barnes GR, Penny WD, Zrinzo L, Hariz MI, Limousin P, Friston KJ, Brown P (2010) Optimized beamforming for simultaneous MEG and intracranial local field potential recordings in deep brain stimulation patients. *Neuroimage* 50:1578–1588.
- Liu JZ, Dai TH, Sahgal V, Brown RW, Yue GH (2002) Nonlinear cortical modulation of muscle fatigue: a functional MRI study. *Brain Res* 957:320–329.
- Liu JZ, Shan ZY, Zhang LD, Sahgal V, Brown RW, Yue GH (2003) Human Brain Activation During Sustained and Intermittent Submaximal Fatigue Muscle Contractions: An fMRI Study. *J Neurophysiol* 90:300–312.
- Locke HS, Braver TS (2008) Motivational influences on cognitive control: behavior, brain activation, and individual differences. *Cogn Affect Behav Neurosci* 8:99–112.
- Loeb G, Ghez C (2000) The motor unit and muscle action. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 674–694. Mc Graw Hill.
- Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453:869–878.

- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- Logothetis NK, Wandell BA (2004) Interpreting the BOLD Signal. *Annu Rev Physiol* 66:735–769.
- Loubinoux I, Tombari D, Pariente J, Gerdelat-Mas A, Franceries X, Cassol E, Rascol O, Pastor J, Chollet F (2005) Modulation of behavior and cortical motor activity in healthy subjects by a chronic administration of a serotonin enhancer. *NeuroImage* 27:299–313.
- Louie K, Glimcher PW (2010) Separating Value from Choice: Delay Discounting Activity in the Lateral Intraparietal Area. *J Neurosci* 30:5498–5507.
- Ludman CN, Cooper TG, Ploutz-Snyder LL, Potchen EJ, Meyer RA (1996) Force of voluntary exercise does not affect sensorimotor cortex activation as detected by functional MRI at 1.5 T. *Nmr Biomed* 9:228–232.
- Ma L, Chablat D, Bennis F, Zhang W (2009) A new simple dynamic muscle fatigue model and its validation. *Int J Ind Ergon* 39:211–220.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Marcora S (2009) Perception of effort during exercise is independent of afferent feedback from skeletal muscles, heart, and lungs. *J Appl Physiol* 106:2060–2062.
- Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* 164:177–190.
- Martin PG, Weerakkody N, Gandevia SC, Taylor JL (2008) Group III and IV muscle afferents differentially affect the motor cortex and motoneurons in humans. *J Physiol* 586:1277–1289.
- Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic theory*, Oxford University Press.
- Mayville JM, Fuchs A, Kelso JAS (2005) Neuromagnetic motor fields accompanying self-paced rhythmic finger movement at different rates. *Exp Brain Res Exp Hirnforsch Expérimentation Cérébrale* 166:190–199.
- Mazzola L, Isnard J, Peyron R, Mauguière F (2012) Stimulation of the human cortex and the experience of pain: Wilder Penfield’s observations revisited. *Brain J Neurol* 135:631–640.
- Mazzoni P, Hristova A, Krakauer JW (2007) Why don’t we move faster? Parkinson’s disease, movement vigor, and implicit motivation. *J Neurosci Off J Soc Neurosci* 27:7105–7116.
- McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004) Separate neural systems value immediate and delayed monetary rewards. *Science* 306:503–507.
- McGuire JT, Botvinick MM (2010) Prefrontal cortex, cognitive control, and the registration of decision costs. *Proc Natl Acad Sci U S A* 107:7922–7926.
- McIntyre CC, Hahn PJ (2010) Network perspectives on the mechanisms of deep brain stimulation. *Neurobiol Dis* 38:329–337.

- McNeil CJ, Butler JE, Taylor JL, Gandevia SC (2013) Testing the excitability of human motoneurons. *Front Hum Neurosci* 7:152.
- Merton PA (1954) Voluntary strength and fatigue. *J Physiol* 123:553–564.
- Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. *Proc Natl Acad Sci U S A* 110:2641–2646.
- Millet GY, Lepers R (2004) Alterations of neuromuscular function after prolonged running, cycling and skiing exercises. *Sports Med Auckl Nz* 34:105–116.
- Minamimoto T, Hori Y, Richmond BJ (2012) Is working more costly than waiting in monkeys? *Plos One* 7:e48434.
- Mohr C, Binkofski F, Erdmann C, Büchel C, Helmchen C (2005) The anterior cingulate cortex contains distinct areas dissociating external from self-administered painful stimulation: a parametric fMRI study. *Pain* 114:347–357.
- Montague P (2002) Neural Economics and the Biological Substrates of Valuation. *Neuron* 36:265–284.
- Moore A, Collins S, Carroll D, McQuay H, Edwards J (2000) Single dose paracetamol (acetaminophen), with and without codeine, for postoperative pain. *Cochrane Database Syst Rev Online*:CD001547.
- Mori Y (1999) The optimal allocation of time and respiratory metabolism over the dive cycle. *Behav Ecol* 10:155–160.
- Nachev P, Kennard C, Husain M (2008) Functional role of the supplementary and pre-supplementary motor areas. *Nat Rev Neurosci* 9:856–869.
- Nagamine T, Kajola M, Salmelin R, Shibasaki H, Hari R (1996) Movement-related slow cortical magnetic fields and changes of spontaneous MEG- and EEG-brain rhythms. *Electroencephalogr Clin Neurophysiol* 99:274–286.
- Nakai T, Matsuo K, Kato C, Takehara Y, Isoda H, Moriya T, Okada T, Sakahara H (2000) Post-stimulus response in hemodynamics observed by functional magnetic resonance imaging--difference between the primary sensorimotor area and the supplementary motor area. *Magn Reson Imaging* 18:1215–1219.
- Naqvi NH, Bechara A (2009) The hidden island of addiction: the insula. *Trends Neurosci* 32:56–67.
- Nichols T, Hayasaka S (2003) Controlling the familywise error rate in functional neuroimaging: a comparative review. *Stat Methods Med Res* 12:419–446.
- Niv Y, Daw N, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology (Berl)* 191:507–520.
- Niv Y, Daw ND, Dayan P (2005) How fast to work: Response vigor, motivation and tonic dopamine. *Adv Neural Inf Process Syst* 18:1019–1026.

- Noakes TD (2000) Physiological models to understand exercise fatigue and the adaptations that predict or enhance athletic performance. *Scand J Med Sci Sports* 10:123–145.
- Noakes TD (2011) Time to move beyond a brainless exercise physiology: the evidence for complex regulation of human exercise performance. *Appl Physiol Nutr Metab Physiol Appliquée Nutr Métabolisme* 36:23–35.
- Noakes TD, Peltonen JE, Rusko HK (2001) Evidence that a central governor regulates exercise performance during acute hypoxia and hyperoxia. *J Exp Biol* 204:3225–3234.
- Noonan MP, Walton ME, Behrens TEJ, Sallet J, Buckley MJ, Rushworth MFS (2010) Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U S A* 107:20547–20552.
- Nybo L, Moller K, Volianitis S, Nielsen B, Secher NH (2002) Effects of hyperthermia on cerebral blood flow and metabolism during prolonged exercise in humans. *J Appl Physiol* 93:58–64.
- Nybo L, Secher NH (2004) Cerebral perturbations provoked by prolonged exercise. *Prog Neurobiol* 72:223–261.
- O’Connell RG, Dockree PM, Kelly SP (2012) A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* 15:1729–1735.
- Obata T, Liu TT, Miller KL, Luh WM, Wong EC, Frank LR, Buxton RB (2004) Discrepancies between BOLD and flow dynamics in primary and supplementary motor areas: application of the balloon model to the interpretation of BOLD transients. *Neuroimage* 21:144–153.
- Ogawa S, Menon RS, Tank DW, Kim SG, Merkle H, Ellermann JM, Ugurbil K (1993) Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophys J* 64:803–812.
- Oostenveld R, Fries P, Maris E, Schoffelen J-M (2011) FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Comput Intell Neurosci* 2011:1–9.
- Ostrowsky K, Magnin M, Ryvlin P, Isnard J, Guenot M, Mauguière F (2002) Representation of pain and somatic sensation in the human insula: a study of responses to direct electrical cortical stimulation. *Cereb Cortex New York N* 1991 12:376–385.
- Özkurt TE, Butz M, Homburger M, Elben S, Vesper J, Wojtecki L, Schnitzler A (2011) High frequency oscillations in the subthalamic nucleus: a neurophysiological marker of the motor state in Parkinson’s disease. *Exp Neurol* 229:324–331.
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226.
- Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci* 11:95–102.
- Palminteri S, Justo D, Jauffret C, Pavlicek B, Dauta A, Delmaire C, Czernecki V, Karachi C, Capelle L, Durr A, Pessiglione M (2012) Critical Roles for Anterior Insula and Dorsal Striatum in Punishment-Based Avoidance Learning. *Neuron* 76:998–1009.

- Palminteri S, Lebreton M, Worbe Y, Grabli D, Hartmann A, Pessiglione M (2009) Pharmacological modulation of subliminal learning in Parkinson's and Tourette's syndromes. *Proc Natl Acad Sci U S A* 106:19179–19184.
- Pasquereau B, Turner RS (2013) Limited Encoding of Effort by Dopamine Neurons in a Cost-Benefit Trade-off Task. *J Neurosci Off J Soc Neurosci* 33:8288–8300.
- Pearson K, Gordon J (2000) Spinal Reflexes. In: *Principles of neural science*, 4th Edition. (Kandel E, Schwartz J, Jessel T, eds), pp 713–736. Mc Graw Hill.
- Pedersen JR, Johannsen P, Bak CK, Kofoed B, Saermark K, Gjedde A (1998) Origin of Human Motor Readiness Field Linked to Left Middle Frontal Gyrus by MEG and PET. *NeuroImage* 8:214–220.
- Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, Schofield TM, Leff AP (2010) Comparing families of dynamic causal models. *Plos Comput Biol* 6:e1000709.
- Pessiglione M, Schmidt L, Draganski B, Kalisch R, Lau H, Dolan RJ, Frith CD (2007) How the brain translates money into force: a neuroimaging study of subliminal motivation. *Science* 316:904–906.
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD (2006) Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442:1042–1045.
- Peters A, Schweiger U, Pellerin L, Hubold C, Oltmanns KM, Conrad M, Schultes B, Born J, Fehm HL (2004) The selfish brain: competition for energy resources. *Neurosci Biobehav Rev* 28:143–180.
- Peters J, Büchel C (2010) Neural representations of subjective reward value. *Behav Brain Res* 213:135–141.
- Peters J, Büchel C (2011) The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends Cogn Sci* 15:227–239.
- Petrovic P, Kalso E, Petersson KM, Andersson J, Fransson P, Ingvar M (2010) A prefrontal non-opioid mechanism in placebo analgesia. *Pain* 150:59–65.
- Petrovic P, Kalso E, Petersson KM, Ingvar M (2002) Placebo and opioid analgesia-- imaging a shared neuronal network. *Science* 295:1737–1740.
- Petrovic P, Pleger B, Seymour B, Klöppel S, Martino BD, Critchley H, Dolan RJ (2008) Blocking Central Opiate Function Modulates Hedonic Impact and Anterior Cingulate Response to Rewards and Losses. *J Neurosci* 28:10509–10516.
- Peyron R, Laurent B, García-Larrea L (2000) Functional imaging of brain responses to pain. A review and meta-analysis (2000). *Neurophysiol Clin Clin Neurophysiol* 30:263–288.
- Pine A, Seymour B, Roiser JP, Bossaerts P, Friston KJ, Curran HV, Dolan RJ (2009) Encoding of marginal utility across time in the human brain. *J Neurosci Off J Soc Neurosci* 29:9575–9581.
- Plassmann H, O'Doherty JP, Rangel A (2010) Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *J Neurosci Off J Soc Neurosci* 30:10799–10808.

- Pogosyan A, Gaynor LD, Eusebio A, Brown P (2009) Boosting cortical activity at Beta-band frequencies slows movement in humans. *Curr Biol* 19:1637–1641.
- Praamstra P, Schmitz F, Freund H-J, Schnitzler A (1999) Magneto-encephalographic correlates of the lateralized readiness potential. *Cogn Brain Res* 8:77–85.
- Prévost C, Pessiglione M, Météreau E, Cléry-Melin M-L, Dreher J-C (2010) Separate valuation subsystems for delay and effort decision costs. *J Neurosci Off J Soc Neurosci* 30:14080–14090.
- Price J, Cole V, Goodwin GM (2009) Emotional side-effects of selective serotonin reuptake inhibitors: qualitative study. *Br J Psychiatry* 195:211–217.
- Pyke GH (1984) Optimal Foraging Theory: A Critical Review. *Annu Rev Ecol Syst* 15:523–575.
- Pyke GH, Pulliam HR, Charnov E (1977) Optimal foraging: A selective review of theory and tests. Available at: <https://repository.unm.edu/handle/1928/1688> [Accessed May 31, 2013].
- Rainville P, Carrier B, Hofbauer RK, Bushnell MC, Duncan GH (1999) Dissociation of sensory and affective dimensions of pain using hypnotic modulation. *Pain* 82:159–171.
- Ramsøy TZ, Skov M (2010) How genes make up your mind: Individual biological differences and value-based decisions. *J Econ Psychol* 31:818–831.
- Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9:545–556.
- Rangel A, Clithero JA (2012) Value normalization in decision making: theory and evidence. *Curr Opin Neurobiol* 22:970–981.
- Rangel A, Hare T (2010) Neural computations associated with goal-directed choice. *Curr Opin Neurobiol* 20:262–270.
- Rasmussen P, Dawson EA, Nybo L, van Lieshout JJ, Secher NH, Gjedde A (2006) Capillary-oxygenation-level-dependent near-infrared spectrometry in frontal lobe of humans. *J Cereb Blood Flow Metab* 27:1082–1093.
- Rasmussen P, Nielsen J, Overgaard M, Krogh-Madsen R, Gjedde A, Secher NH, Petersen NC (2010) Reduced muscle activation during exercise related to brain oxygenation and metabolism in humans. *J Physiol* 588:1985–1995.
- Raz A, Fan J, Posner MI (2005) Hypnotic suggestion reduces conflict in the human brain. *Proc Natl Acad Sci U S A* 102:9978–9983.
- Richter M, Gendolla GHE (2009) The heart contracts to reward: monetary incentives and prejection period. *Psychophysiology* 46:451–457.
- Rigoux L, Guigon E (2012) A Model of Reward- and Effort-Based Optimal Decision Making and Motor Control. *Plos Comput Biol* 8:e1002716.
- Robbins TW (2007) Shifting and stopping: fronto-striatal substrates, neurochemical modulation and clinical implications. *Philos Trans R Soc B Biol Sci* 362:917–932.

- Rosa MJ, Daunizeau J, Friston KJ (2010) EEG-fMRI integration: a critical review of biophysical modeling and data analysis approaches. *J Integr Neurosci* 9:453–476.
- Ross EZ, Middleton N, Shave R, George K, Nowicky A (2007) Corticomotor excitability contributes to neuromuscular fatigue following marathon running in man. *Exp Physiol* 92:417–426.
- Rudebeck PH, Walton ME, Smyth AN, Bannerman DM, Rushworth MFS (2006) Separate neural pathways process different decision costs. *Nat Neurosci* 9:1161–1168.
- Rushton DN (2002) Electrical stimulation in the treatment of pain. *Disabil Rehabil* 24:407–415.
- Rushworth MFS, Kolling N, Sallet J, Mars RB (2012) Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr Opin Neurobiol* 22:946–955.
- Rushworth MFS, Noonan MP, Boorman ED, Walton ME, Behrens TE (2011) Frontal cortex and reward-guided learning and decision-making. *Neuron* 70:1054–1069.
- Rutledge RB, Dean M, Caplin A, Glimcher PW (2010) Testing the reward prediction error hypothesis with an axiomatic model. *J Neurosci Off J Soc Neurosci* 30:13525–13536.
- Salamone JD, Correa M (2012) The mysterious motivational functions of mesolimbic dopamine. *Neuron* 76:470–485.
- Salamone JD, Correa M, Farrar A, Mingote SM (2007) Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology (Berl)* 191:461–482.
- Satow T, Matsushashi M, Ikeda A, Yamamoto J, Takayama M, Begum T, Mima T, Nagamine T, Mikuni N, Miyamoto S, Hashimoto N, Shibasaki H (2003) Distinct cortical areas for motor preparation and execution in human identified by Bereitschaftspotential recording and ECoG-EMG coherence analysis. *Clin Neurophysiol* 114:1259–1264.
- Scheidegger O, Hess CW, Rösler KM (2010) General features of motor fatigue – a review. *Schweiz Arch Für Neurol Psychiatr* 161:150–153.
- Schmidt L, Cléry-Melin M-L, Lafargue G, Valabregue R, Fossati P, Dubois B, Pessiglione M (2009) Get Aroused and Be Stronger: Emotional Facilitation of Physical Effort in the Human Brain. *J Neurosci* 29:9450–9457.
- Schmidt L, d' Arc BF, Lafargue G, Galanaud D, Czernecki V, Grabli D, Schüpbach M, Hartmann A, Lévy R, Dubois B, Pessiglione M (2008) Disconnecting force from money: effects of basal ganglia damage on incentive motivation. *Brain J Neurol* 131:1303–1310.
- Schmidt L, Lebreton M, Cléry-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural mechanisms underlying motivation of mental versus physical effort. *Plos Biol* 10:e1001266.
- Schnitzler A, Gross J (2005) Normal and pathological oscillatory communication in the brain. *Nat Rev Neurosci* 6:285–296.
- Schoffelen J-M, Oostenveld R, Fries P (2005) Neuronal coherence as a mechanism of effective corticospinal interaction. *Science* 308:111–113.
- Schonberg T, Fox CR, Poldrack RA (2011) Mind the gap: bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends Cogn Sci* 15:11–19.

- Schultz W (2000) Multiple reward signals in the brain. *Nat Rev Neurosci* 1:199–207.
- Schultz W (2006) Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol* 57:87–115.
- Schultz W, Dayan P, Montague PR (1997) A Neural Substrate of Prediction and Reward. *Science* 275:1593–1599.
- Schurger A, Sitt JD, Dehaene S (2012) An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci U S A* 109:E2904–2913.
- Schweighofer N, Shishida K, Han CE, Okamoto Y, Tanaka SC, Yamawaki S, Doya K (2006) Humans can adopt optimal discounting strategy under real-time constraints. *Plos Comput Biol* 2:e152.
- Schweimer J, Saft S, Hauber W (2005) Involvement of catecholamine neurotransmission in the rat anterior cingulate in effort-related decision making. *Behav Neurosci* 119:1687–1692.
- Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* 5:532–546.
- Selen LPJ, Shadlen MN, Wolpert DM (2012) Deliberation in the motor system: reflex gains track evolving evidence leading to a decision. *J Neurosci Off J Soc Neurosci* 32:2276–2286.
- Serences J (2008) Value-Based Modulations in Human Visual Cortex. *Neuron* 60:1169–1181.
- Serretti A, Calati R, Goracci A, Di Simplicio M, Castrogiovanni P, De Ronchi D (2010) Antidepressants in healthy subjects: What are the psychotropic/psychological effects? *Eur Neuropsychopharmacol* 20:433–453.
- Sescousse G, Caldú X, Segura B, Dreher J-C (2013) Processing of primary and secondary rewards: a quantitative meta-analysis and review of human functional neuroimaging studies. *Neurosci Biobehav Rev* 37:681–696.
- Seymour B, Daw N, Dayan P, Singer T, Dolan R (2007a) Differential encoding of losses and gains in the human striatum. *J Neurosci Off J Soc Neurosci* 27:4826–4831.
- Seymour B, McClure SM (2008) Anchors, scales and the relative coding of value in the brain. *Curr Opin Neurobiol* 18:173–178.
- Seymour B, Singer T, Dolan R (2007b) The neurobiology of punishment. *Nat Rev Neurosci* 8:300–311.
- Shadlen MN, Hanks TD, Churchland AK, Kiani R, Yang T (2006) The speed accuracy of a simple perceptual decision: a mathematical primer. In: *Bayesian Brain: Probabilistic Approaches to Neural Coding*, MIT Press. (Doya K, Ishii S, Rao R, Pouget A, eds). Cambridge.
- Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. *Exp Brain Res Exp Hirnforsch Expérimentation Cérébrale* 185:359–381.
- Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih T-Y (2010) Temporal discounting of reward and the cost of time in motor control. *J Neurosci Off J Soc Neurosci* 30:10507–10516.
- Shephard RJ (2009) Is it time to retire the “central governor”? *Sports Med Auckl Nz* 39:709–721.

- Shibasaki H, Hallett M (2006) What is the Bereitschaftspotential? *Clin Neurophysiol* 117:2341–2356.
- Shmuel A, Augath M, Oeltermann A, Logothetis NK (2006) Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. *Nat Neurosci* 9:569–577.
- Sirotin YB, Das A (2009) Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. *Nature* 457:475–479.
- Skoubis PD, Lam HA, Shoblock J, Narayanan S, Maidment NT (2005) Endogenous enkephalins, not endorphins, modulate basal hedonic state in mice. *Eur J Neurosci* 21:1379–1384.
- Smirmaul B de PC (2012) Sense of effort and other unpleasant sensations during exercise: clarifying concepts and mechanisms. *Br J Sports Med* 46:308–311.
- Smith HS (2009) Potential analgesic mechanisms of acetaminophen. *Pain Physician* 12:269–280.
- Spraker MB, Yu H, Corcos DM, Vaillancourt DE (2007) Role of Individual Basal Ganglia Nuclei in Force Amplitude Generation. *J Neurophysiol* 98:821–834.
- St Clair Gibson A, Lambert ML, Noakes TD (2001) Neural control of force output during maximal and submaximal exercise. *Sports Med Auckl Nz* 31:637–650.
- Staaal C, Olesen AE, Andresen T, Arendt-Nielsen L, Drewes AM (2009) Assessing analgesic actions of opioids by experimental pain models in healthy volunteers – an updated review. *Br J Clin Pharmacol* 68:149–168.
- Stahl SM (2008) Chapter 12 - Antidepressants. In: *Stahl's essential psychopharmacology. Neuroscientific basis and practical applications*, 2nd ed. Cambridge University Press.
- Stancák A Jr, Riml A, Pfurtscheller G (1997) The effects of external load on movement-related changes of the sensorimotor EEG rhythms. *Electroencephalogr Clin Neurophysiol* 102:495–504.
- Starkstein SE, Mayberg HS, Preziosi TJ, Andrezejewski P, Leiguarda R, Robinson RG (1992) Reliability, validity, and clinical correlates of apathy in Parkinson's disease. *J Neuropsychiatry Clin Neurosci* 4:134–139.
- Steel P (2007) The nature of procrastination: a meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychol Bull* 133:65–94.
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46:1004–1017.
- Stephens DW (2008) Decision ecology: foraging and the ecology of animal decision making. *Cogn Affect Behav Neurosci* 8:475–484.
- Sugrue LP, Corrado GS, Newsome WT (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat Rev Neurosci* 6:363–375.
- Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009) Intracranial EEG Reveals a Time- and Frequency-Specific Role for the Right Inferior Frontal

- Gyrus and Primary Motor Cortex in Stopping Initiated Responses. *J Neurosci* 29:12675–12685.
- Tachibana Y, Hikosaka O (2012) The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron* 76:826–837.
- Takahashi M, Watanabe Y, Haraguchi T, Kawai T, Yamane G, Abe S, Sakiyama K, Hiraide Y, Lee W, Ide Y, Ishikawa T (2004) Neuromagnetic analysis of the late phase of the readiness field for precise hand movements using magnetoencephalography. *Bull Tokyo Dent Coll* 45:9–17.
- Tallon-Baudry C, Meyniel F, Bourgeois-Gironde S (2011) Fast and automatic activation of an abstract representation of money in the human ventral visual pathway. *Plos One* 6:e28229.
- Talmi D, Dayan P, Kiebel SJ, Frith CD, Dolan RJ (2009) How humans integrate the prospects of pain and reward during choice. *J Neurosci Off J Soc Neurosci* 29:14617–14626.
- Talmi D, Pine A (2012) How costs influence decision values for mixed outcomes. *Front Neurosci* 6:146.
- Tanaka H, Krakauer JW, Qian N (2006) An optimization principle for determining movement duration. *J Neurophysiol* 95:3875–3886.
- Tanaka M, Watanabe Y (2012) Supraspinal regulation of physical fatigue. *Neurosci Biobehav Rev* 36:727–734.
- Taylor JJ, Borckardt JJ, Canterberry M, Li X, Hanlon CA, Brown TR, George MS (2013) Naloxone-Reversible Modulation of Pain Circuitry by Left Prefrontal rTMS. *Neuropsychopharmacol Off Publ Am Coll Neuropsychopharmacol* 38:1189–1197.
- Taylor JL, Gandevia SC (2008) A comparison of central aspects of fatigue in submaximal and maximal voluntary contractions. *J Appl Physiol Bethesda Md* 1985 104:542–550.
- Taylor JL, Todd G, Gandevia SC (2006) Evidence for a supraspinal contribution to human muscle fatigue. *Clin Exp Pharmacol Physiol* 33:400–405.
- Thickbroom GW, Phillips BA, Morris I, Byrnes ML, Mastaglia FL (1998) Isometric force-related activity in sensorimotor cortex measured with functional MRI. *Exp Brain Res* 121:59–64.
- Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive coding of reward value by dopamine neurons. *Science* 307:1642–1645.
- Todorov E (2000) Direct cortical control of muscle activation in voluntary arm movements: a model. *Nat Neurosci* 3:391–398.
- Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7:907–915.
- Todorov E (2006) Optimal control theory. In: *Bayesian brain: probabilistic approaches to neural coding* (Doya K, ed), pp 269–298. MIT Press.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.

- Toma K, Mima T, Matsuoka T, Gerloff C, Ohnishi T, Koshy B, Andres F, Hallett M (2002) Movement rate effect on activation and functional coupling of motor cortical areas. *J Neurophysiol* 88:3377–3385.
- Toms L, McQuay HJ, Derry S, Moore RA (2008) Single dose oral paracetamol (acetaminophen) for postoperative pain in adults. *Cochrane Database Syst Rev Online*:CD004602.
- Torrealba F, Riveros ME, Contreras M, Valdes JL (2012) Histamine and motivation. *Front Syst Neurosci* 6:51.
- Tucker R (2009) The anticipatory regulation of performance: the physiological basis for pacing strategies and the development of a perception-based model for exercise performance. *Br J Sports Med* 43:392–400.
- Turner RS, Desmurget M (2010) Basal ganglia contributions to motor control: a vigorous tutor. *Curr Opin Neurobiol* 20:704–716.
- Tversky A, Kahneman D (1981) The Framing of Decisions and the Psychology of Choice. *Science* 211:453–458.
- Tzagarakis C, Ince NF, Leuthold AC, Pellizzer G (2010) Beta-band activity during motor planning reflects response uncertainty. *J Neurosci Off J Soc Neurosci* 30:11270–11277.
- Uhlhaas PJ, Singer W (2006) Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology. *Neuron* 52:155–168.
- Vaillancourt DE, Mayka MA, Thulborn KR, Corcos DM (2004) Subthalamic nucleus and internal globus pallidus scale with the rate of change of force production in humans. *NeuroImage* 23:175–186.
- Vaillancourt DE, Thulborn KR, Corcos DM (2003) Neural Basis for the Processes That Underlie Visually Guided and Internally Guided Force Control in Humans. *J Neurophysiol* 90:3330–3340.
- Van de Vijver I, Ridderinkhof KR, Cohen MX (2011) Frontal oscillatory dynamics predict feedback learning and action adjustment. *J Cogn Neurosci* 23:4106–4121.
- Van Wijk BCM, Beek PJ, Daffertshofer A (2012) Neural synchrony within the motor system: what have we learned so far? *Front Hum Neurosci* 6:252.
- Vartiainen J, Liljeström M, Koskinen M, Renvall H, Salmelin R (2011) Functional magnetic resonance imaging blood oxygenation level-dependent signal and magnetoencephalography evoked responses yield different neural functionality in reading. *J Neurosci Off J Soc Neurosci* 31:1048–1058.
- Wager TD, Atlas LY, Lindquist MA, Roy M, Woo C-W, Kross E (2013) An fMRI-based neurologic signature of physical pain. *N Engl J Med* 368:1388–1397.
- Wager TD, Rilling JK, Smith EE, Sokolik A, Casey KL, Davidson RJ, Kosslyn SM, Rose RM, Cohen JD (2004) Placebo-Induced Changes in fMRI in the Anticipation and Experience of Pain. *Science* 303:1162–1167.
- Walsh DM, Howe TE, Johnson MI, Sluka KA (2009) Transcutaneous electrical nerve stimulation for acute pain. *Cochrane Database Syst Rev Online*:CD006142.

- Walton ME, Bannerman DM, Alterescu K, Rushworth MFS (2003) Functional Specialization within Medial Frontal Cortex of the Anterior Cingulate for Evaluating Effort-Related Decisions. *J Neurosci* 23:6475–6479.
- Walton ME, Bannerman DM, Rushworth MFS (2002) The role of rat medial frontal cortex in effort-based decision making. *J Neurosci Off J Soc Neurosci* 22:10996–11003.
- Walton ME, Kennerley SW, Bannerman DM, Phillips PEM, Rushworth MFS (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Netw* 19:1302–1314.
- Wang AY, Miura K, Uchida N (2013) The dorsomedial striatum encodes net expected return, critical for energizing performance vigor. *Nat Neurosci* 16:639–647.
- Wang X-J (2012) Neural dynamics and circuit mechanisms of decision-making. *Curr Opin Neurobiol* 22:1039–1046.
- Wark DM (2008) What we can do with hypnosis: a brief note. *Am J Clin Hypn* 51:29–36.
- Way BM, Taylor SE, Eisenberger NI (2009) Variation in the mu-opioid receptor gene (OPRM1) is associated with dispositional and neural sensitivity to social rejection. *Proc Natl Acad Sci U S A* 106:15079–15084.
- Weir JP, Beck TW, Cramer JT, Housh TJ (2006) Is fatigue all in your head? A critical review of the central governor model. *Br J Sports Med* 40:573–586; discussion 586.
- Wikenheiser AM, Stephens DW, Redish AD (2013) Subjective costs drive overly patient foraging strategies in rats on an intertemporal foraging task. *Proc Natl Acad Sci U S A* 110:8308–8313.
- Williams D, Kühn A, Kupsch A, Tijssen M, van Bruggen G, Speelman H, Hotton G, Loukas C, Brown P (2005) The relationship between oscillatory activity and motor reaction time in the parkinsonian subthalamic nucleus. *Eur J Neurosci* 21:249–258.
- Williamson JW (2010) The relevance of central command for the neural cardiovascular control of exercise. *Exp Physiol* 95:1043–1048.
- Wolpert DM, Ghahramani Z (2000) Computational principles of movement neuroscience. *Nat Neurosci* 3:1212–1217.
- Wu S-W, Delgado MR, Maloney LT (2009) Economic decision-making compared with an equivalent motor task. *Proc Natl Acad Sci U S A* 106:6088–6093.
- Wyart V, de Gardelle V, Scholl J, Summerfield C (2012) Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* 76:847–858.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.
- Zubieta J-K, Bueller JA, Jackson LR, Scott DJ, Xu Y, Koeppe RA, Nichols TE, Stohler CS (2005) Placebo effects mediated by endogenous opioid activity on mu-opioid receptors. *J Neurosci Off J Soc Neurosci* 25:7754–7762.

Zubieta J-K, Smith YR, Bueller JA, Xu Y, Kilbourn MR, Jewett DM, Meyer CR, Koeppe RA, Stohler CS (2001) Regional Mu Opioid Receptor Regulation of Sensory and Affective Dimensions of Pain. *Science* 293:311–315.



## 6 Appendices

### Article 1

Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. *Proc Natl Acad Sci USA* 110:2641–2646.

### Article 2

Meyniel F, Pessiglione M, Better get back to work: a role for motor beta de-synchronization in incentive motivation (in revision)

### Article 3

Meyniel F, Safra L, Pessiglione M, How the brain decides when to work and when to rest: evidence for implicit cost-evidence monitoring (submission)



# Neurocomputational account of how the human brain decides when to have a break

Florent Meynier<sup>a,b,c</sup>, Claire Sergent<sup>b,c,d</sup>, Lionel Rigoux<sup>a,b,c</sup>, Jean Daunizeau<sup>a,b,c,e</sup>, and Mathias Pessiglione<sup>a,b,c,e,1</sup>

<sup>a</sup>Motivation, Brain and Behavior Laboratory, <sup>d</sup>Physiological Investigations of Clinical Normal and Impaired Cognition Laboratory, <sup>e</sup>Centre de NeuroImagerie de Recherche, <sup>b</sup>Brain and Spine Institute, Hôpital de la Pitié-Salpêtrière, 75013 Paris, France; and <sup>c</sup>Centre National de la Recherche Scientifique Unité Mixte de Recherche 7225, Institut National de la Santé et de la Recherche Médicale Unité Mixte de Recherche 975, Université Pierre et Marie Curie–Paris 6, 75005 Paris, France

Edited by Marcus E. Raichle, Washington University in St. Louis, St. Louis, MO, and approved December 11, 2012 (received for review July 12, 2012)

**No pain, no gain: cost–benefit trade-off has been formalized in classical decision theory to account for how we choose whether to engage effort. However, how the brain decides when to have breaks in the course of effort production remains poorly understood. We propose that decisions to cease and resume work are triggered by a cost evidence accumulation signal reaching upper and lower bounds, respectively. We developed a task in which participants are free to exert a physical effort knowing that their payoff would be proportional to their effort duration. Functional MRI and magnetoencephalography recordings conjointly revealed that the theoretical cost evidence accumulation signal was expressed in proprioceptive regions (bilateral posterior insula). Furthermore, the slopes and bounds of the accumulation process were adapted to the difficulty of the task and the money at stake. Cost evidence accumulation might therefore provide a dynamical mechanistic account of how the human brain maximizes benefits while preventing exhaustion.**

accumulation model | combined fMRI–MEG | decision making | effort/benefit trade-off | fatigue

Should we have a break? The question of when we should stop ongoing work and when we should resume work again has to be solved every day. The problem can be reduced to a trade-off between the costs and benefits of effort exertion, which has been extensively investigated in the decision-making literature (1–4). However, standard decision theory only considers the question of whether to engage an action at a specific time point and does not account for the temporal dynamics of effort allocation. The temporal dynamics are difficult to determine beforehand because prior information about costs is usually limited. For instance, when people have to move a refrigerator up some stairs, they rarely decide in advance the number and duration of breaks. Although central for understanding high-level control of behavior, the issue of how the human brain monitors effort production online so as to make decisions about breaks is virtually unexplored.

William James (p. 323 in ref. 5) provided insightful intuitions about the underlying mechanisms: “Ordinarily, we stop when we meet the first effective layer, so to call it, of fatigue.... But, if an unusual necessity forces us to press onward, a surprising thing occurs. The fatigue gets worse up to a certain critical point, when gradually or suddenly it passes away.... We have evidently tapped a level of new energy.” According to James, writing in the early 20th century, “psychologists ignore (this) conception altogether”; a century later, we intended to take it seriously. We retained three key features of James’ conception: (i) there is a signal analog to fatigue that accumulates during effort exertion, (ii) the decision to stop is triggered by the signal reaching an upper bound, and (iii) the bound can be shifted depending on circumstances. To explain not only effort but also rest termination, we added the hypothesis that the same signal triggers the decision to resume work when reaching a lower bound.

The first aim of the present studies was to examine whether such signal, with both the waning and waxing components, is indeed represented in the human brain. We called this signal “cost evidence” to avoid making any strong assumption about how it relates to conscious subjective sensations such as fatigue or pain.

This appellation also has the advantage of drawing an explicit link with both the literature on cost-based decision making (1–4) and the literature on evidence accumulation in the perceptual domain (6, 7). We note, however, that our signal is more complex than in classical perception studies, because it includes not only the accumulation but also the dissipation process. The second aim was to uncover how the putative cost evidence signal would be adapted to expected benefit and task difficulty, so as to maximize payoff while preventing exhaustion. In particular, we tested James’ idea that motivation can push back the bounds against alternative models where it plays on accumulation slopes.

To investigate the process of cost evidence accumulation, we developed a paradigm where subjects are free to exert or not exert an effort over a long period (30 s). The paradigm was adapted from previous experiments that demonstrated the implication of the ventral striato-pallidal complex in incentive motivation, i.e., in energizing behavioral performance as a function of the reward at stake (8, 9). The task (Fig. 1) involves participants squeezing a handgrip to win a given amount of money, which was manipulated to vary the expected benefit of effort exertion. The key difference to previous tasks is that the payoff was proportional to the time spent above a given force level, which was manipulated to vary the cost of effort exertion. Thus, we had two orthogonal factors forming a three by three design: monetary incentive ( $I$ : 10, 20, or 50 cents) and effort difficulty ( $D$ : 70%, 80%, or 90% of the subject’s maximal force). The incentive was displayed before each trial but the difficulty was not made explicit, such that subjects had to experience cost and adjust their behavior online.

We first examined at the behavioral level how these factors impacted the free parameters of the accumulation model: accumulation slope during effort ( $S_e$ ), dissipation slope during rest ( $S_r$ ), and amplitude between bounds ( $A$ ). Then we used two complementary functional neuroimaging techniques to search for brain activity signaling cost evidence: functional magnetic resonance imaging (fMRI) for an accurate localization of the putative cost evidence signal and magnetoencephalography (MEG) for a precise characterization of its temporal dynamics.

## Results

**Behavioral Results.** Grip force was analyzed conjointly for the two groups who participated in the fMRI ( $n = 19$ ) and MEG ( $n = 19$ ) studies, because they exhibited similar behavior. As expected, subjects alternated periods of rest and effort in the course of a trial (see example in Fig. 24). The duration of every single effort and rest period (not the total over the trial) was extracted to assess incentive and difficulty effects. Incentives significantly

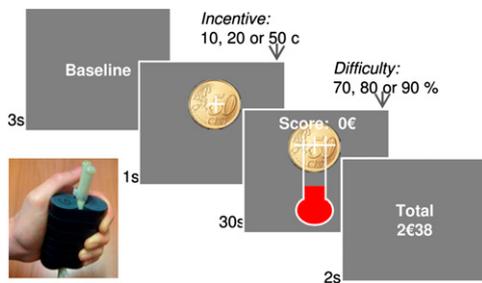
Author contributions: F.M. and M.P. designed research; F.M. performed research; L.R. contributed new reagents/analytic tools; F.M., C.S., J.D., and M.P. analyzed data; and F.M. and M.P. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

<sup>1</sup>To whom correspondence should be addressed. E-mail: mathias.pessiglione@gmail.com.

This article contains supporting information online at [www.pnas.org/lookup/suppl/doi:10.1073/pnas.1211925110/-DCSupplemental](http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1211925110/-DCSupplemental).



**Fig. 1.** Task. The illustrated screenshots were successively presented every trial. When the thermometer image was displayed, participants had to squeeze a handgrip to win as much money as possible. Subjects were provided online feedback on force level and cumulative payoff. The payoff was only increased when force level was above the target bar, at a constant rate proportional to the monetary incentive. Two factors were manipulated over trials: the incentive (10, 20, or 50 cents), which was explicitly indicated as a coin image, and the difficulty, i.e., the force required to reach the target bar (70%, 80%, or 90% of maximal force), which remained implicit. The last screen indicated the money won so far, summed over all preceding trials.

affected the duration of both effort [ $F_{(2,72)} = 25.3$ ;  $P = 1.6 \times 10^{-6}$ ] and rest [ $F_{(2,72)} = 25.2$ ;  $P = 4.9 \times 10^{-7}$ ] periods, whereas difficulty only affected effort duration [ $F_{(2,72)} = 42.8$ ;  $P = 2.8 \times 10^{-12}$ ], not rest duration [ $F_{(2,72)} = 0.1$ ;  $P = 0.86$ ]. There was no significant incentive by difficulty interaction, neither for effort [ $F_{(4,144)} = 1.7$ ;  $P = 0.18$ ] nor for rest [ $F_{(4,144)} = 1.4$ ;  $P = 0.26$ ] duration. Thus, subjects spent more time squeezing and less time resting for higher incentives, and less time squeezing with higher difficulty (Fig. 2B).

We then examined which free parameters of the accumulation model best explained the observed effects on effort and rest durations (effort time,  $T_e = A/Se$ , and rest time,  $T_r = A/Sr$ ). Each free parameter ( $Se$ ,  $Sr$ , and  $A$ ) was written as a linear combination of the experimental factors ( $I$  and  $D$ ), for instance  $A = A_{\text{mean}} + A_I \cdot I + A_D \cdot D$ . Note that we did not include any interaction term because there was no significant incentive by difficulty interaction in our data. Each of the three parameters could in principle be modulated or not modulated by each of the two factors, leading to a total of 64 linear models (Fig. S1, Upper). Only 20 of the 64 possible models could a priori produce the observed behavioral pattern. These 20 models were estimated and compared using Bayesian model selection (BMS). The same best model was identified separately in the fMRI and MEG groups, with exceedance probability  $x_p = 0.97$  and  $x_p = 0.84$ , respectively (Fig. S1, Lower). In this model, incentives play on both  $Sr$  and  $A$ —they accelerate cost evidence dissipation during rest and expend the amplitude between bounds—whereas difficulty only plays on  $Se$ —only accelerates cost evidence accumulation during effort.

**fMRI Results.** We used fMRI during task performance to assess (i) whether some brain activity is correlated with cost evidence accumulation during effort and dissipation during rest and (ii) whether the amplitude of this putative brain signal is modulated by incentives.

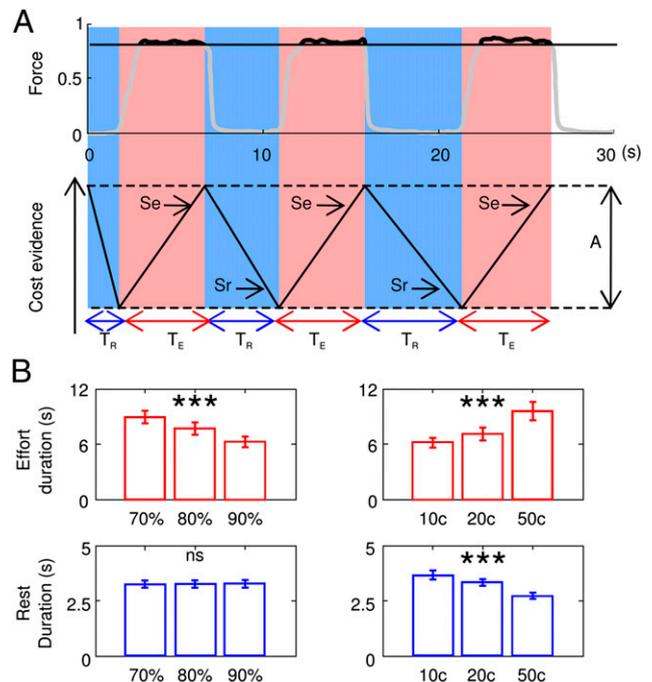
We estimated a first general linear model (GLM1) including cost evidence as a parametric modulator of neural activity at every time point, with constant amplitude throughout all conditions (Fig. 3A). This cost evidence signal was significantly expressed (surviving both voxel- and cluster-wise whole-brain correction) in the bilateral posterior insula (secondary somatosensory cortex SII) and the ventromedial thalamus (Fig. 3B and Table S1). These two brain regions are considered components of the so-called pain matrix and more generally of the proprioception network (10, 11). We extracted them together to form a single region of interest (ROI) that we used in all of following analyses. Note that the same maps were obtained whether or not the exerted force

level was included as an additional regressor in the GLM (compare Tables S1 and S2).

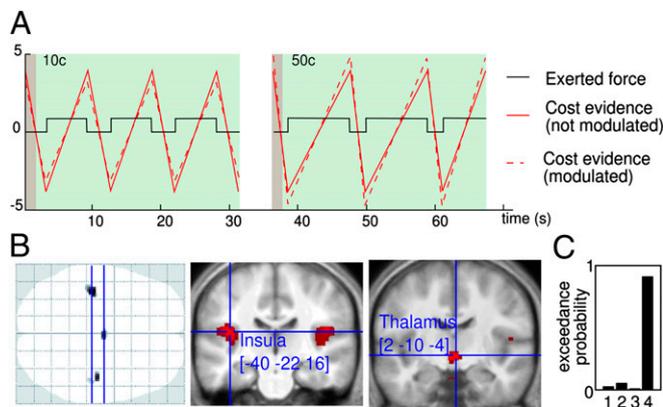
Three additional GLMs were built to account for amplitude modulation by incentives with changes in the upper bound (GLM2), lower bound (GLM3), or both bounds (GLM4). BMS analysis (Fig. 3C) indicated that the activity extracted from the ROI was best explained (with  $x_p > 0.99$ ) by GLM including amplitude modulation (GLM2-4 versus GLM1), i.e., GLMs that were not used to identify the ROI. Among the three possible modulations, changing both bounds (as illustrated in Fig. 3A) was the most probable ( $x_p > 0.96$  for GLM4). Thus, fMRI data revealed that proprioceptive regions continuously signal cost evidence over effort and rest periods. Additionally, they showed that the two bounds triggering effort cessation and return are moved apart when incentives are increased, an effect that could not be inferred from behavioral data.

**MEG Results.** We used MEG to confirm the conclusions drawn from the fMRI study with a reverse approach: instead of a model-driven approach showing that a theoretical signal fits brain activity, we followed a data-driven approach to show that brain activity fulfills theoretical predictions. More specifically we assessed (i) whether scalp activity arising from the ROI sources was ramping up and down during effort and rest periods and (ii) whether incentive and difficulty effects on bounds and slopes conformed to the model optimized on behavioral and fMRI data.

MEG time series were epoched into behaviorally defined rest and effort periods, which were resampled to a same duration and averaged over conditions and subjects. The principal component analysis (PCA) performed on this grand average revealed that



**Fig. 2.** Behavioral results. (A) Example recording of the force exerted during one trial. Three effort (red shading) and three rest (blue shading) epochs could be defined. Force level is shown in black (not gray) when rewarded, i.e., when above the target level (here, 80% of maximal force). The theoretical cost evidence that was hypothesized to underpin effort production is shown below.  $A$ , amplitude between bounds;  $Sr$ , dissipation slope during rest;  $Se$ , accumulation slope during effort;  $T_e$ , effort time;  $T_r$ , rest time. (B) Average data pooled over the MEG and fMRI studies, sorted by incentive and difficulty levels. The bars are mean effort and rest epoch durations, and error bars are the intersubject SEs. Significance of group-level ANOVA main effects: \*\*\* $P < 0.0005$ , \*\* $P < 0.005$ , \* $P < 0.05$ .



**Fig. 3.** fMRI results. (A) Example of two successive trials with the corresponding cost evidence modeled for fMRI data analysis. The green and gray shading indicate incentive display and effort exertion periods. The incentive was 10 cents in the left trial and 50 cents in the right trial. The exerted force is shown in black and the cost evidence in red. Two alternative cost evidence regressors are illustrated: one with a constant amplitude (solid line) and one with both bounds modulated by the incentive (dashed line). (B) Neural correlates of cost evidence. The statistical parametric maps show brain regions where activity was significantly correlated with cost evidence with constant amplitude. Statistical threshold was set at  $P < 0.05$  with voxel-wise (axial projection plan on the left) or cluster-wise (coronal slices on the right) family-wise error correction for multiple comparisons over the entire brain. The coronal slices were taken along the planes indicated by the blue lines on the glass brain. The  $[x\ y\ z]$  peak coordinates refer to the Montreal Neurological Institute (MNI) space. (C) Modulation of cost evidence amplitude. The bar graph represents the result of a Bayesian model selection comparing the fit of different cost evidence regressors to the activity recorded in the significant clusters shown on slices (bilateral posterior insula and ventromedial thalamus). The cost evidence regressors differed on which bound was impacted by incentives: 1 = none, 2 = upper bound, 3 = lower bound, and 4 = both bounds.

both effort and rest periods were dominated by ramping activity (Fig. 4A). Indeed, the first components were linear variations that accounted for most of the signal variance (91% for effort and 72% for rest periods). The sources of these first components

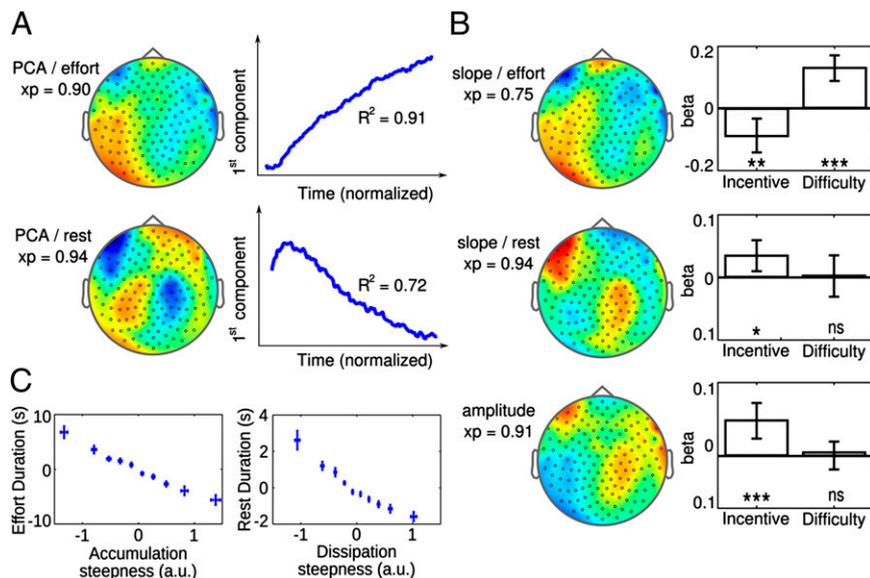
were reconstructed subject-wise with a minimum-norm procedure, either unconstrained or informed by setting priors on the fMRI-based ROI. BMS showed that the reconstruction using the anatomical priors had a much higher exceedance probability ( $x_p = 0.90$  for effort and  $x_p = 0.94$  for rest). Thus, the effort and rest ramping signals that dominated scalp activity were most likely to arise from the same sources: the ROI identified with fMRI. This is consistent with the idea that these regions (posterior insula and ventromedial thalamus) generate an accumulation signal throughout task performance.

We then returned to scalp raw time series (no resampling and averaging) and estimated accumulation and dissipation slopes using linear regression for every single effort and rest epoch on each channel (Fig. 4B). In each subject, slopes were averaged over epochs and conditions, for effort and rest separately. The sources of slope topography were reconstructed subject-wise, using either unconstrained or informed minimum-norm procedure as above. The BMS showed again that including fMRI-based priors largely improved reconstruction plausibility ( $x_p = 0.75$  for effort and  $x_p = 0.94$  for rest).

Individual reconstruction matrices were then used to estimate accumulation and dissipation slopes in the source space for each epoch and subject. The variations of these slopes were fitted with a linear model including incentive and difficulty levels as regressors. Incentives affected both accumulation slope during effort and dissipation slope during rest ( $P < 0.004$  and  $P < 0.008$ , respectively), whereas difficulty only impacted accumulation, not dissipation slope ( $P < 4 \times 10^{-6}$  and  $P > 0.9$ , respectively). Thus, the modulations identified from behavior (incentive increasing  $S_r$  and difficulty increasing  $S_e$ ) were confirmed and an additional modulation (incentive decreasing  $S_e$ ) was revealed. The effect of this additional modulation is to prolong effort periods for higher incentives, which was so far entirely imputed to larger amplitude between bounds.

To assess whether the amplitude of the accumulation signal identified in MEG activity was modulated, we fitted a single V shape to every contiguous rest-effort duplet. Regression coefficients were taken as amplitude estimates and were submitted to the same analysis as slopes. The source reconstruction for these amplitudes as well was much more plausible when including the fMRI-based ROI as priors ( $x_p = 0.91$ ). The linear regression

**Fig. 4.** MEG results. (A) PCA performed on average activity for effort and rest epochs separately. The scalp topography, time series, and proportions of variance explained ( $R^2$  statistics) correspond to the first component. (B) Regression analysis performed on every single event to estimate accumulation slope (by fitting a line to rest and effort epochs separately) and amplitude (by fitting a V shape to two consecutive rest and effort epochs). For each analysis, the scalp topography shows the slopes and amplitudes averaged over trials and subjects. The bar graph represents the coefficients ( $\beta$  values) obtained for the two experimental factors (incentive and difficulty) with a linear regression model fitted on the slopes and amplitudes reconstructed within the regions identified with fMRI. Error bars are intersubject 5% confidence intervals. Significance of group-level  $t$  tests: \*\*\* $P < 0.0005$ , \*\* $P < 0.005$ , \* $P < 0.05$ . For each scalp topography, the  $x_p$  value denotes the exceedance probability of the source reconstruction model that included as priors the regions identified with fMRI. (C) Correlation between residual durations and accumulation or dissipation steepness, obtained by regressing out incentive and difficulty effects. Durations were defined from behavioral data and slopes from MEG data reconstructed in the source space. Effort duration was plotted against accumulation steepness (Left) and rest duration against dissipation steepness (Right). For illustration, data were binned into deciles in every subject. The dots represent intersubject means  $\pm$  SEs for the 10 bins. Note that we use steepness instead of slope for the accumulation signal (power) reconstructed in the source space because it is not signed.



performed on amplitudes in the source space showed a significant increase with higher incentives but no significant modulation by difficulty ( $P < 3 \times 10^{-4}$  and  $P > 0.65$ , respectively). Thus, MEG data revealed that incentives in fact affected all three parameters of the accumulation model ( $Se$ ,  $Sr$ , and  $A$ ), whereas the difficulty effect remained relatively specific (only impacting  $Se$ ).

Additionally, we tested whether incentive and difficulty effects would also be observed across subjects. Between-subject correlations between mean durations and accumulation parameters replicated all four effects that were found to explain the modulation of behavior by incentive and difficulty levels (Fig. S2). Thus, subjects with greater effects on accumulation/dissipation slope or amplitude had greater effects on effort/rest duration.

The above analyses suggest that incentive and difficulty effects on effort and rest durations are underpinned by modulation of cost evidence accumulation slopes and bounds. However, a strong prediction of the model is that decisions to engage and terminate effort are triggered by the cost evidence signal reaching a predetermined threshold. This implies that effort and rest durations should be correlated with accumulation and dissipation slopes, respectively, even when the correlation induced by our experimental manipulation is removed. To assess this prediction, we regressed the variance related to incentive and difficulty levels out of the reconstructed signal and duration, and tested the correlation between residual steepness and duration (Fig. 4C). The correlation across trials was negative for every subject and for both effort and rest periods: the steeper the slope, the shorter the duration. Correlation coefficients were highly significant at the group level for both effort ( $P = 1.4 \times 10^{-8}$ ) and rest periods ( $P = 5.9 \times 10^{-11}$ ).

Taken together, fMRI and MEG findings demonstrated that cost evidence is indeed tracked in proprioceptive brain regions and that the impact of potential benefits on the accumulation process is more complex than suggested solely on the basis of behavior (see illustration in Fig. S3).

## Discussion

We addressed the issue of how the human brain dynamically allocates effort over time, depending on costs and benefits. We found a brain signal that linearly accumulates cost evidence during effort production and dissipates at rest. The observed decisions to stop and restart effort production corresponded to this cost evidence signal reaching upper and lower bounds. We argue that such a mechanism is adaptive because, contrary to benefit estimates, cost estimates can be refined during the course of action, using the information about how much work has been effectively exerted.

In addition, it constitutes an application to the proprioception domain of the accumulation principle that has been implicated in other processes, such as visual perception (6, 7) or subjective valuation (12–14), and interpreted as a process extracting information from a noisy input. However, there are several important differences between standard perceptual evidence accumulation and the process investigated in our study. Here, the functional role of the accumulation process would consist in adjusting effort and rest durations, so as to maximize the payoff while avoiding the peripheral and central damages that can result from prolonged exercise (15–17). It remains unclear whether our accumulation signal reflects an integration process because we do not know the sensory input. It could be a stationary signal related to muscle contraction (i.e., force level), or a signal related to any metabolic or physiological variable that would itself fluctuate over effort and rest periods. Also, the question of how the accumulation signal is reinitialized for subsequent trials is generally not addressed, or at least not interpreted functionally, in perception studies. In our case, the dissipation process would have a functional significance: it would indicate when the body is available for the next effort, and hence when effort should be resumed to maximize benefits. Because of this dissipation, our cost evidence signal cannot be considered as reflecting a pure accumulation (which would only increase). One interpretation of the dissipation

process is that the accumulation is leaky, meaning that the effort is progressively forgotten during rest. Another interpretation is that a control input is subtracted to the sensory input to reset the system. In any case, the global cost evidence signal could be interpreted in terms of predictive coding, as an estimate of exhaustion probability (i.e., temporal proximity).

This complex neural signal with precise spatiotemporal characteristics was identified with both fMRI and MEG using standard analytical tools, which validates the usual assumptions about the relationships between electromagnetic and hemodynamic activity (18–22). Due to its high temporal resolution, MEG brought evidence that the electromagnetic activity emitted by our regions of interest indeed followed the neural dynamics that was modeled and convolved with hemodynamic response to fit fMRI data. In addition, due to its high spatial resolution, fMRI confirmed that our theoretical cost evidence signal was indeed represented in proprioceptive brain regions. The sign of the correlation observed with fMRI, denoting a signal ramping up during effort and down during rest, favored an interpretation in terms of cost evidence accumulation. This was important because, in principle, balancing effort and rest could rely on a signal representing energy dissipation (i.e., the available resource), not cost accumulation (i.e., the expended resource).

Our interpretation was confirmed by the signal being encoded in regions pertaining to the pain matrix, such as the posterior insula and ventral thalamus (10, 11). More precisely, the main activation foci were located in the operculum parietale area 1, which has been implicated in high-order somatosensory processing, in connection with both the ventral thalamus and parietal network (23). The classical pain matrix additionally includes midcingulate regions (24, 25), which were also activated in relation to cost evidence in our results but slightly below the statistical threshold for significance. Interestingly, direct electrical stimulation of the posterior insula was shown to induce painful sensations (26). However, we cannot infer from brain localization alone that the cost evidence variable can be equated to a subjective pain sensation. Indeed, different functions have been assigned to this brain network and particularly to the posterior insula, such as the monitoring of bodily states (27–29).

Let us emphasize that the signal labeled here as cost evidence did not mirror the behavioral output. Its linear fluctuations, dipping when effort starts and peaking when effort ends, were decorrelated from the force produced, which followed a boxcar dynamics. Thus, contrary to the behavioral output, the cost evidence signal spanned the same range of values in effort and rest periods. We believe that the dissipation at rest arises from an active process rather than from a passive relaxation. This is because, at the beginning of trials when subjects had not yet squeezed the handgrip, the cost evidence signal was first brought down to the lower bound in anticipation of effort exertion. This resembles the motor readiness potential that is known to precede limb movements by a few seconds. However, this readiness potential has been localized in motor cortical areas and not in deep proprioceptive regions (30, 31). In addition, the complex modulation of our cost evidence signal by monetary incentives suggests that it is not merely related to motor output.

The model selected from our data suggests that difficulty and incentive levels have computationally distinct impacts on the accumulation process that underpins effort allocation. Task difficulty probably impacted the behavior on the fly, as it was not explicitly mentioned to subjects, whereas the effects of monetary incentives, which were explicitly indicated at trial start, could be regarded as a strategic adjustment. Consistently, difficulty effects specifically manifested as an increased accumulation rate during effort, leaving unchanged the dissipation rate during rest. This was related to the nontrivial behavioral observation that rest duration did not change with task difficulty.

Incentive effects were twofold: (i) they slowed accumulation and speeded dissipation of cost evidence and (ii) they moved the bounds within which cost evidence fluctuates. The first process

might reflect the intervention of an opponent motivation signal that would be continuously subtracted to cost evidence throughout effort and rest periods. This signal might come from brain regions involved in reward processing, or in the trade-off between reward and effort, such as the ventral striatum, the anterior cingulate cortex, or the ventromedial prefrontal cortex (9, 32, 33). The second process might implement the intuitive psychological phenomenon that, when motivated, we literally push back our limits, allowing our body to work closer from exhaustion. It could be explained by reward-related regions adjusting decision thresholds in regions that are downstream to the posterior insula in the chain leading to motor outputs. However, these speculative mechanisms that might adjust accumulation parameters to expected benefits would require further investigation.

Note that incentive effects argue against the possibility that the signal might encode money and not cost accumulation, because the slope observed during effort decreased, not increased, with higher incentives. These effects are also consistent with reports that placebo analgesia reduce insular and thalamic responses to pain stimulation (34), and therefore that the brain can indeed adjust the sensitivity of these regions depending on expectations. More generally, our conclusions are in line with the “central governor” model, which supposes an anticipatory regulation of exercise performance, as opposed to catastrophic models, which attribute performance cessation to homeostatic failure (35, 36). They extend the theory, which was meant to explain how athletes pace their running on a treadmill, to the problem of when people have breaks during work.

Thus, our findings provide empirical evidence supporting the intuition, partly specified by William James a century ago (5), that effort allocation might be controlled online using an accumulation signal reaching bounds. James formulated this idea in psychological terms based on his own introspection; here, we propose a paradigm that allows probing this psychological concept and we reveal a possible implementation at the neural level. The relationship between the two description levels remains to be specified: we have not established yet whether the cost evidence signal is related to what we subjectively perceive as fatigue or pain. Another unresolved issue is whether this signal is indeed causally involved in the decision to cease and resume effort production. Further experiments would be needed to address these remaining issues. For instance, analgesic treatments may help establish causal links between cost evidence accumulation, subjective pain, and online effort allocation.

## Materials and Methods

See *SI Materials and Methods* for behavioral and imaging data acquisition settings.

**Subjects.** The study was approved by the Pitié-Salpêtrière Hospital ethics committee. All subjects were recruited via e-mail within an academic database and gave informed consent before participation in the study. They were right-handed, between 20 and 30 y old, and had normal vision, no history of neurological or psychiatric disease, and no contraindication to MRI. Twenty subjects (eight males; age,  $23.6 \pm 0.6$  y) were included in the fMRI study and 19 in the MEG study (eight males; age,  $24.9 \pm 0.7$  y). One subject in the fMRI study was excluded from the analysis because of calibration issues. Subjects believed that the money won while performing the task would be their remuneration for participating, but eventually their payoff was rounded up to a fixed amount (100€).

**Behavioral Task.** Before starting task performance, we measured the maximal force for each hand, following published guidelines (37). Participants were verbally encouraged to squeeze continuously as hard as they could, until a line growing in proportion to their force reached a target displayed on a computer screen. Maximal force was defined as the average, over the last half of the squeezing period, of data points exceeding the median force. Then subjects were provided a real-time feedback about the force produced on the handgrip, which appeared as a fluid level moving up and down within a thermometer, the maximal force being indicated as a horizontal bar at the top. Subjects were asked to try outreaching the bar and state whether it truly

corresponded to their maximal force. If not, the calibration procedure was repeated.

Task sessions included nine trials corresponding to the nine cells of the factorial design (three incentive by three difficulty conditions), which were presented in a random order. Subjects performed eight sessions in total, switching hands as instructed between sessions to avoid muscle exhaustion. After baseline measurement of the pressure at rest, each trial started by revealing the monetary incentive with a coin image (10, 20, or 50 cents) displayed for 1 s. Then subjects had 30 s to win as much money as possible. They knew that the payoff was proportional to both the incentive and the time spent above the target bar, which was always placed at the same height in the thermometer. The force needed to reach the bar (70%, 80%, or 90% of subject's maximal force), i.e., trial difficulty, was not indicated to subjects. Subjects only knew that task difficulty would vary across trials. They were provided with online feedback on both the exerted force (with a fluid level moving up and down within a thermometer) and the trial-wise cumulative payoff (with a counter displayed above the thermometer). The counter was only started when fluid level was above the target bar, with a rate proportional to the current incentive. The fluid had the same luminance as the background to avoid confounding force level with basic visual features. Each trial ended with a 2-s display of the session-wise cumulative payoff.

**Behavioral Data Analysis.** Effort onsets and offsets were identified offline with an algorithm using the same two criteria for all conditions: (i) force temporal derivative higher than 1 SD and (ii) force level lower (for effort onset) or higher (for effort offset) than half the maximal force. The first rest period started with coin presentation and the subsequent effort and rest periods were defined by force onsets and offsets. Effort and rest period durations were separately analyzed using a repeated-measure ANOVA (R software; Companion to Applied Regression library of John Fox, McMaster University, Hamilton, ON, Canada), with incentive and difficulty as factors of interest. The *P* values reported for these repeated-measure ANOVAs integrate the conservative Greenhouse–Geisser correction.

The accumulation model was formalized with the following observation equations:  $Te = \frac{A}{Se}$ ,  $Tr = \frac{A}{Sr}$  with  $\begin{cases} A = A_{\text{mean}} + A_I + A_D D \\ Se = Se_{\text{mean}} + Se_I + Se_D D \\ Sr = Sr_{\text{mean}} + Sr_I + Sr_D D \end{cases}$

where cost evidence variations have an amplitude *A*, a slope *Sr* during rest and *Se* during effort; *Te* and *Tr* are the effort and rest period durations; *I* and *D* are the z-scored incentive and difficulty levels. Each experimental factor (incentive and difficulty) could in principle affect each free parameter (*A*, *Se*, and *Sr*). We used BMS to test whether each possible modulation (*A*, *A<sub>D</sub>*, *Se<sub>I</sub>*, *Se<sub>D</sub>*, *Sr<sub>I</sub>*, and *Sr<sub>D</sub>*) was significant. There were  $2^6 = 64$  possible models but only 20 combinations could a priori reproduce the three behavioral results: increased *Te* with incentive, decreased *Tr* with incentive, and decreased *Te* with difficulty (Fig. S1). These 20 models were inverted using a variational Bayes approach under the Laplace approximation (38, 39), implemented in a toolbox by Jean Daunizeau (available at <http://sites.google.com/site/jeandaunizeauswebsite/>). This algorithm not only inverts nonlinear models but also estimates their evidence, which represents a trade-off between accuracy (goodness of fit) and complexity (degrees of freedom) (40). The log-evidences estimated for each subject and model were submitted to a group-level random-effect analysis (41) using SPM8 (Statistical Parametric Mapping, Wellcome Department of Imaging Neuroscience, London, UK).

**fMRI Data Analysis.** All GLMs included realignment parameters as covariates of no interest to correct for movement artifacts. Regressors of interest were specified at the 125-ms scale, and convolved with the canonical hemodynamic response function (HRF) and its first temporal derivative. The GLM included two categorical regressors: one modeling coin display onset with a delta function and one modeling the entire session with a boxcar function. There were also two parametric regressors: one modulating coin display by its value (10, 20, or 50 cents) and one modeling cost evidence variation over the entire session. Cost evidence was continuously modeled over effort and rest periods defined from the behavior, with linear increases and decreases between constant minimum and maximum, and then z-scored. Thus, the parametric cost evidence regressor was ramping up and down, between positive and negative values, during task trial and put to zero between trials (Fig. 3A).

Regression coefficients were estimated at the subject level using the standard restricted minimum-likelihood (ReML) estimation implemented in SPM8. Individual linear contrasts of HRF regressors were then taken to a group-level random-effect analysis using one-sample *t* tests. Statistical thresholds corresponding to correction for multiple comparisons over the entire brain were determined using the randomization ( $n = 10,000$  permutations) technique implemented in FSL (Centre for Functional MRI of the

Brain, Oxford, United Kingdom). Two thresholds were used: a voxel-wise family-wise error (FWE) rate of  $P < 0.05$  and a cluster-wise FWE rate of  $P < 0.05$  (defined for voxels surviving  $P < 0.001$ , uncorrected). The cost evidence regressor mapped onto three regions (bilateral posterior insula and ventromedial thalamus) that survived both voxel-wise and cluster-wise corrections. These three clusters formed at  $P < 0.001$ , uncorrected, were grouped together to form a single region of interest, which was used for all subsequent analyses.

Three other GLMs were built that differed only on the cost evidence regressor, which now had an amplitude modulated by incentives in the same proportion as the model optimized on behavior. This modulation could in principle rely on the upper bound only (GLM2), or the lower bound only (GLM3), or be shared between both bounds (GLM4). Variational Bayes estimation procedure implemented in SPM8 was used to estimate for each subject the GLM log-evidences, which were summed over all voxels included in the ROI. Individual log-evidences were then submitted to a group-level random-effect BMS to identify the most probable cost evidence model given the ROI activity.

**MEG Data Analysis.** For the PCA, MEG data were epoched into rest and effort periods, resampled to 1,250 points, and averaged over conditions. A PCA was computed on the grand average (over subjects) to estimate the  $R^2$  statistic of each component for rest and effort periods, separately. PCAs were also computed in each subject to reconstruct the sources of the first component with SPM8, using a minimum-norm algorithm that could include or not

include the fMRI-based ROI as priors. The best reconstruction method (with or without priors) was determined using group-level BMS (42).

For the slope and amplitude analyses, three regressions were performed on the scalp raw time series. A linear trend was first fitted separately on rest and effort epochs to estimate accumulation and dissipation slopes, respectively. Then a V shape was fitted on contiguous rest-effort epochs to estimate the signal amplitude. Sources of the mean slopes and amplitudes were reconstructed subject-wise using the same procedure as for PCA. The estimated scalp-to-source reconstruction matrices were then used to invert each epoch. Activity in the source space was rectified (absolute value), log-normal transformed, and averaged within the ROI. The resulting activity was analyzed using a linear model including three regressors: incentive (10, 20, 50) and difficulty (70, 80, 90) levels plus a constant. The significance of regression coefficients was estimated at the group level using a two-tailed  $t$  test.

**ACKNOWLEDGMENTS.** We thank A. Ducorps, C. Gitton, L. Hugueville, and V. Marchand-Pauvert for critical technical support with electrophysiological recordings; the Centre de NeuroImagerie de Recherche team for their help with fMRI recordings; T. Andriillon and R. Le Bouc for their help in behavioral pilots and data acquisition; and C. Tallon-Baudry for her comments on the MEG analysis. This study was funded by a Starting Grant from the European Research Council. F.M. received a fellowship from the French Ministère de l'Éducation Nationale. C.S. was supported by the Institut National de la Santé et de la Recherche Médicale and L.R. by the Emergence Program from the Ville de Paris.

- Walton ME, Kennerley SW, Bannerman DM, Phillips PEM, Rushworth MFS (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Netw* 19(8):1302–1314.
- Salamone JD, Correa M, Farrar A, Mingote SM (2007) Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology (Berl)* 191(3):461–482.
- Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9(7):545–556.
- Boksem MAS, Tops M (2008) Mental fatigue: Costs and benefits. *Brain Res Brain Res Rev* 59(1):125–139.
- James W (1907) The energies of men. *Science* 25(635):321–332.
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci* 9(6):467–479.
- Pessiglione M, et al. (2007) How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science* 316(5826):904–906.
- Schmidt L, Lebreton M, Cléry-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural mechanisms underlying motivation of mental versus physical effort. *PLoS Biol* 10(2): e1001266.
- Peyron R, Laurent B, García-Larrea L (2000) Functional imaging of brain responses to pain. A review and meta-analysis (2000). *Neurophysiol Clin* 30(5):263–288.
- Friebel U, Eickhoff SB, Lotze M (2011) Coordinate-based meta-analysis of experimentally induced and chronic persistent neuropathic pain. *Neuroimage* 58(4): 1070–1080.
- Krajibich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13(10):1292–1298.
- Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci USA* 107(50):21767–21772.
- Hunt LT, et al. (2012) Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* 15(3):470–476, S1–S3.
- Nybo L, Secher NH (2004) Cerebral perturbations provoked by prolonged exercise. *Prog Neurobiol* 72(4):223–261.
- Subudhi AW, Miramon BR, Granger ME, Roach RC (2009) Frontal and motor cortex oxygenation during maximal exercise in normoxia and hypoxia. *J Appl Physiol* 106(4): 1153–1158.
- Amann M, Dempsey JA (2008) Locomotor muscle fatigue modifies central motor drive in healthy humans and imposes a limitation to exercise performance. *J Physiol* 586(1): 161–173.
- Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453 (7197):869–878.
- Lee JH, et al. (2010) Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature* 465(7299):788–792.
- Gutschalk A, Hämäläinen MS, Melcher JR (2010) BOLD responses in human auditory cortex are more closely related to transient MEG responses than to sustained ones. *J Neurophysiol* 103(4):2015–2026.
- Vartiainen J, Liljeström M, Koskinen M, Renvall H, Salmelin R (2011) Functional magnetic resonance imaging blood oxygenation level-dependent signal and magnetoencephalography evoked responses yield different neural functionality in reading. *J Neurosci* 31(3):1048–1058.
- Rosa MJ, Daunizeau J, Friston KJ (2010) EEG-fMRI integration: A critical review of biophysical modeling and data analysis approaches. *J Integr Neurosci* 9(4):453–476.
- Eickhoff SB, et al. (2010) Anatomical and functional connectivity of cytoarchitectonic areas within the human parietal operculum. *J Neurosci* 30(18):6409–6421.
- Mohr C, Binkofski F, Erdmann C, Büchel C, Helmchen C (2005) The anterior cingulate cortex contains distinct areas dissociating external from self-administered painful stimulation: A parametric fMRI study. *Pain* 114(3):347–357.
- Beckmann M, Johansen-Berg H, Rushworth MFS (2009) Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci* 29(4):1175–1190.
- Mazzola L, Isnard J, Peyron R, Mauguière F (2012) Stimulation of the human cortex and the experience of pain: Wilder Penfield's observations revisited. *Brain* 135(Pt 2): 631–640.
- Craig ADB (2009) How do you feel—now? The anterior insula and human awareness. *Nat Rev Neurosci* 10(1):59–70.
- Nagvi NH, Bechara A (2009) The hidden island of addiction: The insula. *Trends Neurosci* 32(1):56–67.
- Jones CL, Ward J, Critchley HD (2010) The neuropsychological impact of insular cortex lesions. *J Neurol Neurosurg Psychiatry* 81(6):611–618.
- Deecke L, Scheid P, Kornhuber HH (1969) Distribution of readiness potential, pre-motor positivity, and motor potential of the human cerebral cortex preceding voluntary finger movements. *Exp Brain Res* 7(2):158–168.
- Colebatch JG (2007) Bereitschaftspotential and movement-related potentials: Origin, significance, and application in disorders of human movement. *Mov Disord* 22(5): 601–610.
- Croxson PL, Walton ME, O'Reilly JX, Behrens TEJ, Rushworth MFS (2009) Effort-based cost-benefit valuation and the human brain. *J Neurosci* 29(14):4531–4541.
- Kolling N, Behrens TEJ, Mars RB, Rushworth MFS (2012) Neural mechanisms of foraging. *Science* 336(6077):95–98.
- Wager TD, et al. (2004) Placebo-induced changes in FMRI in the anticipation and experience of pain. *Science* 303(5661):1162–1167.
- Tucker R (2009) The anticipatory regulation of performance: The physiological basis for pacing strategies and the development of a perception-based model for exercise performance. *Br J Sports Med* 43(6):392–400.
- Noakes TD (2011) Time to move beyond a brainless exercise physiology: The evidence for complex regulation of human exercise performance. *Appl Physiol Nutr Metab* 36 (1):23–35.
- Gandevia SC (2001) Spinal and supraspinal factors in human muscle fatigue. *Physiol Rev* 81(4):1725–1789.
- Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007) Variational free energy and the Laplace approximation. *Neuroimage* 34(1):220–234.
- Daunizeau J, Friston KJ, Kiebel SJ (2009) Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D* 238(21): 2089–2118.
- Robert CP (2001) *The Bayesian Choice: From Decision Theoretic Foundations to Computational Implementation* (Springer, New York), 2nd Ed.
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46(4):1004–1017.
- Daunizeau J, et al. (2005) Assessing the relevance of fMRI-based prior in the EEG inverse problem: A Bayesian model comparison approach. *IEEE Trans Signal Process* 53(9):3461–3472.

# Supporting Information

Meyniel et al. 10.1073/pnas.1211925110

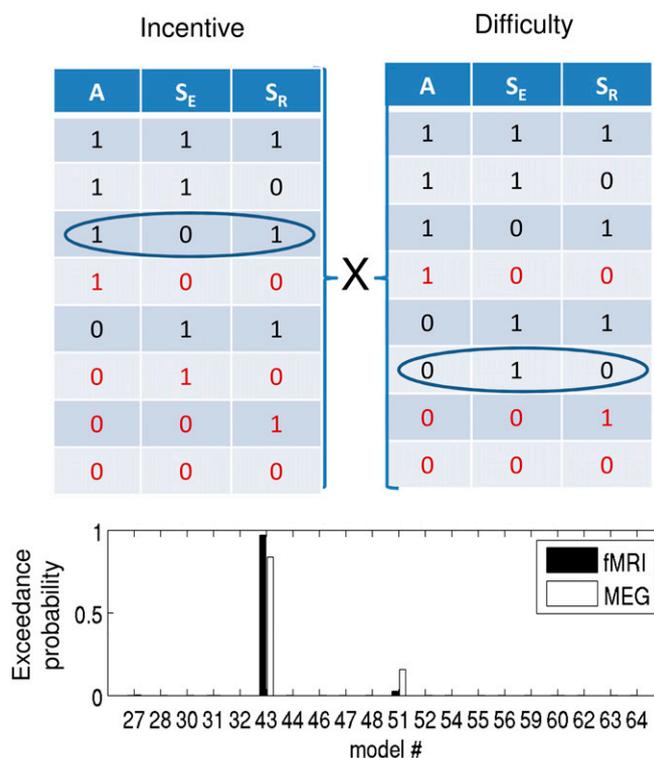
## SI Materials and Methods

**Experimental Setting.** Before scanning, participants were given written instructions to the task, which were repeated step by step orally. Subsequently, they were escorted inside the MRI or magnetoencephalography (MEG) scanning room and invited to find a comfortable body position that they could keep throughout the experiment. The only change was passing the power grip from one hand to the other between sessions. We used homemade power grips composed of two plastic cylinders compressing an air tube when squeezed. The tube led to the control room, where it was connected to a transducer converting air pressure into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to the force exerted. The signal was fed to the stimuli presentation PC via a signal conditioner (CED 1401; Cambridge Electronic Design) and read by a MATLAB program (MathWorks). Stimuli presentation was also programmed with MATLAB using Cogent 2000 (Wellcome Department of Imaging Neuroscience, London, UK).

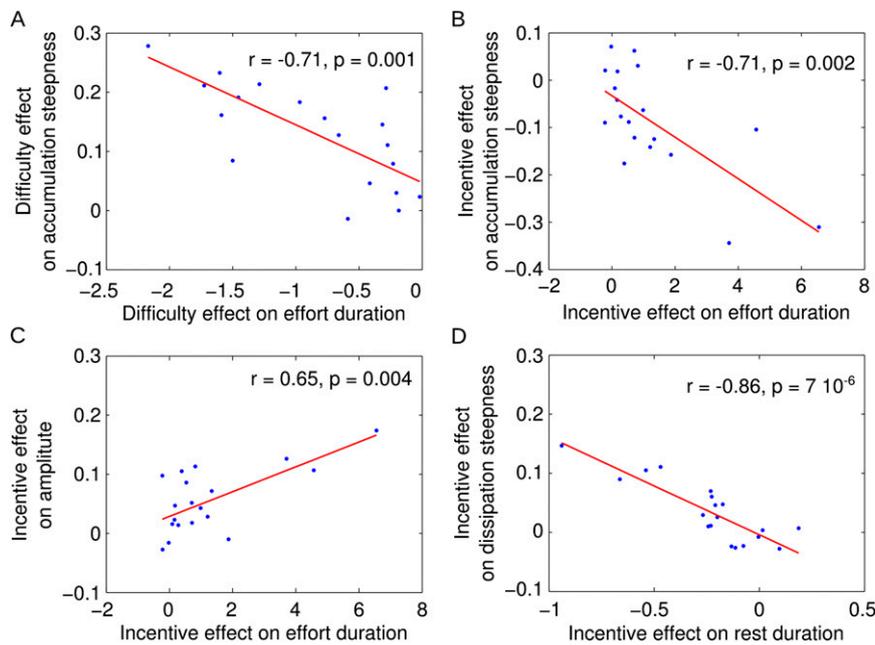
**fMRI Data Acquisition.** Subject's head was constrained using foam and sand packs to limit movements. Functional echo-planar images (EPIs) were acquired with a T2\*-weighted contrast on a 3-T scanner (Siemens Trio). Interleaved 2-mm slices separated by a 1.5-mm gap and oriented along a 30° tilted plane were acquired to cover the whole brain with a repetition time of 2 s. The first five scans were discarded to allow for equilibration effects. All

preprocessing steps were performed using SPM8. Structural T1-weighted images were also acquired, coregistered with the mean EPI, segmented, and normalized to SPM standard Montreal Neurological Institute (MNI) T1 template. Normalized T1-images were averaged between subjects to localize group-level functional activations. EPIs were spatially realigned and normalized (using the same transformation as for structural images), and spatially smoothed with a 8-mm full-width at half-maximum Gaussian kernel.

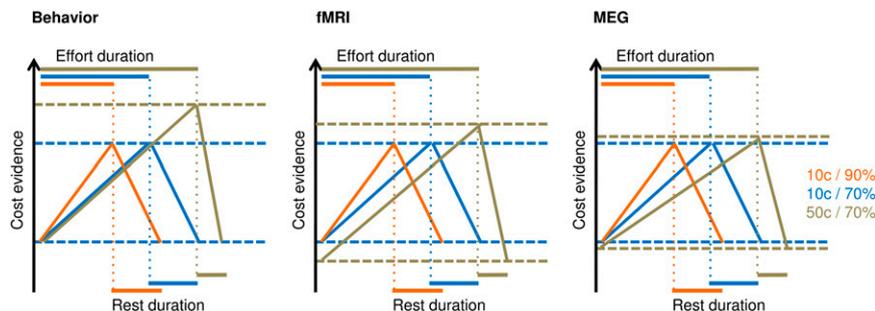
**MEG Data Acquisition.** A whole-head MEG system comprising 151 axial gradiometers (CTF Systems) was used to sample brain activity at 1,250 Hz with online low-pass filter of 300 Hz. Head position was determined using marker coils at fiducial points (nasion, left and right ears). Ocular artifacts were marked manually and removed using the Gratton method with DataHandler (Cogimage, Centre de Recherche de l'Institut du Cerveau et de la Moelle Épineuse, Paris, France). Data were imported in MATLAB and displayed using FieldTrip (Donders Institute, Nijmegen University, Nijmegen, The Netherlands). MEG signal was low-pass filtered offline at 30 Hz. Effort onsets and offsets were detected manually based on the electromyogram. A template mesh (8,196 tessels) and individual fiducials were used in SPM8 to compute a single shell head model and a lead field matrix per subject and session.



**Fig. S1.** Model comparison performed on behavioral data. The table illustrates all possible models obtained by varying the modulation of amplitude and slopes ( $A$ ,  $S_E$ , and  $S_R$ ) by monetary incentive (left columns) and/or task difficulty (right columns). Allowed (versus excluded) modulations are noted 1 (versus 0). “Allowed” means that there is a term for this modulation in the model. The red models are the combinations discarded a priori because they cannot reproduce the behavioral results. For instance, a model in which neither  $A$  or  $S_E$  has an incentive term cannot reproduce the modulation by incentives and therefore must be discarded. The surrounded combination corresponds to the model with the highest exceedance probability, calculated using Bayesian selection separately for fMRI and MEG subjects.



**Fig. S2.** Across-subject correlations between behavioral and computational effects of experimental factors (monetary incentive and task difficulty). Behavioral effects are as follows: shortened effort duration with higher task difficulty, prolonged effort duration with higher monetary incentive, and shortened rest duration with higher monetary incentive. Correlations were searched with all four computational effects that were found to explain these behavioral effects in MEG recordings. Computational effect refers to modulation of one parameter (accumulation steepness, dissipation steepness, or amplitude between bounds) of the signal reconstructed in our region of interest. Note that we use the term “steepness” and not “slope” because this power signal has no sign. The four computational effects are steeper accumulation with higher difficulty (A), slower accumulation with higher incentives (B), larger amplitude with higher incentives (C), and steeper dissipation with higher incentives (D).  $R$  values are Pearson  $\rho$  correlation coefficients;  $P$  values indicate the significance of robust-fit regressions (which underweight potential outliers).



**Fig. S3.** Model refinement with fMRI and MEG findings. The diagrams illustrate how the experimental factors (monetary incentive and task difficulty) affect the accumulation of cost evidence. Compared with the blue line, the brown line shows the effect of increasing the incentive and the orange line the effect of increasing the difficulty. In the model that best explained the behavior (Left), the amplitude and dissipation slope are only impacted by incentives, whereas the accumulation slope is only impacted by difficulty. The fMRI results (Center) revealed that both the lower and upper bounds of the cost evidence signal encoded in brain activity are modulated by incentives. The MEG results (Right) revealed that the accumulation slope of this neural cost evidence is additionally impacted by incentives. Note that the three models produce the same pattern of effort and rest durations across conditions. In the final rightmost model, increasing the incentive (*i*) augments effort duration by inflating the amplitude and lowering the accumulation slope and (*ii*) reduces rest duration by accentuating the dissipation slope. In contrast, increasing the difficulty only shortens effort duration by enhancing the accumulation slope, without affecting the amplitude or dissipation slope.

**Table S1. Brain regions parametrically modulated by the cost evidence signal in fMRI data analysis**

Anatomical label	Peak <i>t</i>	Peak uncorrected <i>P</i>	Peak FWE <i>P</i>	Cluster FWE <i>P</i>	No. of voxels	Peak coordinates
Left posterior insula	5.825	0.000	0.011	0.006	226	[-40 -22 16]
Ventromedial thalamus	5.693	0.000	0.013	0.018	69	[2 -10 -4]
Right posterior insula	5.469	0.000	0.018	0.006	254	[42 -16 10]
Hypothalamus	4.819	0.000	0.053	0.031	36	[0 0 -20]
Fusiform gyrus	4.774	0.000	0.057	0.018	68	[22 -52 -10]
Cerebellum	4.422	0.000	0.100	0.041	24	[4 -46 -40]
Fusiform gyrus	4.234	0.001	0.132	0.038	27	[18 -70 -6]
Fusiform gyrus	4.227	0.001	0.133	0.035	30	[-10 -68 -4]

All clusters are listed that were observed using a voxel-wise threshold of  $P < 0.001$ , uncorrected, and a cluster-wise threshold of  $P < 0.05$ , family-wise error (FWE) corrected. The [x y z] peak coordinates in millimeters refer to the Montreal Neurological Institute (MNI) space.

**Table S2. Brain regions parametrically modulated by the cost evidence signal when including the motor output (force) as a covariate in the GLM**

Anatomical label	Peak <i>t</i>	Peak uncorrected <i>P</i>	Peak FWE <i>P</i>	Cluster FWE <i>P</i>	No. of voxels	Peak coordinates
Ventromedial thalamus	5.490	0.000	0.017	0.022	88	[2 -10 -4]
Right posterior insula	5.463	0.000	0.018	0.007	367	[42 -16 8]
Left posterior insula	5.410	0.000	0.019	0.008	286	[-42 -22 18]
Fusiform gyrus	4.661	0.000	0.064	0.005	616	[22 -52 -10]
Hypothalamus	4.578	0.000	0.073	0.033	59	[0 0 -20]
Cerebellum	4.540	0.000	0.078	0.032	61	[4 -48 -40]

All clusters are listed that were observed using a voxel-wise threshold of  $P < 0.001$ , uncorrected, and a cluster-wise threshold of  $P < 0.05$ , family-wise error (FWE) corrected. The [x y z] peak coordinates in millimeters refer to the Montreal Neurological Institute (MNI) space. The contrast tested is the cost evidence signal.



## Title Page

### Title

Better get back to work: a role for motor beta de-synchronization in incentive motivation

### Abbreviated title

Incentive motivation and motor oscillations

### Authors & Affiliations

Florent Meyniel<sup>a, c, d</sup> & Mathias Pessiglione<sup>a, b, c, d</sup>

<sup>a</sup> Motivation, Brain & Behavior (MBB) lab

<sup>b</sup> Centre de NeuroImagerie de Recherche (CENIR)

<sup>c</sup> Brain & Spine Institute (ICM – Institut du Cerveau et de la Moelle épinière), Hôpital de la Pitié-Salpêtrière, 47 Bd de l'Hôpital, Paris, 75013 France

<sup>d</sup> CNRS UMR 7225, Inserm UMR 975, Université Pierre et Marie Curie (UPMC – Paris 6), 4 Place Jussieu, Paris, 75005 France

### Corresponding author

Mathias Pessiglione, Bâtiment ICM, Hôpital de la Pitié-Salpêtrière, 47 Bd de l'Hôpital, Paris, 75013 France.

E-mail: mathias.pessiglione@gmail.com

Number of pages: 31

Number of figures: 5, tables: 0, other: none.

Number of words: Abstract (227), Introduction (495), Discussion (1465).

Acknowledgements: The study was funded by a Starting Grant from the European Research Council. FM benefited from a doctoral fellowship by the French Ministère de l'Éducation Nationale. We thank Catherine Tallon-Baudry for her valuable discussions on the analysis, the CENIR MEG team and Raphael Lebouc for help during recordings, and Véronique Marchand-Pauvert for her technical support with EMG recordings. The authors declare no competing financial interests.

## Abstract

Much research has been devoted to characterizing brain representations of reward and movement. However, the mechanisms allowing expected rewards to influence motor commands remain poorly understood. Unraveling such mechanisms is crucial to providing explanations of how behavior can be driven by goals, hence accounting for apathy cases in clinics. Here, we propose that the reduction of motor beta synchrony (MBS) prior to movement onset could participate in this incentive motivation process. To test this hypothesis, we recorded brain activity using magnetoencephalography (MEG), while participants were exerting physical effort in order to win monetary incentives. Knowing that the payoff was proportional to the time spent above a target force, subjects spontaneously took breaks when exhausted and resumed effort production when repleted. Behavioral data indicated that the rest periods were shorter when higher incentives were at stake. MEG data showed that the amplitude of MBS reduction correlated to both incentive level and rest duration. Moreover, the time of effort initiation could be predicted by MBS reduction measured at the beginning of rest periods. Incentive effects on MBS reduction and rest duration were also correlated across subjects. Finally, Bayesian comparison between possible causal models suggested that MBS reduction mediates the impact of incentive level on rest duration. We conclude that MBS reduction could represent a neural mechanism that speeds the initiation of effort production when the effort is more rewarded.

## Introduction

Typical explanations of behavior involve reward prospect: we engage in actions in order to attain more valuable states. The question of how expected value can induce behavioral activation - a process termed incentive motivation - has been extensively investigated in neuroscience (Berridge, 2004; Schultz, 2006). A central issue is how the value of a world state can be translated into a motor code. The basal ganglia system seems a likely candidate for mediating this incentive motivation process, as revealed by electrophysiology or functional MRI in healthy subjects (Schmidt et al., 2012; Tachibana and Hikosaka, 2012), and by the effects of focal lesions or degenerative disease in patients (Czernecki et al., 2002; Schmidt et al., 2008).

The case of Parkinson's disease (PD) is particularly enlightening. PD is primarily due to degeneration of dopaminergic neurons and manifests as a set of motor symptoms (akinesia, rigidity, tremor). However, the reduction of movement (or hypokinesia) could easily be reframed in terms of dysfunctional motivation, i.e., as a difficulty in activating motor plans leading to better states (Mazzoni et al., 2007; Baraduc et al., 2013). This idea accords well with an abundant literature that has implicated dopamine in incentive motivation (Berridge, 2004; Salamone and Correa, 2012). Interestingly, one electrophysiological hallmark of PD is the high level of synchronous oscillations in the beta band (Schnitzler and Gross, 2005; Brown, 2006; Uhlhaas and Singer, 2006). Dopamine replacement therapy, as well as deep brain stimulation, was shown to simultaneously reduce beta synchrony and alleviate hypokinesia (Brown et al., 2001; Kühn et al., 2008). More precisely, delay in voluntary movement initiation was linked to insufficient reduction of beta synchrony in the subthalamic nucleus (Kühn et al., 2004; Williams et al., 2005). Furthermore, simultaneous deep and surface recordings showed that beta oscillatory activities were coherent across basal ganglia nuclei and motor cortical areas (Klostermann et al., 2007; Litvak et al., 2010). Taken together, these findings suggest that reducing motor beta synchronization (MBS) represent a neural mechanism through which expected rewards may facilitate action initiation.

In healthy people, there is indeed a progressive reduction of MBS (typically in the 13-30Hz range and centered on the precentral cortex), preceding movement initiation (Jasper and Penfield, 1949; Feige et al., 1996; van Wijk et al., 2012). It is thought to play a "gating role", meaning that high beta synchrony maintains the motor status quo, whereas low beta synchrony allows for a motor change (Engel and Fries, 2010; Jenkinson and Brown, 2011). Direct evidence of this idea is the finding that driving motor cortical activity in the beta rhythm slows motor performance (Pogosyan et al.,

2009; Joundi et al., 2012). Therefore, we reasoned that in incentive motivation paradigms, the reward magnitude should modulate MBS, which, in turn, should have an impact on movement initiation. To test this hypothesis, we analyzed MEG activity previously recorded in healthy participants while they were trying to minimize the duration of breaks during force production, in order to maximize their payoff (Meyniel et al., 2013).

## Materials and methods

The experimental procedures were designed for other purposes: tracking accumulation and dissipation of cost signals while subjects alternate effort and rest periods. Nevertheless, they seem well suited to address our question about the role of MBS in incentive motivation. The task involved subjects squeezing a handgrip in order to accumulate as much money as possible. The payoff was calculated as the monetary incentive multiplied by the time spent above a target force level (which indexed task difficulty). Both the incentive and difficulty levels were varied across trials such that we could assess their effects on effort allocation. Incentive levels were sufficient for subjects making the effort and reach the target, but difficulty levels were too demanding for subjects to sustain their effort throughout trials, which lasted 30 seconds. Instead, they freely alternated effort and rest periods within trials (see Figure 1B for an illustration). Note that any time spent resting corresponded to potential money that was not earned. Thus, subjects had no instructions about when to start force production, but were motivated with various incentives to hurry up. We reproduce here the methodological aspects that are relevant to the present question; further details can be found in our previous publication (Meyniel et al., 2013).

### Subjects

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. Subjects ( $n=19$ , 8 males,  $24.9\pm 0.7$  years) were recruited within an academic database via email, and gave informed consent prior to participating in the study. They were right-handed (according to self-report), between 20 and 30 years old, free from magnetic artifacts, and with normal vision and no history of neurological or psychiatric disease. They believed that the money won while performing the task would be their remuneration for participating, but their payoff was eventually rounded up to a fixed amount (100€).

### Experimental setting

Prior to scanning, participants were given written instructions on the task, which were repeated step-by-step orally. Subsequently, they were escorted inside the MEG scanning room and encouraged to find a comfortable body position that they could keep throughout the experiment. The only change in position was passing the power grip from one hand to the other between sessions. We used homemade power grips composed of two plastic cylinders compressing an air tube when squeezed. The tube led to the control room, where it was connected to a transducer converting air pressure

into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to the force exerted. The signal was fed to a stimuli presentation PC via a signal conditioner (CED 1401, Cambridge electronic design, UK) and read by a Matlab program (The MathWorks Inc., USA). Stimuli presentation was also programmed with Matlab using Cogent 2000 (Wellcome Department of Imaging Neuroscience, London, UK).

### **Behavioral task**

Before starting task performance, we measured the maximal force for each hand, following published guidelines (Gandevia, 2001). Participants were verbally encouraged to squeeze continuously as hard as they could, until a line growing in proportion to their force reached a target displayed on a computer screen. Maximal force was defined as the average, over the last half of the squeezing period, of data points exceeding the median force. Then subjects were provided real-time feedback about the force produced on the handgrip, which appeared as a fluid level moving up and down a thermometer (the maximal force being indicated as a horizontal bar at the top). Subjects were asked to try outreaching the bar and state whether it truly corresponded to their maximal force. If not, the calibration procedure was repeated.

Task sessions included 9 trials corresponding to the 9 cells of the factorial design (3 incentive x 3 difficulty conditions), which were presented in random order. Subjects performed 8 sessions in total, switching hands as instructed between sessions to avoid muscle exhaustion. After baseline measurement of the pressure at rest, each trial started by revealing the monetary incentive with a coin image (10, 20 or 50 cents) displayed for 1s. Then subjects had 30s to win as much money as possible. They knew that the payoff was proportional to both the incentive and the time spent above the target bar, which was always placed at the same height in the thermometer. The force needed to reach the bar (70, 80 or 90% of subject's maximal force), i.e. trial difficulty, was not indicated to subjects. Subjects only knew that task difficulty would vary across trials. They were provided with online feedback on both the exerted force (with a fluid level moving up and down a thermometer) and the trial-wise cumulative payoff (with a counter displayed above the thermometer). The counter was only started when fluid level was above the target bar, with a rate proportional to the current incentive. The fluid had the same luminance as the background to avoid confounding force level with basic visual features. Each trial ended with a 2s display of the session-wise cumulative payoff.

### **Behavioral data analysis**

Effort onsets and offsets were manually identified offline based on the electromyograms and force level time series. The first rest period started with coin presentation and ended with the first effort onset. The last period, which was interrupted by trial ending, was not analyzed. Effort initial speed was estimated as the mean derivative of force level within the 500ms following effort onset. Effort and rest period durations were separately analyzed using a repeated-measure ANOVA (R software, John Fox's CAR library), with incentive and difficulty as factors of interest. The p-values reported for these ANOVA integrate the conservative Greenhouse-Geisser correction.

### **MEG data acquisition**

A whole-head MEG system with 151 axial gradiometers (CTF Systems) was used to sample brain activity at 1250 Hz with an online low-pass filter of 300 Hz. Two bad sensors (MLT42 & MRT32) were excluded because of high noise levels. Head position was determined using marker coils at fiducial points (nasion, left and right ears) before each session. The first session served as a reference to control that head displacement in the 7 remaining sessions never exceeded 5 mm. The first two sessions were excluded from the analyses in two participants, due to excessive change in head position.

Electromyograms were recorded simultaneously with two pairs of disposable surface electrodes on each hand, which were placed to target the flexor digitorum superficialis (on the forearm) and the first dorsal interossei (between thumb and index). The two EMG traces and the force time series were recorded along with MEG data (hence sampled at 1250 Hz), and were conjointly used to detect effort and rest onsets by visual inspection. For this offline manual detection, the EMG traces were high pass filtered ( $> 100$  Hz) and the force levels low pass filtered ( $< 10$  Hz).

### **MEG spectral decomposition**

Data was imported into Matlab and analyzed using Fieldtrip toolbox (<http://fieldtrip.fcdonders.nl>, (Oostenveld et al., 2011)). For each session (lasting around 320s), the whole data set was decomposed into power over time and frequency, using a product with a set of Morlet wavelets after fast-Fourier transform. A product in the frequency space is equivalent, but computationally faster, to a convolution in the time space. The Morlet wavelets trade temporal against spectral resolution ( $\sigma_t$  against  $\sigma_f$ ), such that  $\sigma_f \sigma_t = 1/(2\pi)$ , and scale this tradeoff according to the frequency  $f$ , such that  $f/\sigma_f$  is constant. This ensures that finer temporal resolution is achieved for higher frequencies at the expense of a lower spectral resolution. We set the  $f/\sigma_f$  ratio to 7, which is standard for the frequency range investigated here. For each unit frequency between 5 and 40 Hz, the product between the

wavelets and the data was computed for every 50-ms step. This is well below the original sampling rate (1250Hz) but sufficient considering the time scale of the process investigated (seconds). The length of each wavelet used for the computation was 3 times its temporal resolution ( $\sigma_t$ ). These power data were used for every analysis, except the source reconstruction.

### **MEG source reconstruction**

We used source reconstruction to check that the desynchronization of beta oscillations during rest was arising from the motor cortex. Templates of brain anatomy (single\_subj\_T1.nii) and meshes of cortical surface and head envelope (cortex\_5124.surf.gii) were taken from SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>). For every subject, sensor positions were co-registered in the MNI space using the mean fiducial positions over sessions as landmarks. The volume conductor model that served to compute dipole orientation was based on the template tessellation of cortical surface. The normalized leadfield, i.e. the forward solution from the distribution of dipoles over the grid space to the set of scalp sensors, was estimated using a realistic single shell head model. This corresponds to the 'Nolte' method in the Fieldtrip software. The grid size was 10x10x10mm, which covered the entire brain. Time series were epoched into pairs of temporal windows, first with the 2s of rest preceding effort onset and second with the 2s of baseline preceding the corresponding trial. All epochs were analyzed together in the 13-30 Hz frequency domain, using multi-tappers to compute the cross-spectrum matrix between sensor pairs (in which the diagonal corresponds to the power spectrum at each sensor). A spatial filter was jointly computed for rest and baseline data, using DICS beamformer (Dynamic Imaging of Coherent Sources (Gross et al., 2001)), without regularization of the solution. Then, we projected the signal of each epoch separately through the filter to estimate the power in the source space, over the 13-30 Hz range. A common filter allows contrasting rest and baseline power levels, hence computing the percentage of signal change. The group mean was interpolated onto the anatomical template for display.

### **MEG statistical analysis**

Multiple regressions of power level against the various factors of interest (Figure 3A) were estimated using Matlab statistical toolbox. The significance of all regression coefficients was estimated at the group level using a non-parametric procedure, which was repeated for every frequency tested. To estimate uncorrected p-values, the null distribution of t-values was estimated by flipping the sign of regression coefficients over participants. To estimate family-wise error (FWE) corrected p-values, the null distribution of the maximal t-value was estimated in a similar fashion (Nichols and Hayasaka, 2003). In both cases, 200 000 distinct sign changes were used among the  $19^2 = 524\ 288$  possible

changes. This correction accounts for multiple testing over frequencies, but not for regressing against multiple factors. We did not correct for the number of regressors in the linear models, since they were tested against each other.

For between-subject correlations (Figure 3B), uncorrected p-values were calculated from Spearman correlations, and FWE-corrected p-values were derived from Holm's step-down adaptive method, which strongly controls FWE without making any assumption (Nichols and Hayasaka, 2003).

Cluster statistics in Figure 4B were calculated using a cluster-mass permutation scheme described in (Maris and Oostenveld, 2007). First, to determine uncorrected statistical threshold ( $T_{\text{thd}}$ ) corresponding to  $p=0.05$ , we used the same randomization procedure as described above for multiple regression analyses. These non-parametric thresholds (incentive:  $T_{\text{thd}}=2.092$ ; rest duration:  $T_{\text{thd}}=2.088$ ) were very close to their parametric counterpart ( $t_{0.975,18}=2.101$ ). Second, cluster-mass FWE-corrected p-values were estimated from the null distribution of maximal cluster mass (sum of cluster t-values) formed at  $T_{\text{thd}}$ . This distribution was approximated using 200 000 distinct randomizations (again among 524 288) of regression coefficient signs over participants. All p-values reported throughout this study correspond to bilateral tests.

### Bayesian model selection

For each model depicted in Figure 5B, rest duration and z-power were separately modeled as linear combinations of the variables supposed to impact them (through the arrows). We fitted the models by adjusting the beta coefficients so as to minimize the residuals (denoted  $\varepsilon$ ) in the following equations (with  $I$  being the incentive level,  $F$  the first eigenvector of MBS and  $D$  the rest duration):

$$\begin{cases} F = \beta_1 I + \varepsilon \\ D = \beta_2 F + \varepsilon \end{cases} \text{ (model 1)}$$

$$\begin{cases} F = \beta_1 I + \varepsilon \\ D = \beta_2 F + \beta_3 I + \varepsilon \end{cases} \text{ (model 2)}$$

$$\begin{cases} F = \beta_1 I + \varepsilon \\ D = \beta_2 I + \varepsilon \end{cases} \text{ (model 3)}$$

All variables were z-scored beforehand, such that their offset and scaling could not bias model comparison. Models were estimated using a variational-Bayes approach under the Laplace approximation (Friston et al., 2007; Daunizeau et al., 2009), which was implemented in a toolbox by Jean Daunizeau (available at <http://sites.google.com/site/jeandaunizeauswebsite/>). This algorithm

estimates model evidence and parameters, using the free-energy approximation (Friston et al., 2007). This approximation of model evidence is computationally tractable and usually more accurate than the Akaike Information Criterion or the Bayesian Information Criterion (Penny, 2012). A model evidence is the probability of observing the data given this model. It corresponds to the marginal likelihood, i.e. the integral over the parameter space of the model likelihood weighted by the prior on its parameters. Therefore, model evidence increases with likelihood (accuracy of the fit) but is penalized by the dimension of the parameter space (complexity of the model). In other words, it quantifies the complexity versus accuracy trade-off, which is mandatory for model comparison (Stephan et al., 2009). Model selection was performed at the group level using a random effect: log-evidence values were passed through Gibbs sampling (provided in SPM8) to approximate the models posterior densities and hence the models exceedance probabilities. Given the data acquired across participants, the model exceedance probability quantifies the belief that this model is more frequently implemented in the general population than any other model of the tested set (Stephan et al., 2009).

## Results

### Factors affecting rest behavior

In our task, participants could freely allocate their effort production within the 30s of each trial (Figure 1A), knowing that the payoff was proportional to both the effort duration and the incentive level. All subjects spontaneously alternated effort and rest periods during the course of trials (see example on Figure 1B), suggesting that incentive levels were high enough to induce effort production, and difficulty levels high enough to impose breaks. Two repeated-measure ANOVA were performed to characterize the effect of the manipulated factors (incentive and difficulty levels) on the durations of rest and effort periods, separately (see Figure 1B). For higher incentives, participants prolonged effort duration ( $F_{2,36}=11.1$ ,  $p=2.7 \cdot 10^{-3}$ ) and shortened rest duration ( $F_{2,36}=10.5$ ,  $p=3.2 \cdot 10^{-4}$ ). These two effects contributed to increase the payoff, as they augmented the total time spent squeezing the grip when more money is at stake. Higher difficulty shortened effort duration ( $F_{2,36}=14.0$ ,  $p=4.3 \cdot 10^{-5}$ ) but did not significantly affect rest duration ( $F_{2,36}=1.0$ ,  $p=0.35$ ). In the following, we focus on rest duration, which was only affected by incentive level. The aim of the following MEG data analysis was to test whether incentive effect on rest duration was mediated by MBS reduction.

### Spatio-temporal characteristics of MBS

To be qualified as MBS, our signal had to exhibit three critical features: 1) a dip around movement initiation, 2) in a specific frequency band (13-30 Hz), 3) from a source located over the central sulcus, lateralized with respect to the hand used.

After spectral decomposition of oscillatory activity, the power was normalized by a z-score transformation. This 'z-power' was calculated at each frequency by subtracting the mean and dividing it by the standard deviation of the trial baseline, which was defined as the 2-s window preceding incentive display. To verify the scalp topography of MBS at effort initiation, we averaged z-power over time (within 1s centered on effort onset) and frequency (within a range of 13 to 30 Hz). Sensors showing the lowest synchrony level were indeed localized around central brain surface (Figure 2A, right). In order to analyze variations of MBS level, we selected as sensors of interest the 10% sensors with lowest beta synchrony level. The selection procedure resulted in a similar list of sensors when applied on left-hand, right-hand, or pooled data (as can be expected from

topographies in Figure 2B). In order to simplify subsequent analyses, we selected a single set of sensors across hands and subjects, based on pooled data. The time course of z-power, averaged across sensors of interest, showed the classical progressive reduction in the beta band, over the 2s preceding effort initiation, dipping just after effort onset (Figure 2A, left).

To confirm that beta synchrony reduction arose from motor regions, we reconstructed the sources of power in the 13-30 Hz range using a beamformer (see Methods). Contrasting rest to baseline revealed that two main sources underpinned beta power reduction in the left and right sensorimotor cortex (MNI peak coordinates: [30 -30 60] and [-50 -40 50]). When contrasting sessions using left and right hand, a significant asymmetry was observed in favor of the contralateral sensorimotor cortex (see Figure 2B). Note that the sources of desynchronization in the alpha band (10 – 13 Hz) were very close, with peaks at [40 -40 50] and [-60 -30 40]. The asymmetry between hemispheres could reflect the fact that slightly different motor regions are recruited when subjects (who were all right-handed in our study) move their dominant versus their non-dominant hand.

## Factors modulating MBS

The above results confirmed that our task elicits MBS reduction prior to effort, using classical fix-window analyses. Since we aimed at explaining rest duration in the following analyses, we now considered variable windows: the entire epochs between effort offset and onset. To identify which factors have an impact on MBS, we averaged, for each rest epoch, z-power over time points and sensors of interest, which were selected on the basis of independent criteria (see above). This mean z-power was fitted with a linear model including both manipulated and observed variables: incentive level, difficulty level, duration of the preceding effort, rest duration, initiation speed of the following effort, and duration of the following effort. The model also included factors of no interest that would capture fatigue or adaptation effects at different time scales: the ordinal position of the considered rest period within a trial, that of the trial within a session, and the session number. Regression coefficients were estimated independently for every frequency and their significance was calculated after family-wise error correction for multiple comparisons at the group level (see methods). Z-power was specifically reduced in the beta band by two factors: higher incentive levels and shorter rest durations (Figure 3A). Among the motor execution parameters (previous and following effort duration, rest duration, effort initiation speed, and effort difficulty), MBS only varied with rest duration. The size of this effect was significantly higher than the others (paired t-test over frequencies and variables, all  $p < 0.05$ , except for 'next effort' at 14-16Hz and 'Difficulty' at 39-40Hz

for which there was a weak trend  $p < 0.1$ ). Thus, the regression analysis supports the idea that MBS reduction may mediate incentive effects on a specific movement-related parameter: rest duration.

This idea also predicts that subjects who exhibit strong incentive effects on MBS reduction should as well exhibit strong incentive effects on rest duration. Note that at the subject level, MBS level and rest duration are not independent, which may bias the between-subject correlation between incentive effects on MBS reduction and rest duration. We therefore orthogonalized the two variables by regressing the linear effect of rest duration out of MBS level prior to estimating the parametric effect of incentive levels. As predicted, the between-subject correlation was positive and significant in the beta range (Figure 3B), surviving family-wise error correction for testing multiple frequencies (see Methods).

### **Comparison of causal models linking incentives, MBS and behavior**

In principle, the statistical dependencies between incentives, MBS, and behavior could result from rest duration mediating incentive effects on MBS. In other words, shorter duration would artifactually reduce measures of MBS level. Even though we do not see what mechanism could support this scenario, we intended to rule this out formally. As consequences cannot precede their causes, we simply tested whether MBS reduction would anticipate rest shortening. Z-power was averaged within a limited time window starting with effort offset and independent from rest duration (see Figure 4A). This early z-power was then regressed against the subsequent rest duration. The end of the time window was set for each subject at the 20<sup>th</sup> percentile of rest durations, to ensure both a sufficient amount of z-power samples (the minimum being 24 samples at 20Hz, corresponding to 1.2s) and a sufficient amount of rest periods included (the 20% rest periods shorter than the time window being excluded). We also included incentive level in the regression model, since it is correlated with rest duration. The regression was done for each subject and frequency and then tested for significance at the group level (see Figure 4B). Several clusters of frequencies formed at a  $p < 0.05$  uncorrected threshold survived FWE correction, revealing that z-power in the beta range was correlated positively with rest duration, and negatively with incentive level. Thus, the early MBS reduction was enhanced by incentive level and predicted the upcoming rest duration. Therefore, it seems unlikely that shorter rest duration may be the cause of MBS reduction.

However, the causal links between incentives, MBS, and behavior remained to be specified. Our hypothesis posits that incentives impact MBS reduction, which in turn controls effort initiation

(model #1). A more complicated possibility is that incentives also impact rest duration, independently from their effect on MBS (model #2). A last alternative is that incentive level is a common cause of both MBS reduction and rest duration, which would induce a spurious correlation between MBS level and rest duration (model #3). To reduce dimensionality of z-power, we ran a singular value decomposition (SVD). Results showed that z-power could reasonably be reduced to its first mode, which captured most of the variance (Figure 5A). In all subjects, the first eigenvector mirrored closely the MBS pattern over frequencies at effort initiation. For display purpose (Figure 5B), we oriented this first eigenvector such that its mean value in the beta range (13-30Hz) was positive in each subject. Indeed, the orientations of eigenvectors over frequencies and observations (i.e. left and right singular vectors) are arbitrary and depend on each other. Thus, using the first mode of z-power, we had one vector of observations for all three variables (incentive level, MBS reduction, and rest duration). Models were specified by linear dependencies between these variables (depicted by arrows in Figure 5C). Given the data, we estimated these models and took their respective evidence to perform Bayesian model selection with a random-effect analysis at the group-level (see Methods). Model #1, with an expected frequency of 0.85, obtained a high exceedance probability ( $x_p > 0.99$ ), which is the confidence that this model is more frequently implemented than the two other models in the general population. Thus, the model comparison provided evidence for MBS reduction mediating the effect of incentive level onto rest duration. A prediction of this mediation model is that the estimates of incentive effect onto rest duration should be reduced when MBS level is included as a second regressor in the analysis. Indeed, the regression coefficients assigned to incentive level when alone ( $-0.074 \pm 0.022$  s.e.m.) were greater (more negative) than when accompanied by MBS level ( $-0.058 \pm 0.020$  s.e.m.) with a significant difference ( $p=0.007$ , bilateral paired t-test).

## Discussion

In this study, participants were not instructed when to start exerting physical effort. Instead, they were motivated by monetary incentives to spend more time working. They spontaneously adjusted effort allocation to these incentives, trading off benefits against costs. In particular, they shortened breaks when work paid more. We found evidence that such an effect of incentive motivation was underpinned by a reduction of motor beta synchronization (MBS), relative to trial baseline. Indeed, MBS reduction was correlated across trials with both incentive level and rest duration. In addition, subjects who exhibited stronger incentive effects on rest duration also exhibited stronger incentive effects on MBS reduction. Finally, direct Bayesian model comparison suggested that the most likely interpretation of statistical dependencies between our three variables of interest is that incentive effect on rest duration was mediated by the amplitude of MBS reduction. In the following, we first discuss the modulation of rest duration by MBS reduction, and then the modulation of MBS reduction by incentives.

Effort onset could be predicted by MBS reduction measured in a fixed time window at the beginning of rest periods. This observation discards the possibility that rest duration per se may bias MBS measurements. On the contrary, it suggests that MBS reduction favors the initiation of effort production. This is in line with the general idea that a high MBS level represents an ‘idling rhythm’ maintaining the motor status quo and that decreasing MBS allows for a motor change (Engel and Fries, 2010). More precisely, the corticospinal pathway might be less excitable during high MBS states, preventing any motor program from triggering movement initiation. Interestingly, we observed that MBS reduction only impacted initiation time and no other effort-related parameter, such as speed or duration. This specificity echoes numerous reports that MBS reduction observed before the action onset is not linked to any movement parameter (van Wijk et al., 2012).

Growing evidence suggests that MBS reduction indeed plays on movement initiation. In PD patients, a higher MBS level in the STN was correlated across trials with longer initiation delay (Kühn et al., 2004), and with successful inhibition of the prepotent response in a Stroop task (Swann et al., 2009; Brittain et al., 2012). In healthy participants, faster finger tapping (with reduced intervals between movements) resulted in a lower MBS level (Toma et al., 2002), but prolonged movement duration (with constant intervals between movements) did not (Cassim et al., 2000), suggesting that MBS is specifically modulated by rest (not effort) duration. Interestingly, increasing response uncertainty (by augmenting the number of possible movements) also enhances both reaction time

and MBS level (Tzagarakis et al., 2010). Accordingly, in a visual detection task, the progressive reduction of MBS correlated with the gradual commitment to a motor response, which was distinct from the confidence in the perceptual decision (Donner et al., 2009; O'Connell et al., 2012).

The specific link of pre-effort MBS to initiation time should be contrasted to the case of readiness potential or field (RP/RF), which also manifests as a slow ramping signal that precedes voluntary movement (Pedersen et al., 1998; Praamstra et al., 1999; Leuthold and Jentzsch, 2002; Shibasaki and Hallett, 2006). Many movement-related factors affect the RP/RF such as the force load, the effector used and the movement complexity (Lang, 2003). Thus, while MBS reduction may reflect motor gating in general (Engel and Fries, 2010), with a main source in the contralateral motor cortex (Jurkiewicz et al., 2006; Donner et al., 2009; Tzagarakis et al., 2010), the RP/RF seems to reflect the preparation of a specific motor program (Shibasaki and Hallett, 2006), with main sources in the supplementary motor area in addition to the primary motor cortex (Ball et al., 1999; Cunnington et al., 2005). Such a clear-cut distinction between MBS-motor gating and RP/RF-motor preparation should however be tempered: some authors argue that the RP denotes the transition from intention to action (Lang, 2003), and others have recently proposed that the RP reflects the passive stochastic accumulation of a 'go' signal (Schurger et al., 2012).

Little is known about how these signals relate to reward processing, which has been overlooked by the EEG-MEG literature until recently. On the one hand, MBS reduction was characterized as a gating signal in the domain of motor control, the movement being directly instructed or related to a perceptual decision, without bearing any particular value for the subject. On the other hand, MBS was characterized in the domain of motor disorders as a pathological signal, which should be eliminated in order to alleviate symptoms such as hypokinesia in Parkinson's disease (PD). To our knowledge, the intuitive idea that MBS could represent a normal process adjusting motor behavior to subjective goals in healthy conditions has not been directly investigated. Here we provide evidence that MBS reduction may speed up effort initiation proportional to expected rewards in healthy subjects. Other neural mechanisms have been suggested for underlying such incentive motivation process. For instance, reward representation may influence motor output through cortico-cortical connections, implementing a top-down regulation of behavior (Locke and Braver, 2008; Kounieher et al., 2009). Another possibility is that the interaction between reward and motor circuits occurs within the basal ganglia, with the ventral parts boosting the dorsal parts (Knutson et al., 2008; Schmidt et al., 2012; Tachibana and Hikosaka, 2012). An alternative suggestion is that dopamine release facilitates the expression of motor programs, either at the cortical or sub-cortical level (Berridge, 2004; Robbins, 2007; Salamone and Correa, 2012).

These possibilities are not mutually exclusive and could be articulated with the phenomenon of MBS reduction. It is known that degeneration of dopaminergic neurons in animal models of PD, as well as in human patients, results in abnormally high beta oscillations that can be reduced by dopamine replacement medications (Schnitzler and Gross, 2005; Uhlhaas and Singer, 2006; Hammond et al., 2007). This is possibly because dopamine release in the striatum helps filter cortical input and desynchronize basal ganglia output signals (McIntyre and Hahn, 2010). Thus, one could speculate that reward prospects represented in limbic circuits may amplify dopamine release, which in turn may facilitate movement initiation by lowering beta oscillations in the motor circuits. This view would invite reconsidering the status of bradykinesia or akinesia as motor symptoms. They would instead represent dysfunction of motivational processes that are in between pure reward and motor representations. Several computational accounts of dopamine depletion support this interpretation, as they attributed slowness to a shift in the movement cost/benefit ratio, rather than to sub-optimal control of the movement spatiotemporal trajectory (Mazzoni et al., 2007; Niv et al., 2007; Baraduc et al., 2013). Yet the link between dopamine release and MBS reduction in healthy conditions remains to be established. We also note that abnormal oscillations in PD are not restricted to the beta band: excessive power was also found at lower frequency (in the alpha band), and deficient power at high frequency oscillations (in the gamma band), both anomalies being correlated to apathy scores (Airaksinen et al., 2012; Özkurt et al., 2011). This is line with the idea that beta and gamma motor oscillations have anti- and pro-kinetic properties, respectively (Brown et al., 2001; Schoffelen et al., 2005).

How the brain machinery for incentive motivation interacts with brain systems signaling effort costs remains poorly understood. In a previous study (Meyniel et al., 2013), we investigated a cost signal conveyed by proprioceptive regions (posterior insula) during the same task. In principle, this signal, which accumulated during effort and dissipated at rest, could be sufficient for making 'stop' and 'go' decisions. In particular, the decision to resume effort exertion was triggered when the insular cost signal reached down to a fixed lower bound. It is possible that insular dissipation is causally linked to MBS reduction or that they simultaneously favor effort initiation. Interestingly, the dissipation rate in the posterior insula increased with incentive level, as if a tonic reward signal was negatively integrated into the cost signal. An open question is whether insular dissipation and MBS reduction feed one another or integrate the same tonic incentive-related signal, possibly coming from a common input.

Thus, we conclude that MBS reduction may represent a neural process translating expected reward into motor activation. There are however some limitations that should be acknowledged. First, our conclusion is based on statistical dependencies between variables, which suggest - but not

prove - a causal pathway from incentive level to MBS reduction to effort initiation. Direct manipulation of MBS level, through dopaminergic medication or electrical stimulation, could provide more conclusive evidence for causality, by affecting subjects' sensitivity to incentives, and perhaps patients' apathetic symptoms. Second, the frequencies that were correlated to incentive level and rest duration were slightly different, even if they could all be labeled as 'beta'. It remains to be understood whether these differences in frequency are functionally significant for the incentive motivation process.

## References

- Airaksinen K, Butorina A, Pekkonen E, Nurminen J, Taulu S, Ahonen A, Schnitzler A, Mäkelä JP (2012) Somatomotor mu rhythm amplitude correlates with rigidity during deep brain stimulation in Parkinsonian patients. *Clin Neurophysiol* 123:2010–2017.
- Ball T, Schreiber A, Feige B, Wagner M, Lücking CH, Kristeva-Feige R (1999) The Role of Higher-Order Motor Areas in Voluntary Movement as Revealed by High-Resolution EEG and fMRI. *NeuroImage* 10:682–694.
- Baraduc P, Thobois S, Gan J, Broussolle E, Desmurget M (2013) A common optimization principle for motor execution in healthy subjects and parkinsonian patients. *J Neurosci* 33:665–677.
- Berridge KC (2004) Motivation concepts in behavioral neuroscience. *Physiology & Behavior* 81:179–209.
- Brittain J-S, Watkins KE, Joundi RA, Ray NJ, Holland P, Green AL, Aziz TZ, Jenkinson N (2012) A role for the subthalamic nucleus in response inhibition during conflict. *J Neurosci* 32:13396–13401.
- Brown P (2006) Bad oscillations in Parkinson's disease. *J Neural Transm Suppl*:27–30.
- Brown P, Oliviero A, Mazzone P, Insola A, Tonali P, Di Lazzaro V (2001) Dopamine dependency of oscillations between subthalamic nucleus and pallidum in Parkinson's disease. *J Neurosci* 21:1033–1038.
- Cassim F, Szurhaj W, Sediri H, Devos D, Bourriez J, Poirot I, Derambure P, Defebvre L, Guieu J (2000) Brief and sustained movements: differences in event-related (de)synchronization (ERD/ERS) patterns. *Clin Neurophysiol* 111:2032–2039.
- Cunnington R, Windischberger C, Moser E (2005) Premovement activity of the pre-supplementary motor area and the readiness for action: Studies of time-resolved event-related functional MRI. *Human Movement Science* 24:644–656.
- Czernecki V, Pillon B, Houeto JL, Pochon JB, Levy R, Dubois B (2002) Motivation, reward, and Parkinson's disease: influence of dopatherapy. *Neuropsychologia* 40:2257–2267.
- Daunizeau J, Friston KJ, Kiebel SJ (2009) Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D* 238:2089–2118.
- Donner TH, Siegel M, Fries P, Engel AK (2009) Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology* 19:1581–1585.
- Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? *Curr Opin Neurobiol* 20:156–165.

Feige B, Kristeva-Feige R, Rossi S, Pizzella V, Rossini P-M (1996) Neuromagnetic study of movement-related changes in rhythmic brain activity. *Brain Research* 734:252–260.

Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007) Variational free energy and the Laplace approximation. *NeuroImage* 34:220–234.

Gandevia SC (2001) Spinal and Supraspinal Factors in Human Muscle Fatigue. *Physiol Rev* 81:1725–1789.

Gross J, Kujala J, Hamalainen M, Timmermann L, Schnitzler A, Salmelin R (2001) Dynamic imaging of coherent sources: Studying neural interactions in the human brain. *Proc Natl Acad Sci USA* 98:694–699.

Hammond C, Bergman H, Brown P (2007) Pathological synchronization in Parkinson's disease: networks, models and treatments. *Trends in Neurosciences* 30:357–364.

Jasper H, Penfield W (1949) Electrocorticograms in man: Effect of voluntary movement upon the electrical activity of the precentral gyrus. *Arch F Psychiatr U Z Neur* 183:163–174.

Jenkinson N, Brown P (2011) New insights into the relationship between dopamine, beta oscillations and motor function. *Trends Neurosci* 34:611–618.

Joundi RA, Jenkinson N, Brittain J-S, Aziz TZ, Brown P (2012) Driving oscillatory activity in the human cortex enhances motor performance. *Curr Biol* 22:403–407.

Jurkiewicz MT, Gaetz WC, Bostan AC, Cheyne D (2006) Post-movement beta rebound is generated in motor cortex: Evidence from neuromagnetic recordings. *NeuroImage* 32:1281–1289.

Klostermann F, Nikulin VV, Kühn AA, Marzinzik F, Wahl M, Pogosyan A, Kupsch A, Schneider G-H, Brown P, Curio G (2007) Task-related differential dynamics of EEG alpha- and beta-band synchronization in cortico-basal motor structures. *Eur J Neurosci* 25:1604–1615.

Knutson B, Wimmer GE, Kuhnen CM, Winkielman P (2008) Nucleus accumbens activation mediates the influence of reward cues on financial risk taking. *Neuroreport* 19:509–513.

Kouneiher F, Charron S, Koechlin E (2009) Motivation and cognitive control in the human prefrontal cortex. *Nat Neurosci* 12:939–945.

Kühn AA, Kempf F, Brücke C, Gaynor Doyle L, Martinez-Torres I, Pogosyan A, Trottenberg T, Kupsch A, Schneider G-H, Hariz MI, Vandenberghe W, Nuttin B, Brown P (2008) High-frequency stimulation of the subthalamic nucleus suppresses oscillatory beta activity in patients with Parkinson's disease in parallel with improvement in motor performance. *J Neurosci* 28:6165–6173.

Kühn AA, Williams D, Kupsch A, Limousin P, Hariz M, Schneider G-H, Yarrow K, Brown P (2004) Event-related beta desynchronization in human subthalamic nucleus correlates with motor performance. *Brain* 127:735–746.

Lang W (2003) Surface recordings of the Bereitschaftspotential in normals. In: *The Bereitschaftspotential : movement-related cortical potentials*, pp 19–34. New York: Kluwer Academic

Publishers. Available at: <http://www.loc.gov/catdir/toc/fy036/2002040791.html> [Accessed February 19, 2010].

Leuthold H, Jentsch I (2002) Distinguishing neural sources of movement preparation and execution: An electrophysiological analysis. *Biological Psychology* 60:173–198.

Litvak V, Eusebio A, Jha A, Oostenveld R, Barnes GR, Penny WD, Zrinzo L, Hariz MI, Limousin P, Friston KJ, Brown P (2010) Optimized beamforming for simultaneous MEG and intracranial local field potential recordings in deep brain stimulation patients. *Neuroimage* 50:1578–1588.

Locke HS, Braver TS (2008) Motivational influences on cognitive control: behavior, brain activation, and individual differences. *Cogn Affect Behav Neurosci* 8:99–112.

Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* 164:177–190.

Mazzoni P, Hristova A, Krakauer JW (2007) Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. *J Neurosci* 27:7105–7116.

McIntyre CC, Hahn PJ (2010) Network perspectives on the mechanisms of deep brain stimulation. *Neurobiology of Disease* 38:329–337.

Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. *Proc Natl Acad Sci USA* 110:2641–2646.

Nichols T, Hayasaka S (2003) Controlling the familywise error rate in functional neuroimaging: a comparative review. *Stat Methods Med Res* 12:419–446.

Niv Y, Daw N, Joel D, Dayan P (2007) Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology* 191:507–520.

O'Connell RG, Dockree PM, Kelly SP (2012) A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* 15:1729–1735.

Oostenveld R, Fries P, Maris E, Schoffelen J-M (2011) FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience* 2011:1–9.

Özkurt TE, Butz M, Homburger M, Elben S, Vesper J, Wojtecki L, Schnitzler A (2011) High frequency oscillations in the subthalamic nucleus: a neurophysiological marker of the motor state in Parkinson's disease. *Exp Neurol* 229:324–331.

Pedersen JR, Johannsen P, Bak CK, Kofoed B, Saermark K, Gjedde A (1998) Origin of Human Motor Readiness Field Linked to Left Middle Frontal Gyrus by MEG and PET. *NeuroImage* 8:214–220.

Penny WD (2012) Comparing Dynamic Causal Models using AIC, BIC and Free Energy. *NeuroImage* 59:319–330.

Pogosyan A, Gaynor LD, Eusebio A, Brown P (2009) Boosting cortical activity at Beta-band frequencies slows movement in humans. *Curr Biol* 19:1637–1641.

Praamstra P, Schmitz F, Freund H-J, Schnitzler A (1999) Magneto-encephalographic correlates of the lateralized readiness potential. *Cognitive Brain Research* 8:77–85.

Robbins TW (2007) Shifting and stopping: fronto-striatal substrates, neurochemical modulation and clinical implications. *Phil Trans R Soc B* 362:917–932.

Salamone JD, Correa M (2012) The mysterious motivational functions of mesolimbic dopamine. *Neuron* 76:470–485.

Schmidt L, d' Arc BF, Lafargue G, Galanaud D, Czernecki V, Grabli D, Schüpbach M, Hartmann A, Lévy R, Dubois B, Pessiglione M (2008) Disconnecting force from money: effects of basal ganglia damage on incentive motivation. *Brain* 131:1303–1310.

Schmidt L, Lebreton M, Cléry-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural mechanisms underlying motivation of mental versus physical effort. *PLoS Biol* 10:e1001266.

Schnitzler A, Gross J (2005) Normal and pathological oscillatory communication in the brain. *Nat Rev Neurosci* 6:285–296.

Schoffelen J-M, Oostenveld R, Fries P (2005) Neuronal coherence as a mechanism of effective corticospinal interaction. *Science* 308:111–113.

Schultz W (2006) Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol* 57:87–115.

Schurger A, Sitt JD, Dehaene S (2012) An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci USA* 109:E2904–2913.

Shibasaki H, Hallett M (2006) What is the Bereitschaftspotential? *Clinical Neurophysiology* 117:2341–2356.

Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *NeuroImage* 46:1004–1017.

Swann N, Tandon N, Canolty R, Ellmore TM, McEvoy LK, Dreyer S, DiSano M, Aron AR (2009) Intracranial EEG Reveals a Time- and Frequency-Specific Role for the Right Inferior Frontal Gyrus and Primary Motor Cortex in Stopping Initiated Responses. *J Neurosci* 29:12675–12685.

Tachibana Y, Hikosaka O (2012) The primate ventral pallidum encodes expected reward value and regulates motor action. *Neuron* 76:826–837.

Toma K, Mima T, Matsuoka T, Gerloff C, Ohnishi T, Koshy B, Andres F, Hallett M (2002) Movement rate effect on activation and functional coupling of motor cortical areas. *J Neurophysiol* 88:3377–3385.

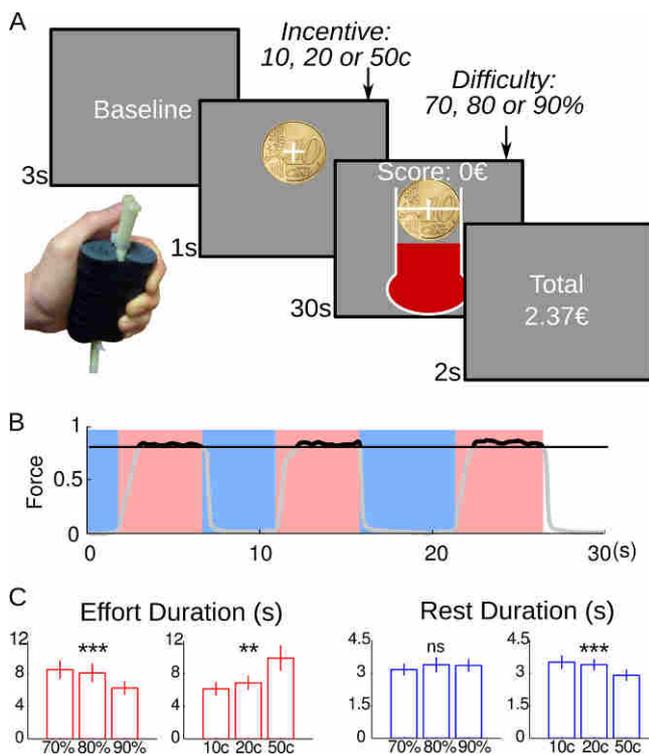
Tzagarakis C, Ince NF, Leuthold AC, Pellizzer G (2010) Beta-band activity during motor planning reflects response uncertainty. *J Neurosci* 30:11270–11277.

Uhlhaas PJ, Singer W (2006) Neural synchrony in brain disorders: relevance for cognitive dysfunctions and pathophysiology. *Neuron* 52:155–168.

Van Wijk BCM, Beek PJ, Daffertshofer A (2012) Neural synchrony within the motor system: what have we learned so far? *Front Hum Neurosci* 6:252.

Williams D, Kühn A, Kupsch A, Tijssen M, van Bruggen G, Speelman H, Hotton G, Loukas C, Brown P (2005) The relationship between oscillatory activity and motor reaction time in the parkinsonian subthalamic nucleus. *Eur J Neurosci* 21:249–258.

## Figures & Legends

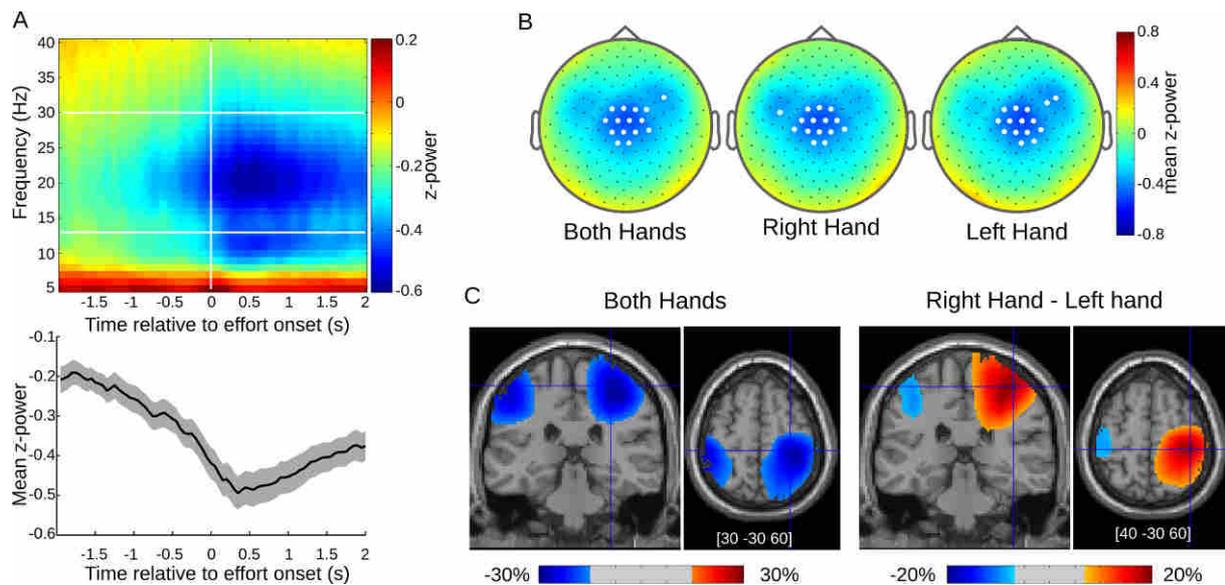


**Figure 1: Behavioral task and results**

A: The illustrated screenshots were successively presented every trial. When the thermometer image was displayed, participants could squeeze a handgrip in order to win money. Subjects were provided online feedback on force level and cumulative payoff. The payoff was only increased when force level was above the target bar, at a constant rate proportional to the monetary incentive. Two factors were manipulated over trials: the incentive (10, 20 or 50 cents), which was explicitly indicated as a coin image, and the difficulty, i.e. the force required to reach the target bar (70, 80 or 90% of maximal force), which remained implicit. The last screen indicated the money won so far, summed over all preceding trials.

B: Example recording of the force level produced during one trial. Three rest (blue shading) and effort (red shading) epochs could be defined. Force production was only rewarded when above the target threshold (here, 80% of maximal force), i.e. when plotted in black (not gray) on the graph.

C: Average data sorted by incentive and difficulty levels. Bars are mean effort and rest epoch durations and error bars the inter-subject standard errors. Significance of repeated-measure ANOVA main effects: \*\*\*  $p < 0.0005$ , \*\*  $p < 0.005$ , \*  $p < 0.05$ .

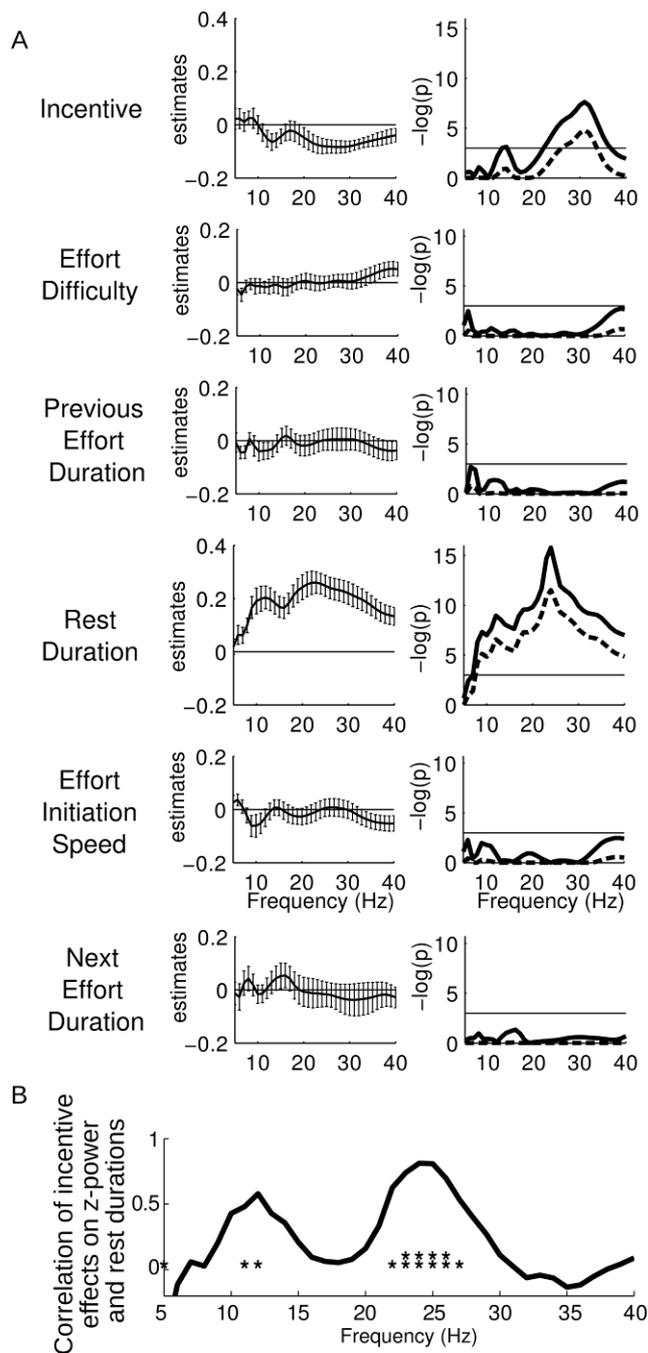


**Figure 2: Characterization of motor beta synchrony**

A: Beta synchrony level recorded from central sensors. The time-frequency map shows synchrony level around effort onset, averaged over central sensors (white dots on both hands topography). The mean ( $\pm$  s.e.m in grey) time series over the beta band (13-30 Hz) is plotted underneath. Z-power means that the power is z-scored relative to the trial baseline.

B: Scalp topography of beta synchrony level. Maps show synchrony level averaged over beta-range frequencies (13 to 30 Hz) in a time window centered on effort onset (-0.5 to 0.5 s), for pooled sessions and for left and right hand sessions, separately. White dots correspond to the 10% sensors with lowest synchrony level.

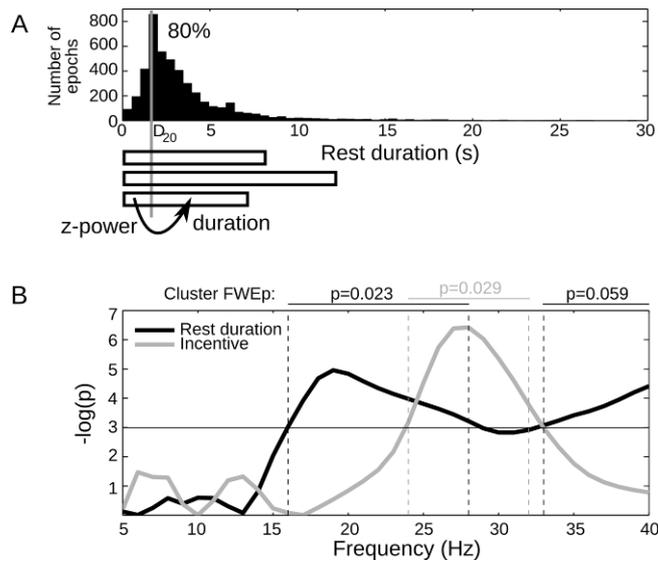
C: Sources of beta synchrony reduction. The color code indicates percentage of beta synchrony reduction relative to the baseline. Coronal and sagittal sections show regions exhibiting the most pronounced reduction, within the 2s preceding effort onset, irrespective of the hand used (left panels) and contrasted between hands (right panels). To superimpose sources onto brain anatomy, we arbitrarily used a bilateral threshold taken at half the extremes ( $\pm$ 14.5%) for left panels and a more liberal threshold ( $\pm$ 8%) to illustrate lateralization effects in right panels.



**Figure 3: Factors affecting motor beta synchrony**

A: Within-subject effects on z-power (z-scored power change relative to baseline level). The six rows show the effect of the six factors of interest included in the multiple regression analysis. The coefficients (betas) obtained at each frequency are shown in the left panels (as inter-subject mean  $\pm$  s.e.m.). Their significance was estimated at the group level and plotted as p-value logarithms in the right panels (plain lines: uncorrected p-values, dashed lines: FWE-corrected p-values). The horizontal line corresponds to  $p=0.05$ .

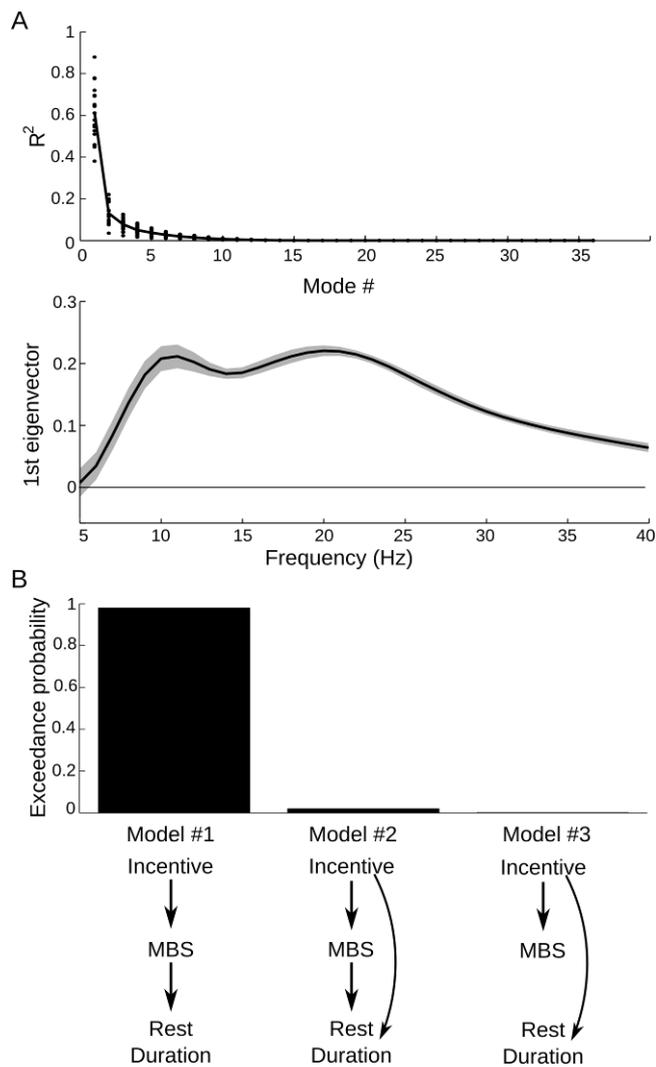
B: Across-subject correlations between incentive effects on rest duration and z-power (the latter being estimated in the null space of rest duration effect). Spearman correlation coefficients are plotted for each frequency in and around the beta band (5-40 Hz). \* <0.05 uncorrected p-value, \*\* <0.05 FWE corrected p-value.



**Figure 4: Early motor beta synchrony predicts rest duration**

A: Distribution of rest durations pooled over all participants. The underneath graph depicts the logic of the analysis: for each subject, the 20<sup>th</sup> percentile of rest duration ( $D_{20}$ ) was estimated and early z-power was defined as the mean z-power between rest onset and  $D_{20}$ . Then rest duration (together with incentive level) was regressed against early z-power, across all epochs that lasted longer than  $D_{20}$  (hence representing 80 % of the total).

B: Simultaneous regression of rest duration and incentive level against early z-power. Regression significance at the group level is shown as logarithm of uncorrected p-values, for each frequency in and around the beta band (5-40 Hz). Cluster-wise FWE-corrected p-values are also indicated, above the graph.



**Figure 5: Evidence for motor beta synchrony reduction mediating incentive effects on rest duration**

A: Singular value decomposition of z-power variations over epochs, in the 5 to 40 Hz range. Top: Dots represent R<sup>2</sup> statistics obtained for each mode and subject; solid line is the mean over subjects. Bottom: The curve indicates the first eigenvector (inter-subject mean +/- sem), for each frequency in and around the beta band. Since the direction of eigenvectors is per se arbitrary, we flipped them subject-wise for their mean value over the beta range to be positive, which allows direct visual comparison of the group average with the other figures.

B: Results of model comparison. The graphs illustrate the 3 models tested to account for statistical dependencies between incentive level, beta z-power and rest duration (with arrows representing linear links). Bars indicate model exceedance probability (i.e. the probability that the model is the most frequently implemented in the population).



## **Title**

How the brain decides when to work and when to rest: evidence for implicit cost-evidence monitoring

## **Authors**

Florent Meyniel<sup>1</sup>, Lou Safra<sup>1</sup> and Mathias Pessiglione<sup>1</sup>

<sup>1</sup>Motivation, Brain & Behavior (MBB) team, Institut du Cerveau et de la Moelle épinière (ICM), Groupe Hospitalier Pitié-Salpêtrière, Université Pierre et Marie Curie (UPMC – Paris 6)

## Abstract

A pervasive case of cost-benefit problem is how to allocate effort over time, i.e. deciding when to work and when to rest. An economic decision perspective would suggest that duration of effort is determined beforehand, depending on expected costs and benefits. However, the literature on exercise performance emphasizes that decisions are made on the fly, depending on physiological variables. Here, we propose and validate a general model of effort allocation that integrates these two views. In this model, a single variable, termed cost evidence, accumulates during effort and dissipates during rest, triggering effort cessation and resumption when reaching bounds. We assumed that such basic mechanism could explain implicit adaptation, whereas the free parameters (slopes and bounds) could be amenable to explicit anticipation. A series of behavioral experiments manipulating effort duration and difficulty was conducted in healthy humans to dissociate implicit from explicit computations. Results show 1) that subjects have no explicit access to the cost-evidence fluctuations that drive their behavior, 2) that effort and rest durations are adapted on the fly to variations in cost-evidence level, and 3) that actual difficulty impacts effort duration whereas expected difficulty impacts rest duration. Taken together, our findings suggest that cost evidence is implicitly monitored online, with an accumulation rate proportional to actual task difficulty. In contrast, cost-evidence bounds and dissipation rate might be adjusted in anticipation, depending on expected task difficulty.

## Introduction

Suppose that you are given a job whose payoff is proportional to the effort made within a limited time, say for instance the number of Christmas cards sold at the end of the day. Maximizing your payoff would require running from house to house, but this effort would induce such fatigue that you decide to walk from time to time. Such situation can be examined through economic decision theory (refs), which would suggest you to write down the expected costs and benefits, and try to figure out whether the effort is worthy. If the cost of a given effort is anticipated to increase with fatigue [1,2], then you will find an optimal duration that can be determined before engaging any action. Yet the literature on exercise performance has developed a different perspective on this issue [3,4], which would suggest that you start by running, and only stop when some physiological variable, for instance in cardiovascular function (such as heart beat rate) or in muscular metabolism (such as lactate concentration), attains a given limit [5,6] . In other words, effort cessation would be a reaction to homeostatic failure, and would not require any explicit anticipation of effort cost.

These two extreme perspectives have obvious limitations. The physiological view does not account for the effect of expectations that might pre-configure behavioral performance [4,7,8] The economic view does not integrate the constraints imposed by physiological reactions, which might be difficult to anticipate [9]. Here, we intend to overcome these limitations by integrating the two perspectives into a same computational model. Furthermore, we have built this model so as to explain the duration of not only effort exertion but also resting (recovery from fatigue). Let us assume that a single waning and waxing variable triggers decisions to stop and restart effort exertion when reaching bounds. As this variable linearly accumulates during effort and dissipates at rest, it can be seen as a simple reflection of physiological reactions that predict the proximity of homeostatic failure. Alternatively, it can be interpreted as tracking cost increase with fatigue, by integrating past effort over time. Thus, the basic architecture of the model (the accumulation-to-bound principle) can account for implicit, online adaptation to actual effort costs, complying with physiological constraints. On top of this basis, the modulation of the model free parameters (slopes and bounds) could allow for anticipatory adjustments, depending on expected costs and benefits.

To dissociate the effects of actual and expected effort costs, we developed seven variants of a paradigm that was employed in a previous study to identify the neural underpinnings of the modeled variable, which we termed cost evidence [10]. The task involved participants squeezing a handgrip with a given force, knowing that their payoff will be proportional to their effort duration.

Cost evidence can be manipulated by varying either an imposed duration or an imposed force (task difficulty). In a first study we demonstrate that subjects have no explicit access to the cost-evidence variable that drives the duration of their effort. In a second study we used three tasks that impose variable durations in order to verify that the behavior is locally adapted under physiological constraints. In a third study we used three other tasks that vary the difficulty in order to dissociate the effects of expected and actual costs.

## Results

### Introspection of cost evidence (study 1)

Our model posits that cost evidence increases during effort at a constant accumulation rate, which depends on task difficulty. Therefore, cost-evidence level should reflect the interaction between task difficulty and effort duration. The logic of this study was first to verify that behavioral choices were driven by the interaction of difficulty and duration, and second to examine whether introspective reports would reflect this interaction as well.

We first re-analyzed the behavioral choices observed in an Effort Allocation Task that involved subjects squeezing a handgrip in order to accumulate as much money as possible [10]. The payoff was calculated as the monetary incentive multiplied by the time spent above a target force level (which indexed task difficulty). Both the incentive (10, 20 or 50 cents) and difficulty levels (70, 80 or 90 % of maximal force) were varied across trials such that we could assess their effects on effort allocation. Incentive levels were sufficient for subjects making the effort and reach the target, but difficulty levels were too demanding for subjects to sustain their effort throughout trials, which lasted 30 seconds. Instead, they freely alternated effort and rest periods within trials (Figure 1A, top). We used the normalized cumulative distribution of effort durations to calculate the probability to stop the effort after a given duration at a given difficulty level. This probability was then fitted with a sigmoid function of cost-evidence level, which accounts for higher cost evidence making effort cessation more likely. Cost evidence was modeled as a linear combination of regressors meant to capture the impact of duration and difficulty. We considered three possibilities: main effects of duration and difficulty, non-linear effects (power functions) of duration and difficulty, and interaction between duration and difficulty. Including or not each of this possibility in the linear combination made a total of eight models, which we compared using Bayesian model selection (Figure 1B, top). The most plausible model was pure interaction, with an expected frequency  $ef=0.62$  (which is much higher than chance level -  $1/8$ ) and exceedance probability  $xp=0.988$  (confidence that the model is more frequently followed than the others).

For introspective reports we had a different group of participants performing a Cost Rating Task, in which both effort duration (from 3 to 7 s) and task difficulty (from 40 to 60 % of maximal force) were imposed and varied experimentally (Figure 1A, bottom). To keep similarity with the Effort Allocation Task, we also manipulated the incentive level. On each trial, the payoff was calculated as

the incentive multiplied by the fraction of the imposed duration that subjects spent squeezing at the required target force level or higher. As participants were asked to be as accurate as possible, this fraction was almost 100% (mean over subjects: 98.7%, extreme subjects: 94.6% and 99.9%). The difference between required and produced force levels did not vary significantly across conditions (incentive:  $p = 0.28$ ; duration:  $p=0.99$ ; difficulty:  $p=0.59$ ; interactions between these factors: all  $p>0.21$ ), suggesting that effort production was well controlled by the experimental design. The quantity that subjects had to rate was framed in terms of exhaustion, which makes the introspection of cost more intuitive. The precise question was ‘How much are you exhausted now?’ and the response scale was ranging from ‘not at all’ to ‘completely’. Cost ratings were not impacted by incentives ( $p=0.1$ ), and marginally by the initial position of the cursor on the scale ( $p=0.056$ ). Critically, cost ratings increased with both duration (beta:  $1.9 \pm 0.79$  s.e.m.,  $p=0.028$ ) and difficulty (beta:  $3.2 \pm 0.49$  s.e.m.  $p=5 \cdot 10^{-6}$ ), without significant interaction between these factors ( $p=0.96$ ). We then fitted cost ratings with the same linear combinations as used for fitting choice probability (Figure 1B, bottom). Bayesian model selection confirmed the absence of significant interaction between duration and difficulty, since the best model was simply additive (chance level is  $1/8$ ,  $ef=0.48$ ;  $xp=0.93$ ). In principle, this additive effect could arise from half the subjects reporting duration and the other half reporting difficulty. This would imply that the effect sizes of these factors are anti-correlated across subjects. We found the opposite result (Pearson rho:  $0.82$ ,  $p=3 \cdot 10^{-5}$ ), suggesting that subjects who were good at perceiving duration were also good at perceiving difficulty. Yet they reported the addition of the two dimensions, and not their product, as should be the case if they were simply introspecting cost evidence.

To directly compare the form of cost evidence inferred from behavioral choices and from introspective reports, we fitted a model with constant elasticity of substitution (CES) between duration and difficulty (see methods). This model has a free parameter that captures the curvature of cost in the duration by difficulty space. This curvature is theoretically convex (parameter below one) in the case of interaction and null (parameter equal to one) in the absence of interaction. We found that the curvature parameter was significantly below one in the Effort Allocation Task (median:  $0.52$ , s.e.m.:  $0.06$ , bilateral sign-test of the median against 1:  $p=6.7 \cdot 10^{-8}$ ), but not in the Cost Rating Task (median:  $1.01$ , s.e.m.:  $0.12$ :  $p=1$ ), with a significant difference between tasks ( $p=.3 \cdot 10^{-6}$ , bilateral Wilcoxon rank sum test for equal medians).

When debriefing the Cost Rating Task, participants unambiguously reported having noticed variations in both difficulty and duration. When asked whether one of these two factors had a greater impact on their ratings, 13 subjects favored the duration, 3 favored the difficulty, and 2 could

not favor one or the other, describing something like an interaction. However, comparison of standardized effect size revealed a greater impact of difficulty on ratings (paired t-test,  $p=0.016$ ). Among the 16 subjects who favored a main effect, 12 got it wrong (the other factor had a higher impact on their ratings), which is more than expected by chance (binomial test,  $p=0.028$ ).

To summarize, the costs reported in subjective ratings do not have the same shape as the costs inferred from behavioral choices. What subjects report is an addition of duration and difficulty, whereas what drives their behavior is an interaction between the two. Additionally, at a meta-cognitive level, subjects have poor insight into the factors that modulate their feeling of exhaustion.

## **Behavioral adaptation to cost evidence (study 2)**

In our previous study [10], we suggested that the alternation of effort and rest periods observed in the Effort Allocation Task was well explained by a waning and waxing accumulation signal. However, this cost-evidence signal that we localized in the brain could be epiphenomenal, in the sense that it would not reflect any causal mechanism triggering the decisions to stop and restart effort. In this second study, we wished to verify that the level of cost evidence imposes actual constraints on subsequent behavior, as predicted by the accumulation-to-bound principle. We therefore tested the predictions of the accumulation model on the behavior that followed an effort whose duration was imposed. The difficulty was not manipulated in this study, for several reasons: firstly, the effect of difficulty was already shown in the previous study and will be further investigated in the next one (study 3 described hereafter), and secondly, manipulation of difficulty only applies to effort periods, whereas manipulation of duration can be equally applied to both effort and rest periods. Predictions of the accumulation model are that 1) prolonging effort should decrease the next effort period (if compensatory resting is not allowed), 2) prolonging rest should increase the next effort period (up to a maximum corresponding to full recovery), prolonging effort should increase the next rest period (if compensatory resting is allowed). These three predictions were tested in different groups of participants, using three variants of the Effort Allocation Task (Figure 3A). The three tasks had the same structure, with first an imposed effort (between start and stop signals), second a rest period (either fixed or free) and third a free effort exertion. Difficulty of both efforts was fixed at 60 % of the maximal force, and payoff was proportional to the duration of the last effort.

### **Effort is adapted to accumulated cost evidence (Task 1)**

In this task, cost evidence was increased by prolonging the first effort period (from 1 to 10s), then the second effort duration was observed after a fixed 2-s rest (Figure 2A, left). To ensure that the rest duration was well controlled, we checked that initiation delay of the second effort after the go signal was not significantly impacted by the duration of the first effort ( $p = 0.09$ ), by cumulated duration of efforts produced in the current session ( $p = 0.76$ ), and by the session number ( $p = 0.25$ ). Critically, the second effort was significantly shortened by prolonging the first effort ( $p = 0.0037$ ). We also estimated in the same model the effect of the initiation delay, which was significant ( $p = 0.001$ ), even after removing the variance explained by the other regressors (imposed effort duration, cumulated past effort duration and session number). As the correlation was negative, it could reflect trial-to-trial fluctuation of fatigue on top of (and orthogonal to) the experimental manipulation: in trials when fatigue is lower, participants initiate the effort faster and sustain it longer.

Next we examined the shape of the transfer function from imposed to observed effort duration. The model predicts that this link should be negative, except if resting is long enough to fully dissipate the accumulated cost. We therefore compared a model with pure negative correlation (no saturation) to models with an upper plateau (over shortest efforts), followed by a decrease. We tried two possibilities for this saturation effect: first a constant followed by a linear decrease and second a negative exponential. The latter was implemented because it provides a better fit of plateau effects when data are noisy (see methods). Bayesian model selection revealed that the pure linear model was far better than the two saturation models (model 1 versus models 2 & 3, family comparison: chance level is  $\frac{1}{2}$ ,  $ef=0.81$ ,  $xp=0.96$ ). Thus, the result supports linear accumulation of cost evidence, which limits subsequent effort production due to the existence of an upper bound. However, we found no evidence for the existence of a lower bound in cost dissipation, probably because our rest period was not long enough. This limitation was overcome in the next task, where rest period was systematically varied.

## **Effort is adapted to dissipated cost evidence (Task 2)**

This task (Figure 2A, middle) was very similar to Task 1, except that effort duration was now fixed (to 7 s) and rest duration was systematically varied (from 1 to 12 s). We checked again that subjects were not delaying effort initiation to compensate for variations in the imposed rest duration ( $p = 0.10$ ). In addition we found that the initiation delay was slightly affected by the cumulated duration of past efforts ( $p = 0.03$ ), but not by the session number (0.46). Critically, observed effort was significantly prolonged by longer rest ( $p = 0.0035$ ). We also found a significant effect of initiation delay (after removing the variance explained by the imposed effort duration, cumulated past effort

duration and session number) on observed effort duration ( $p = 0.026$ ), which could reflect additional trial-to-trial fatigue fluctuations as in Task 1.

Next we tested the existence of a saturation, meaning that beyond a certain rest duration, cost evidence is entirely dissipated and subsequent effort cannot be further prolonged. As was done for the previous task, we compared three models for the link from rest to effort duration: 1) a linear effect (no saturation), 2) a linear effect bounded by an upper plateau (over longest rests), 3) an exponential asymptotic plateau. Bayesian model selection showed that the saturation family was now more plausible (models 2 and 3 versus model 1, chance level is  $\frac{1}{2}$ ,  $ef=0.79$ ,  $xp=0.94$ ). Direct comparison between models 2 and 3 revealed that the asymptotic saturation was more likely than the linear plateau ( $xp=0.98$ ). Thus, the results confirm that prolonging rest after a first effort augments the capacity to produce a second effort, as if cost evidence was dissipated. Moreover, the saturation effect suggests the existence of a threshold after which prolonging rest is useless, which would correspond to a lower bound for cost-evidence dissipation.

### **Rest is adapted to accumulated cost evidence (Task 3)**

This task (Figure 2A, right) was quite similar to Task 2, except that participants were not asked to resume their effort immediately at the go signal, but only when they felt ready to do so. There were therefore two dependent variables of interest: rest duration and subsequent effort duration. Critically, rest duration was significantly increased by prolonging the imposed effort duration ( $p = 5 \cdot 10^{-4}$ ). We also found that rest duration increased across sessions ( $p = 0.028$ ) and with the cumulated duration of the efforts produced within sessions ( $p = 5 \cdot 10^{-5}$ ). This probably reflects more rest being allocated to compensate for the fatigue that accumulates at longer time scales, as was observed in Tasks 1 and 2.

We expected that participants would rest long enough to fully dissipate the first effort cost, which hence would have no impact on the second effort duration. This was not the case: prolonging the first effort significantly shortened the second effort ( $p = 0.006$ ). Note that the rest duration necessary to fully compensate for the 7-s effort imposed in Task 2 was 5.9s (estimated from the plateau fitted with median parameter values). According to the linear fit in Task 3 (Figure 2B, left), subjects rested on average for  $4.8 \text{ s} \pm 0.4$  after 7 s of effort. Thus, subjects did not wait long enough to compensate for the imposed effort cost. This partial recovery might explain why fatigue accumulated, such that effort duration also decreased both across sessions ( $p = 0.008$ ) and within sessions ( $p = 4 \cdot 10^{-4}$ ), as was observed in Tasks 1 and 2.

## Dissociation of implicit from explicit cost processing (study 3)

The two studies presented so far are compatible with a completely implicit and automatic model, in which decisions to cease and resume effort production are controlled by an internal variable fluctuating between bounds that might be determined by physiological constraints. In our previous study, we observed that task difficulty shortened effort duration, which could reflect cost evidence (difficulty times duration) reaching the upper bound, but did not affect rest duration. We hypothesized that the last observation could arise from task difficulty not being made explicit to participants. Indeed, monetary incentives, contrary to the difficulty levels, were explicitly presented with coin images at trial start and affected both effort and rest durations (with longer effort and shorter rest for higher incentive).

We therefore tested whether providing explicit information about difficulty level would change the way participants process cost evidence. We constructed three variants of the Effort Allocation Task, which were administered to three different groups of participants. The Implicit Task (Figure 1A, top) is the task used in our previous study [10], with no visual cue for difficulty level. In the Explicit Task, the only change is that difficulty level (percentage of maximal force: 70, 80 or 90%) was announced before the beginning of trials, on the same screen as incentive level. In the Dissociation Task, we kept the explicit cues, but they were no longer predictive of the actual task difficulty. To maintain sufficient statistical power, only two difficulty levels were used (75 and 85%), in a full factorial design (two cued difficulties crossed by two actual difficulties). This design was meant to disentangle the effects of implicit versus explicit cost processing. Monetary incentives were also manipulated in all tasks and crossed with the three (Implicit and Explicit Tasks) or four (Dissociation task) cells corresponding to variations in difficulty. We only used two incentive levels (10 versus 20c) in the Dissociation task to avoid combinatorial inflation. In every task, the effect of experimental factors (incentive, actual and cued difficulty) on effort and rest durations were estimated in separate multiple linear regressions.

### Comparison of Implicit and Explicit Tasks

As previously shown, in the Implicit Task (Figure 3, left), effort duration was both longer for higher incentive ( $p = 8.1 \cdot 10^{-7}$ ) and shorter for higher difficulty ( $p = 1.6 \cdot 10^{-10}$ ). In contrast, rest duration was shorter for higher incentives ( $p = 2.0 \cdot 10^{-5}$ ) but was not modulated by the difficulty ( $p = 0.32$ ).

Interactions were included in the regression model, but the incentive x difficulty interaction was not significant, neither for effort or for rest duration (all  $p > 0.084$ ).

All significant results were replicated in the Explicit Task (Figure 3, middle): effort duration was both longer for higher incentive ( $p = 1.1 \cdot 10^{-3}$ ) and shorter for higher difficulty ( $p = 6.0 \cdot 10^{-6}$ ), and rest duration was shorter for higher incentive ( $p = 9.7 \cdot 10^{-4}$ ). The novel result is that rest duration was now increased by higher difficulty ( $p = 1.6 \cdot 10^{-3}$ ), which was correctly cued at trial start. The difference in standardized effect sizes between Implicit and Explicit Tasks was also significant ( $p = 1.2 \cdot 10^{-4}$ ). All interactions remained non-significant, neither for effort or rest duration (all  $p > 0.1$ ). Thus, the difficulty in the Explicit Task, which was both expected and experienced during effort exertion, affected both effort and rest durations. We reasoned that the two effects could be dissociable, with the actual difficulty affecting effort duration, and the expected difficulty affecting rest duration.

## **Analysis of the Dissociation Task**

In the Explicit Task (Figure 3, right), the two levels of actual and cued difficulty were manipulated independently. As in the Implicit and Explicit tasks, higher incentive increased effort duration ( $p = 0.022$ ) and shortened rest duration ( $p = 1.5 \cdot 10^{-3}$ ). Effort duration was affected by the actual ( $p = 0.021$ ) but not by the cued difficulty ( $p = 0.64$ ). The difference in standardized effect size was at significance limit ( $p = 0.050$ ). We also verified that the effect of cued difficulty on effort duration in the Dissociation Task was significantly lower than the (actual) difficulty effects observed in the Implicit ( $p = 4.3 \cdot 10^{-7}$ ) and Explicit ( $p = 2.3 \cdot 10^{-6}$ ) tasks. Conversely, rest duration was affected by the cued ( $p = 1.7 \cdot 10^{-3}$ ) but not by the actual ( $p = 0.63$ ) difficulty. The difference in standardized effect size was as well significant ( $p = 0.045$ ). We also verified that the effect of cued difficulty on rest duration was higher in the Dissociation Task than the (actual) difficulty effect observed in the Implicit Task ( $p = 0.002$ ), and that the effect of actual difficulty in the Dissociation Task was lower than the (cued) difficulty effect observed in the Explicit Task ( $p = 0.008$ ). Thus, within- and between-task comparisons both support a double dissociation between the actual and cued difficulty effects on effort and rest durations.

As some critical p-values were near 0.05 type I error rate, we conducted a permutation test to ensure the reliability of the parametric t-distribution in our small sample. This permutation-based t-distribution yielded the same exact p-values up to the 3<sup>rd</sup> decimal. Second and third order interaction terms between incentive, cued and actual difficulty were included in the model, but none

of them was significant neither for rest or effort duration (all  $p > 0.18$ ). We also checked that there was no interaction of cued difficulty with time, which could potentially reflect a progressive discount of the cue effect (which does not predict the actual difficulty). Time was modeled at three nested scales (rest or effort period position within a trial, trial position within a session, and session number). Two-way interactions with cued difficulty were estimated for each time scale: none of them was significant (all  $p > 0.25$ ).

## Bayesian Model Selection

We compared different versions of our accumulation model to identify how the free parameters ( $A$ : amplitude between bounds,  $S_E$ : accumulation slope during effort, and  $S_R$ : dissipation slope during rest) were affected by the experimental factors ( $I$ : Incentive,  $D_a$ : actual difficulty,  $D_c$ : cued difficulty). We started with the formalization that we proposed in our previous publication (Meyniel 2013) to account for the behavior observed in the Implicit Task. All models were built as a set of three equations that defines each parameter as a linear combination of the different factors (see methods). Only models that can produce the behavioral results (significant effect on effort or rest duration) were included in the space covered by Bayesian Model Selection. In the Implicit Task, this left 24 possible models (see Table 2A) with one that was much more plausible than the others (chance level is  $1/24$ ,  $ef = 0.30$ ,  $x_p = 0.90$ ).

For the novel tasks (Explicit and Dissociation), we explored two possibilities for integrating the additional factor (cued difficulty). The first possibility was to integrate it as an additive term, just as was done with actual difficulty (see Table 2B and 1C). Note that these purely linear models do not enable dissociating the effects of actual and expected difficulty in the Explicit Task. The second possibility was to integrate cued difficulty as a hyperbolic discount of incentives, which is quite standard in the literature for capturing not only delay discounting [11,12] [12] but also for effort discounting [13]. Thus, for the novel tasks that manipulate expected difficulty, we included the hyperbolic equivalent of our linear models (see Table 2D). With this hyperbolic version, we can dissociate the effect of actual and expected difficulty (the former is linear, the latter hyperbolic) even in the Explicit Task where the two factors are confounded.

Family comparison revealed that there was far more evidence in favor of a hyperbolic rather than linear discount of incentives by cued difficulty, in both the Explicit and Dissociation tasks (chance level is  $1/2$ ,  $ef > 0.91$ ,  $x_p > 0.999$ ). Among the 78 possible hyperbolic models, a best model was identified with  $x_p = 0.90$  (chance level is  $1/78$ ,  $ef = 0.13$ ) in the Dissociation Task and with  $x_p = 0.82$

(chance level is 1/78,  $ef=0.14$ ,) in the Explicit Task. Crucially, the best hyperbolic model identified in the Explicit and Dissociation tasks was the same model, which also corresponded to the best model identified in the Implicit Task (where modulation by cued difficulty is necessarily absent). This best model is written as follows ( $Te$  and  $Tr$  being effort and rest duration,  $\alpha$ ,  $\beta$ ,  $\gamma$  the coefficients and  $I$ ,  $Da$  and  $Dc$  the incentive, actual difficulty and cued difficulty levels):

$$Te = \frac{A}{Se}; \quad Tr = \frac{A}{Sr}$$

$$\left\{ \begin{array}{l} A = \alpha_m + \alpha_I I \\ Se = \beta_m + \beta_{Da} Da \\ Sr = \gamma_m + \frac{\gamma_I I}{1 + \gamma_{Dc} Dc} \end{array} \right.$$

A graphical interpretation of the model with a summary of the observed effects is provided in Figure 4. In short, incentives impacted both the amplitude between bounds and the dissipation rate, resulting in longer effort and shorter rest for higher incentives. The effect of task difficulty was computationally dissociable: higher actual difficulty accelerated the accumulation, resulting in shorter effort, whereas higher expected difficulty slowed the dissipation, resulting in longer rest.

## Discussion

In our previous study we addressed the issue of how the brain allocates effort production over time, in a situation where the payoff depends on the total effort duration. We found a neural signal that was ramping up and down during effort and rest periods and that could, in principle, trigger the decisions to stop and restart effort production. Here we provide evidence that the core accumulation mechanism is reactive and implicit. Indeed, cost evidence is defined as the product of effort duration and difficulty, but participants were unable to report this product explicitly (study 1). However, they adapted their effort allocation when we implicitly manipulated both the duration (study 2) and the difficulty (study 3). In addition, we suggest that some parameters of the accumulation process are susceptible to anticipatory adjustment. Indeed, we found that expected benefit and difficulty could modulate the distance between bounds as well as the dissipation rate during rest. The dissociation of implicit and explicit cost processing could reconcile the perspectives offered by sport physiology on the one hand, and economic theory of choice on the other hand.

The implicit part of the model - monitoring cost evidence and triggering decisions when bounds are attained, accords well with the literature on exercise performance [4,8]. Although it was developed to explain how athletes pace their running on treadmills, we can borrow the notion that behavioral changes are reactions to physiological variables reaching homeostatic borders. Results of Study 2 show that bounds between which the cost-evidence signal fluctuates are true limits that determine the decisions to stop and restart effort exertion. On the contrary, the explicit part of the model – adjusting the behavior depending on expected benefit and difficulty, is consistent with the literature on value-based decision-making [14,15]. It is quite remarkable in Study 3 that the computational effect of actual (implicit) difficulty during effort was simply linear, as in a passive accumulation, whereas the effect of expected (explicit) difficulty was hyperbolic, as in economic models of discounting [13,16]. It should be acknowledged that mixtures of anticipatory calculations and on-line adaptations are frequently used in motor control theory [17,18], for instance to explain how movement trajectory can be adjusted to internal noise or unexpected target displacement. However, these models have not integrated the conflict between costs and benefits until very recently [19,20]. Finally, we note that the perspectives offered by the literatures on exercise performance and value-based choice only explain the duration of effort, without further specification they say nothing about the duration of rest. Our model accounts for the timing of both effort and rest, within a same accumulation framework.

Examining whether the accumulation mechanism is optimal or not would go beyond the scope of this study. It can nonetheless be seen has a heuristic mechanism that certainly has advantages. Physiologically, it ensures that effort production does not put the body at threat, avoiding for instance damage to the muscles. In this view the signal would indicate the likelihood of physiological damage, and the upper bound would implement a threshold on that risk. Economically, it ensures that costs do not overcome benefits. In this view, the signal would indicate the cost and the upper bound the benefit of the potential effort at the next time point. Mixing predictive and reactive processes also presents advantages. Online monitoring of effort consequences allows adjusting cost estimation, which is usually unknown beforehand, as in our implicit version of the task. Anticipatory estimation allows deciding whether or not to engage the action, and scaling energy expenditure to expected costs and benefits. In our case, this means spending more time at work and less time at rest when the net value of effort is higher.

The two periods, effort and rest, are not equivalent though. While monitoring cost evidence during effort might be a passive process (mechanically integrating difficulty over time), dissipating cost evidence during rest seems more active. Indeed, the dissipation rate was susceptible to modulation by explicit information (monetary incentive and cued difficulty). Moreover, the assertion that subjects have no explicit access to cost evidence was only tested in Study 1 for the effort period. It remains possible that during rest, subjects are fully aware of the cost-evidence level, and hence of how much effort they would be able to produce next. We could have tried to test whether their introspective reports integrate duration with cued difficulty after a given rest, but asking the question in this case would have been awkward.

Using dissipation as well as accumulation in order to explain behavioral choices is a major difference between our model and the standard evidence accumulation models. Classically, accumulation of evidence is meant to improve the estimation of a stationary noisy input, whether external, as in perceptual decision-making [21–23] or internal, as in value-based decision making [24–26]. The fact that the cost evidence variable dissipates at rest rules out the possibility that this signal simply reflects an integration of the force produced throughout the trial (which can only increase). It is plausible that the signal reflects an input that is already dynamical (and not stationary). This might be true not only at the theoretical level, if we interpret it as signaling the potential effort cost or the proximity of exhaustion, but also at the biological level. For instance, our cost evidence signal could relay the accumulation and dissipation of a by-product of effort exertion, which could integrate several variables such as lactate concentration, stretch of muscle fibers or heart beat rate. Alternatively, the cost-evidence signal could reflect increase in the efferent drive needed to

overcome fatigue and maintain motor output [27]. Using combined fMRI and MEG recordings, we localized the cost-evidence signal in proprioceptive areas (posterior insula). This localization would incline us to situate the input in the afferent proprioception or pain pathways coming from the muscles [28]. However, the fact that subjects had a poor introspection into that signal argues against the idea that it represents the neural counterpart of a common and intuitive sensation such as pain or fatigue.

Yet the fact that cost evidence dissipation could be modulated depending on expected benefit and difficulty suggests that other neural processes occur during rest than passive transmission of effort-induced physiological perturbation. First, the dissipation of cost evidence could be linked to the preparation of the next effort. Such preparation is reflected by motor signals such as the readiness potential (Lang 2003; Shibasaki and Hallett 2006) or the de-synchronization of beta oscillations [29,30]. We showed in a previous study (Meyniel submitted) that the last process is modulated by incentive level, it could therefore mediate the effect of motivation on cost dissipation in the posterior insula. Second, the dissipation of cost evidence could be accentuated by analgesic mechanisms. The posterior insula region that signals cost evidence is also involved in pain perception [31,32] and placebo effect [33]. The placebo effect suggests that the brain has an endogenous means to control pain, possibly through the opioid system [34–36] Another possibility would be the serotonin, which participates to the analgesia induced by common pain killers such as acetaminophen [37] and to the sensation of fatigue during effort [1,5]. Thus, through opioids or serotonin, the brain might be able to regulate cost-related signals depending on motivation level.

Before concluding, we must acknowledge a number of limitations and inconsistencies. First, the situation explored in our paradigm is highly restricted. Notably, subjects are only allowed to adjust the duration of their effort and not the intensity, which we can usually adjust in ecological situations. We have conducted a series of studies where the payoff was based on effort intensity [38–40], but we still have to explore the situation where the two dimensions can vary. Second, the model does not account for a number of observations, for instance the fact that fatigue accumulates at longer time scales. Indeed, we observed in some tasks an effect of trial and/or session number on effort duration, which could mean that cost is not fully dissipated after each rest period, or that slower effort-induced perturbations are accumulated elsewhere and imposed constraints on performance.

Despite these limitations, the results of the present studies taken together provide strong evidence that costs are implicitly monitored in order to adapt effort duration on the fly, which can be

dissociated from anticipatory adjustments depending on expected costs and benefits. Moreover, this dissociation was computationally tractable and might be of clinical relevance. It would suggest the existence of two different kinds of apathy: effort could be limited because the expected cost is over-estimated, or because the actual effort-induced cost is inflated. The first category (perhaps in depression disorders) would rest a lot but would encounter little difficulty in maintaining their effort once it is engaged, whereas the other (perhaps in chronic fatigue) would easily initiate efforts but then would rapidly renounce.

## References

1. Gandevia SC (2001) Spinal and Supraspinal Factors in Human Muscle Fatigue. *Physiol Rev* 81: 1725–1789.
2. Ma L, Chablat D, Bennis F, Zhang W (2009) A new simple dynamic muscle fatigue model and its validation. *International Journal of Industrial Ergonomics* 39: 211–220. doi:10.1016/j.ergon.2008.04.004.
3. Noakes TD (2000) Physiological models to understand exercise fatigue and the adaptations that predict or enhance athletic performance. *Scand J Med Sci Sports* 10: 123–145.
4. Abbiss CR, Laursen PB (2005) Models to explain fatigue during prolonged endurance cycling. *Sports Med* 35: 865–898.
5. Boyas S, Guével A (2011) Neuromuscular fatigue in healthy muscle: underlying factors and adaptation mechanisms. *Ann Phys Rehabil Med* 54: 88–108. doi:10.1016/j.rehab.2011.01.001.
6. Shephard RJ (2009) Is it time to retire the “central governor”? *Sports Med* 39: 709–721. doi:10.2165/11315130-000000000-00000.
7. Tucker R (2009) The anticipatory regulation of performance: the physiological basis for pacing strategies and the development of a perception-based model for exercise performance. *Br J Sports Med* 43: 392–400. doi:10.1136/bjism.2008.050799.
8. Noakes TD (2011) Time to move beyond a brainless exercise physiology: the evidence for complex regulation of human exercise performance. *Appl Physiol Nutr Metab* 36: 23–35. doi:10.1139/H10-082.
9. Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7: 907–915. doi:10.1038/nn1309.
10. Meyniel F, Sergent C, Rigoux L, Daunizeau J, Pessiglione M (2013) Neurocomputational account of how the human brain decides when to have a break. *Proc Natl Acad Sci USA* 110: 2641–2646. doi:10.1073/pnas.1211925110.
11. Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10: 1625–1633. doi:10.1038/nn2007.
12. Peters J, Büchel C (2011) The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends Cogn Sci (Regul Ed)* 15: 227–239. doi:10.1016/j.tics.2011.03.002.
13. Prévost C, Pessiglione M, Météreau E, Cléry-Melin M-L, Dreher J-C (2010) Separate valuation subsystems for delay and effort decision costs. *J Neurosci* 30: 14080–14090. doi:10.1523/JNEUROSCI.2752-10.2010.
14. Walton ME, Kennerley SW, Bannerman DM, Phillips PEM, Rushworth MFS (2006) Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks* 19: 1302–1314. doi:10.1016/j.neunet.2006.03.005.

15. Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9: 545–556. doi:10.1038/nrn2357.
16. Botvinick MM, Huffstetler S, McGuire JT (2009) Effort discounting in human nucleus accumbens. *Cognitive, Affective, & Behavioral Neuroscience* 9: 16–27. doi:10.3758/CABN.9.1.16.
17. Wolpert DM, Ghahramani Z (2000) Computational principles of movement neuroscience. *Nat Neurosci* 3: 1212–1217. doi:10.1038/81497.
18. Todorov E (2006) Optimal control theory. In: Doya K, editor. *Bayesian brain: probabilistic approaches to neural coding*. MIT Press. pp. 269–298.
19. Rigoux L, Guigon E (2012) A Model of Reward- and Effort-Based Optimal Decision Making and Motor Control. *PLoS Comput Biol* 8: e1002716. doi:10.1371/journal.pcbi.1002716.
20. Baraduc P, Thobois S, Gan J, Broussolle E, Desmurget M (2013) A common optimization principle for motor execution in healthy subjects and parkinsonian patients. *J Neurosci* 33: 665–677. doi:10.1523/JNEUROSCI.1482-12.2013.
21. Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. *Annu Rev Neurosci* 30: 535–574. doi:10.1146/annurev.neuro.29.051605.113038.
22. Heekeren HR, Marrett S, Ungerleider LG (2008) The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci* 9: 467–479. doi:10.1038/nrn2374.
23. Wang X-J (2012) Neural dynamics and circuit mechanisms of decision-making. *Curr Opin Neurobiol* 22: 1039–1046. doi:10.1016/j.conb.2012.08.006.
24. Krajbich I, Armel C, Rangel A (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13: 1292–1298. doi:10.1038/nn.2635.
25. Basten U, Biele G, Heekeren HR, Fiebach CJ (2010) How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci U S A* 107: 21767–21772. doi:10.1073/pnas.0908104107.
26. Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340: 95–98. doi:10.1126/science.1233912.
27. Marcora S (2009) Perception of effort during exercise is independent of afferent feedback from skeletal muscles, heart, and lungs. *J Appl Physiol* 106: 2060–2062. doi:10.1152/jappphysiol.90378.2008.
28. Eickhoff SB, Jbabdi S, Caspers S, Laird AR, Fox PT, et al. (2010) Anatomical and Functional Connectivity of Cytoarchitectonic Areas within the Human Parietal Operculum. *J Neurosci* 30: 6409–6421. doi:10.1523/JNEUROSCI.5664-09.2010.
29. Van Wijk BCM, Beek PJ, Daffertshofer A (2012) Neural synchrony within the motor system: what have we learned so far? *Front Hum Neurosci* 6: 252. doi:10.3389/fnhum.2012.00252.
30. Engel AK, Fries P (2010) Beta-band oscillations--signalling the status quo? *Curr Opin Neurobiol* 20: 156–165. doi:10.1016/j.conb.2010.02.015.

31. Friebel U, Eickhoff SB, Lotze M (2011) Coordinate-based meta-analysis of experimentally induced and chronic persistent neuropathic pain. *Neuroimage* 58: 1070–1080. doi:10.1016/j.neuroimage.2011.07.022.
32. Wager TD, Atlas LY, Lindquist MA, Roy M, Woo C-W, et al. (2013) An fMRI-based neurologic signature of physical pain. *N Engl J Med* 368: 1388–1397. doi:10.1056/NEJMoa1204471.
33. Wager TD, Rilling JK, Smith EE, Sokolik A, Casey KL, et al. (2004) Placebo-Induced Changes in fMRI in the Anticipation and Experience of Pain. *Science* 303: 1162–1167. doi:10.1126/science.1093065.
34. Benedetti F (2008) Mechanisms of placebo and placebo-related effects across diseases and treatments. *Annu Rev Pharmacol Toxicol* 48: 33–60. doi:10.1146/annurev.pharmtox.48.113006.094711.
35. Staahl C, Olesen AE, Andresen T, Arendt-Nielsen L, Drewes AM (2009) Assessing analgesic actions of opioids by experimental pain models in healthy volunteers – an updated review. *British Journal of Clinical Pharmacology* 68: 149–168. doi:10.1111/j.1365-2125.2009.03456.x.
36. Petrovic P, Kalso E, Petersson KM, Ingvar M (2002) Placebo and opioid analgesia-- imaging a shared neuronal network. *Science* 295: 1737–1740. doi:10.1126/science.1067176.
37. Smith HS (2009) Potential analgesic mechanisms of acetaminophen. *Pain Physician* 12: 269–280.
38. Schmidt L, D' Arc BF, Lafargue G, Galanaud D, Czernecki V, et al. (2008) Disconnecting force from money: effects of basal ganglia damage on incentive motivation. *Brain* 131: 1303–1310. doi:10.1093/brain/awn045.
39. Schmidt L, Cléry-Melin M-L, Lafargue G, Valabregue R, Fossati P, et al. (2009) Get Aroused and Be Stronger: Emotional Facilitation of Physical Effort in the Human Brain. *J Neurosci* 29: 9450–9457. doi:10.1523/JNEUROSCI.1951-09.2009.
40. Schmidt L, Lebreton M, Cléry-Melin M-L, Daunizeau J, Pessiglione M (2012) Neural mechanisms underlying motivation of mental versus physical effort. *PLoS Biol* 10: e1001266. doi:10.1371/journal.pbio.1001266.
41. Borg GA (1982) Psychophysical bases of perceived exertion. *Med Sci Sports Exerc* 14: 377–381.
42. Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007) Variational free energy and the Laplace approximation. *NeuroImage* 34: 220–234. doi:10.1016/j.neuroimage.2006.08.035.
43. Daunizeau J, Friston KJ, Kiebel SJ (2009) Variational Bayesian identification and prediction of stochastic nonlinear dynamic causal models. *Physica D* 238: 2089–2118. doi:10.1016/j.physd.2009.08.002.
44. Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46: 1004–1017. doi:10.1016/j.neuroimage.2009.03.025.
45. Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, et al. (2010) Comparing families of dynamic causal models. *PLoS Comput Biol* 6: e1000709. doi:10.1371/journal.pcbi.1000709.

## Methods

### Participants

The study was approved by the Pitié-Salpêtrière Hospital ethics committee. All subjects were recruited via email within an academic database and gave informed consent prior to participating in the study. There was no restriction of handedness, excepted for the original (Implicit) Effort Allocation Task, in which participants were all right handed for neuroimaging purposes. Other inclusion criteria were: age between 20 and 39 years, absence of self-reported psychiatric or neurological history and of current psycho-active substance consumption.

In all studies, participants were told that they would win the money accumulated during the task. In the previous study (Implicit task), the payoff was eventually rounded up to a fixed amount (100€) credited by bank transfer. In all new studies participants were paid in cash at the end of the experiment. The payoff was partitioned into a fixed amount and variable amount depending on the money won during the task. For the Cost Rating Task, the amount earned during the task was eventually down-scaled (divided by 2.48) to fit in a budget of 30€ per subject while maintaining the correspondence between payoff and incentive during the task. Participants were informed about this correction prior to the experiment.

The original implicit task was performed in a MRI scanner for half the subjects and under a MEG helmet for the other half. One subject in the MRI group was excluded from the analysis because of calibration issues. For the Adaptation Task 3, one participant was excluded because of calibration issues and another for cheating (repeated, direct manipulation of the air tube). For the Dissociation Task, one participant was excluded due to an instruction issue: she could not understand the meaning of the percentage displayed on the screen, which indicated the difficulty level in proportion of the maximal force. Two other participants were excluded due to calibration issues. The task-specific information is summarized in Table 1.

### Experimental set up

We used homemade power grips composed of two plastic or wood cylinders compressing an air tube when squeezed. The tube was connected to a transducer converting air pressure into voltage. Thus, grip compression resulted in the generation of a differential voltage signal, linearly proportional to

the force exerted. The signal was amplified and digitized by a signal conditioner (CED 1401, Cambridge electronic design, UK) for Implicit, Explicit and Dissociation tasks, and by a homemade device for the Adaptation Tasks and Cost Rating task. The digitized signal was read by a Matlab program (The MathWorks Inc., USA).

## **Pre-processing of force data**

In the Adaptation Tasks (1 to 3), the effort onsets were identified on-line and used to update the screen displayed to the participants. The effort onset was determined as the first sample exceeding 20% of the participant maximal force.

In the Effort Allocation Tasks ('implicit', 'explicit' and 'dissociation'), effort onsets and offsets were identified off-line with an algorithm using the same two criteria for all conditions: 1) force temporal derivative higher than one standard deviation and 2) force level lower (for effort onset) or higher (for effort offset) than half the maximal force. The first rest period started with coin presentation and the subsequent effort and rest periods were defined by force onsets and offsets.

## **Maximal force estimate**

For all tasks, we measured the maximal force for each hand before starting task performance, following published guidelines [1]. Participants were verbally encouraged to squeeze continuously as hard as they could, until a growing line displayed on a computer screen reached a target. The growing rate was proportional to the force produced to motivate subject to squeeze hard. Maximal force was set to the average of data points over the last half of the squeezing period exceeding the median. Then subjects were provided a real-time feedback about the force produced on the handgrip, which appeared as a fluid level moving up and down within a thermometer, the maximal force being indicated as a horizontal bar at the top. Subjects were asked to try outreaching the bar and state whether it truly corresponded to their maximal force. If not, the calibration procedure was repeated.

The procedure was slightly simplified for the Adaptation Tasks and Cost Rating Task: 1) the rate of the growing bar was held constant and not indexed on the participants exerted force level, 2) the duration during which the participants had to squeeze as hard as they could was fixed to 5s and 3) all data points were used for the estimate (and not the last half of recorded levels).

## Behavioral tasks

All tasks were presented on a computer screen, and were programmed with Matlab using Cogent 2000 (Wellcome Department of Imaging Neuroscience, London, UK) for the Implicit and Explicit Tasks, and Psychtoolbox (<http://psychtoolbox.org>) for the Dissociation Task, Adaptation Tasks, and Cost Rating Task.

### Cost Rating Task

The task included 7 sessions, using right and left hands alternatively. Each session comprised 21 trials. The design was fully factorial, crossing all factor levels: 3 incentive levels (10c, 20c, 50c), 7 duration levels (equally spaced from 3 to 7s) and 7 difficulty levels (equally spaced from 40 to 60% of the participant maximal force). Each cell was presented only once, as there were 147 cells for 147 trials. The order of presentation was pseudo-randomized such that the different sessions had exactly the same incentive and difficulty level on average, and little variation in mean duration.

Every trial started with baseline (1s), followed by incentive display (1s) and then by the appearance of a thermometer that served as a 'go' signal to trigger effort exertion. The fluid level within the thermometer provided online feedback on the force being exerted, with scaling adjusted such that the target bar corresponded to difficulty level (40% to 60% of the maximal force). The thermometer was displayed as long as the participant had to sustain the effort. The imposed duration was applied starting when the target force, was reached and not when the thermometer appeared. Exhaustion rating was done just following the effort. 'Avez-vous épuisé vos ressources?' ('How much are you exhausted now?') was written on screen, and participants indicated their rating from 'Pas du tout' ('not at all') to 'Totalement' ('completely') with a cursor that could be moved with left/right key press. We framed the question in terms of exhaustion instead of the perceived exertion [41] because the rating occurred after (not during) the effort. The rating scale had 50 steps but no visible graduation. Rating and validation (by pressing the space bar) were self-paced. The last screen lasted 1.5s and summarized the payoff earned in the current trial and the cumulated payoff over all preceding trials. The amount earned during a trial was calculated as the incentive value multiplied by the proportion of the imposed duration spent above the target force level.

### Adaptation Tasks

The display was quite similar in all adaptation tasks, which included a total of 8 sessions. The exerted force level was always displayed as a fluid moving up in a thermometer, the target bar on the top of the thermometer indicating 60% of the participant maximal force. All trials included a first effort (with imposed duration), a rest period, and a second effort (with free duration). The payoff was proportional (with a fixed rate) to the time spent above the target force level during the second effort. The color of the fluid in the thermometer instructed what to do: red for the first effort, blue for rest (with 'STOP!' displayed above the thermometer), green for the second effort, which participants initiated either immediately (in Tasks 1 and 2), or at their convenience (in Task 3). When participants stopped squeezing, more precisely at the first force sample under 50% of their maximal force, the color turned to white, instructing that they should rest until the following trial. In all Tasks, 'PLUS FORT' (meaning 'harder') was displayed above the thermometer during the imposed effort when the force being exerted was under the target level (60%). For Tasks 1 and 2, the color turned to white and the message 'VOUS AVEZ APPUYE TROP TARD' (meaning 'you squeezed too late') was displayed if the participant initiated the trial too late (more than 1s after the color change). In all three tasks, a flickering dollar symbol was displayed when the force was above the target level, during the second effort whose duration was free (green color), to indicate that money was being accumulated. Both the trial payoff and the cumulated payoff were displayed on screen at the end of each trial.

#### **Adaptation Task 1 (variable effort / constant rest / free effort)**

Each trial presented the following events successively: imposed effort (at 60% of maximal force), imposed rest (2s), go signal to initiate an effort of free duration (20s allowed), feedback (2s), inter-trial interval (2s). Imposed effort durations were drawn from a set of 36 points regularly spaced between 1 and 10s. The same 36 durations, divided into 4 sessions of 9 trials each, were presented once to the left hand and once to the right hand in the same randomized order. For session  $S_i$ , effort durations were picked up every 4 points in the randomized sequence of 36 values, starting at a sample  $i$  randomly drawn (without replacement) between 1 and 4. This procedure ensures that over subjects, all sessions have the same average effort duration.

#### **Adaptation Task 2 (constant effort, variable rest, free effort)**

Each trial presented the following events successively: imposed effort (7s at 60% of maximal force), imposed rest (between 1 and 10s), go signal to initiate an effort of free duration (20s allowed), feedback (2s), and no inter-trial interval. Imposed rest durations were defined so as to sample small

durations more than long durations. We simulated a mixture of Gaussians (10000 points), with 75% of points drawn from  $N(3, 2)$ , and 25% drawn from  $N(10, 2)$ , where  $N(m, \sigma)$  denotes a Gaussian distribution with mean  $m$  and standard deviation  $\sigma$ . This distribution was cut off to retain values higher than 1s and divided into 37 bins. The first 36 bins were then retained for our sampling rest durations. The same 36 durations were presented to the left and right hands in the same randomized order, using the same randomization technique as was implemented for Task 1.

### **Adaptation Task 3 (variable effort, free rest and free effort)**

Each trial presented the following events successively: imposed effort (at 60% of maximal force), imposed rest (2s), a signal indicating to the participant that the second effort can be initiated (20s allowed for in total for first resting and then exerting effort), feedback (2s), inter-trial interval (2s). Imposed effort durations were 36 points equally spaced between 1s and 10s. The same 36 durations were presented to the left and right hands in the same randomized order, using the same randomization technique as was implemented for Task 1.

### **Effort Allocation Tasks (Implicit, Explicit and Dissociation Tasks)**

The Implicit task is described in [10]; we reproduce here the relevant details. Participants performed 8 sessions of 9 trials corresponding to the 9 cells of the factorial design (3 incentive x 3 difficulty conditions), which were presented in a random order. Subjects performed 8 sessions in total, switching hands as instructed between sessions to avoid muscle exhaustion. After baseline measurement of the pressure at rest, each trial started by revealing the monetary incentive with a coin image (10, 20 or 50 cents) displayed for 1s. Then subjects had 30s to win as much money as possible. They knew that the payoff was proportional to both the incentive and the time spent above the target bar, which was always placed at the same height in the thermometer. The force needed to reach the bar (70, 80 or 90% of subject's maximal force), i.e. trial difficulty, was not indicated to subjects. Subjects only knew that task difficulty would vary across trials. They were provided with online feedback on both the exerted force (with a fluid level moving up and down within a thermometer) and the trial-wise cumulated payoff (with a counter displayed above the thermometer). Each trial ended with a 2s display of the session-wise cumulated payoff.

The only change from the Implicit to the Explicit Task is that the difficulty level was displayed on the right and left of the coin image, as percentages of maximal force (70%, 80% or 90%).

In the Dissociation Task, monetary incentive (10c or 20c), actual difficulty (75% or 85%) and cued difficulty (75% or 85%) were combined into a factorial design comprising 8 cells. Cued difficulty

level was indicated on the screen as in the Explicit task but was congruent with the actual difficulty level (actual force needed to reach the target bar) in half the trials only. The experiment was divided into 8 sessions presenting one trial for each of the 8 cells in a random order. The randomization avoided to present identical pair of cues (for incentive and difficulty levels) in two consecutive trials. Apart from the potential mismatch between the cued and actual difficulty levels, the trial settings were identical to those of the Explicit Task.

## Statistical analysis

### Cost Rating Task vs. Implicit Task

We first submitted the ratings obtained from the Cost Rating Task to a multiple regression analysis, so as to estimate the effect of several factors. The regressors comprised the manipulated factors (incentive, difficulty and duration level) and covariates (a constant per session to capture the mean, a linear trend per session to capture drift over trials, and the initial position of the rating cursor). Two-way interaction terms were also included. Regressors were z-scored over all trials, except trends that were z-scored within their sessions and padded with 0, and constant regressors. The significance of the parameter estimates was assessed with a random-effect analysis at the group level using a two-tailed t-test.

To compare with the Effort Allocation (implicit) Task, we modeled cost evidence as follows:  $C = \beta_0 + \beta_1 T^{\lambda_1} + \beta_2 D^{\lambda_2} + \beta_3 T^{\lambda_1} D^{\lambda_2}$ , where D and T are difficulty and time (duration). To make the estimation of D and T betas independent of their unit (force versus time), they were divided by their mean value. Setting the  $\lambda$  terms to 1 made the model linear with respect to experimental factors, setting  $\beta_3$  to zero made the model purely additive and setting  $\beta_1$  and  $\beta_2$  to zero made the model purely interactive. All these possibilities (including non-linearities, including additive terms, including interaction) were combined, resulting in a total of 8 models. Note that formally, the linear and non-linear constant models ( $C=\beta_0$ ) are strictly equivalent. In the Cost Rating Task, the dependent variable was subjective cost report (fatigue sensation), which could be directly regressed against the cost evidence model. In the Effort Allocation (implicit) Task, costs had to be inferred from the behavior. The probability to stop the effort after a given exertion duration was derived from the cumulated distribution of effort duration, for each difficulty levels. This probability was regressed against a sigmoid (or probit) function of the modeled cost evidence:  $P = \frac{1}{1 + e^{-C}}$ . This sigmoid was not parameterized (i.e., C was not transformed with scaling and offset parameters), since this would

be redundant with the beta parameters included in the definition of C itself. Apart from the sigmoid transformation, the same procedure was thus applied to the Cost Rating and Effort Allocation Tasks.

A constant elasticity of substitution (CES) model was also fitted to characterize the curvature of cost evidence. The CES model is  $C = (\alpha D^\delta + (1 - \alpha) T^\delta)^{\frac{1}{\delta}}$ , in which  $\alpha$  ranges from 0 to 1 and characterizes the equivalence between D and T (or substitution ratio), and  $\delta$  is strictly positive and characterizes the curvature of this equivalence. We introduced an offset and a scaling factor to the CES model as two additional free parameters, which were independent from the estimation of alpha and delta since D and T had a mean of one. Following the same procedure as in the model comparison, cost evidence was fitted onto subjective cost in the Cost Rating Task, and passed through a probit function to be fitted onto stop probability in the Effort Allocation task.

### **Adaptation Tasks**

Effort and rest durations were analyzed using multiple linear regressions. For Task 1, the dependent variable (second effort duration) was fitted with four regressors: first effort duration, session number, session-wise cumulated effort, and the residual effort initiation delay (i.e., delay between go signal and effort onset, after removing the variance explained by the three other regressors). As before, significance of parameter estimates was assessed using a two-tailed t-test at the group level. For Task 2, the same four regressors were used to explain the dependent variable (second effort duration), except that the manipulated factor was rest duration (not first effort duration). For Task 3, the two dependent variables (rest effort and second effort durations) were fitted with the same linear model as was done for Task 1, except that there was no initiation delay.

We also analyzed the relationship between imposed and observed durations in Tasks 1 and 2, by fitting linear and saturation models, which we compared using Bayesian model selection (BMS). The linear model was  $T_1 = \beta_1 + \beta_2 T_2$ . We tested two models for saturation: a bounded linear model ( $T_1 = \beta_1 + \beta_2 T_2$  if  $T_2 < \gamma$ , and  $T_1 = \beta_1 + \beta_2 \gamma$  otherwise) and a model with exponential saturation ( $T_1 = (\gamma - \beta_1) * (1 - \exp(-\beta_2 T_2)) + \beta_1$ , in which  $\beta_1$  is the intercept,  $\beta_2$  the increase rate and  $\gamma$  the asymptote). The BMS procedure is described in the next section. In principle, the cost-evidence model predicts that recovery during rest should be bounded, such that after a certain time, more rest does not increase the duration of the subsequent effort. For simplicity we assumed a linear dynamics for accumulation and dissipation, which implies that the saturation should manifest as a linear increase followed by a constant plateau. However, we reasoned that white noise in data generation should render this function closer to an exponential saturation. We confirmed this intuition with a simulation,

proceeding as follows. 1) We fitted a bounded linear model to individual data and retained the median parameter estimates and the residuals standard deviation for each subject. 2) With these parameter estimates and the durations manipulated experimentally, we simulated 100 sets of noisy effort data per subject, using a white noise of the same magnitude as residuals standard deviation. 3) We fitted both the bounded linear model and the exponential model to these individually simulated data, and calculated their respective log-evidence. 4) We averaged these log-evidences over simulations and performed a group-level BMS, which favored the exponential model with high confidence (chance level is  $\frac{1}{2}$ , expected frequency  $ef=0.75$ , exceedance probability  $xp=0.97$ ). We noticed, as expected, that this exceedance probability decreased when decreasing the noise magnitude. Therefore, we included the exponential model in the BMS performed on our observed data, since it was more likely to capture the saturation effect at the observed noise magnitude.

### **Effort Allocation Tasks (Implicit, Explicit and Dissociation Tasks)**

Effort and rest durations were submitted to multiple regression analysis. The regressors comprised the manipulated factors (incentive and difficulty levels for the Implicit and Explicit Task; incentive cued and actual difficulty levels for the Dissociation Task), temporal factors (the session number, the trial position within a session and the effort or rest position within a trial), and interaction terms (the two-way interactions between the manipulated factors and the temporal factors, and the two-way interactions between the manipulated factors, which was extended to a third-way interaction between the three manipulated factors in the Dissociation task). All the regressors were z-scored to provide standardized effect size.

The significance of parameter estimates (regression coefficients) was assessed with a random-effect analysis at the group level using a two-tailed t-test. Dissociation between cued and actual difficulty in the Dissociation Task was estimated using a two-tailed paired t-test on the parameter estimates. For non-parametric t-tests, we estimated the null t-distribution using all possible permutations ( $n=2^{15}$ ) between the 'cued' and 'actual' labels, and estimated the probability of t-values more extreme than observed (bilateral test).

### **Bayesian model selection**

To perform model selection, models were first estimated for each subject using a variational Bayes approach under the Laplace approximation [42,43], using a toolbox by Jean Daunizeau (available at <http://sites.google.com/site/jeandaunizeauswebsite/>). This algorithm not only estimates linear and non-linear models but also calculates their evidence based on a free-energy approximation [42]. The

evidence of a model is the probability of observing the data given this model. This probability corresponds to the marginal likelihood, which is the integral over the parameter space of the likelihood of the parameterized model weighted by the prior on its parameters. This probability increases with the likelihood (which is the accuracy of the fit) and is penalized by the integration over the parameter space (which is the complexity of the model). The model evidence thus represents a trade-off between accuracy and complexity and can guide model selection [44]. Model selection was performed with a group-level random-effect analysis of the log-evidences obtained for each model and subject, using Gibbs sampling in SPM8 (Statistical Parametric mapping, Wellcome Department of Imaging Neuroscience, London, UK) [44]. This procedure estimates the expected frequency (denoted  $ef$ ) and the exceedance probability (denoted  $xp$ ) for each model within a set of models, given the data gathered from all subjects. Expected frequency quantifies the posterior probability, i.e. the probability that the model generated the data for any randomly selected subject. This quantity must be compared to chance level (one over the number of models or families in the search space). Exceedance probability quantifies the belief that the model is more likely than all the other models of the set, or in other words, the confidence in the model having the highest expected frequency [44]. Family-level inference was conducted similarly to model-level inference after defining a partition within the model space as described in [45] and implemented in SPM8.

### **Computational models (Implicit, Explicit and Dissociation tasks)**

We first defined a class of models that can a priori produce the results that we intended to explain. These models were then submitted to a Bayesian model selection in order to identify the most plausible model among all the possible models. The model space was defined by simplifying a full model, starting with the Implicit Task. The model is based on accumulation-dissipation processes: cost evidence ramps up during effort to a bound that triggers effort cessation, and ramps down during rest to a bound that triggers effort resumption. As for simplicity the fluctuations were modeled as linear, the effort and rest durations ( $Te$  and  $Tr$ ) are just the ratios between the amplitude  $A$  (distance between bounds) and the accumulation or dissipation slope ( $Se$  and  $Sr$ ). In the full model, the free parameters  $A$ ,  $Se$ ,  $Sr$  can vary across trials around their mean values ( $\alpha_m$ ,  $\beta_m$ ,  $\gamma_m$ ), depending linearly on experimental factors: in this case the incentive  $I$  and the difficulty  $D$ . The full model is thus:

$$Te = \frac{A}{Se}; \quad Tr = \frac{A}{Sr}$$

$$\begin{cases} A = \alpha_m + \alpha_I I + \alpha_D D \\ Se = \beta_m + \beta_I I + \beta_D D \\ Sr = \gamma_m + \gamma_I I + \gamma_D D \end{cases}$$

Simpler models can be designed by setting one or more weights to zero. As there are 6 weights, all combinations give a total of  $2^6 = 64$  models. However, some of these models are not worth considering as they cannot account for the effect that we want to explain. The most extreme case is when all weights are null: such a model cannot produce any of the effect of incentive and difficulty that we observed in the data. After discarding all models with which at least one of the significant results reported in Figure 3 could not be produced, the search space was restricted to 24 models (see red lines in Table 2A). Note that predicting an effect that was not significant in our data was not a criterion for rejection. The same approach was applied for defining linear models of the Explicit Task, leading to a search space of 16 models. For the Dissociation Task, another modulator was included as there were two types of difficulty (cued and actual). The full model has therefore 9 weights, which gives  $2^9=512$  possible models, which were reduced to 144 models after rejection of irrelevant models.

$$Te = \frac{A}{Se}; \quad Tr = \frac{A}{Sr}$$

$$\begin{cases} A = \alpha_m + \alpha_I I + \alpha_{Da} Da + \alpha_{Dc} Dc \\ Se = \beta_m + \beta_I I + \beta_{Da} Da + \beta_{Dc} Dc \\ Sr = \gamma_m + \gamma_I I + \gamma_{Da} Da + \gamma_{Dc} Dc \end{cases}$$

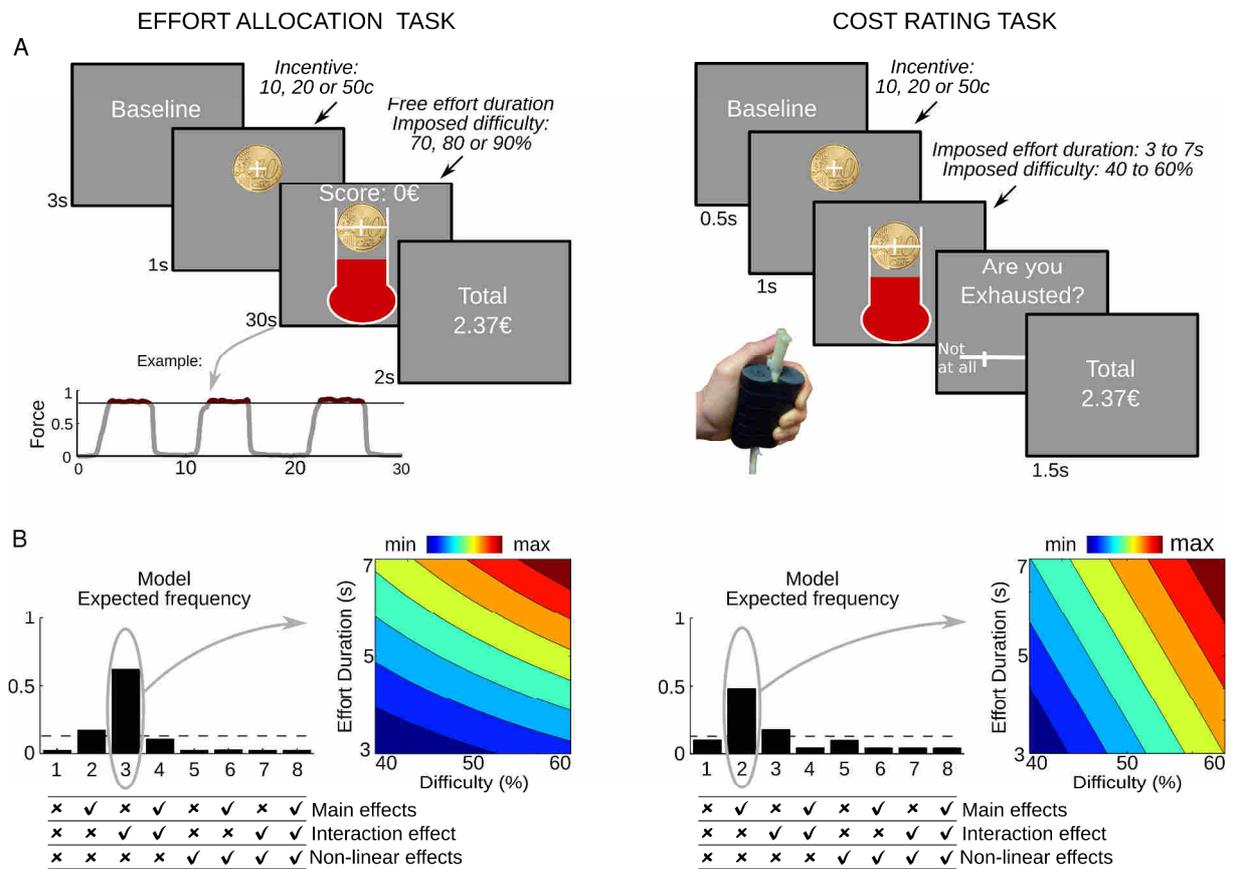
We also tested a class of hyperbolic models for the Explicit and Dissociation Tasks. As opposed to the linear formulation, the discount of incentive by cued difficulty was assumed to be hyperbolic, as in standard economic models. The full hyperbolic model is:

$$Te = \frac{A}{Se}; \quad Tr = \frac{A}{Sr}$$

$$\begin{cases} A = \alpha_m + \frac{\alpha_I I}{1 + \alpha_{Dc} Dc} + \alpha_{Da} Da \\ Se = \beta_m + \frac{\beta_I I}{1 + \beta_{Dc} Dc} + \beta_{Da} Da \\ Sr = \gamma_m + \frac{\gamma_I I}{1 + \gamma_{Dc} Dc} + \gamma_{Da} Da \end{cases}$$

The  $D$  term that denoted difficulty in first linear model has been decomposed into  $Da$  and  $Dc$ , denoting actual and cued difficulty in the Dissociation Task. Note that for in the Explicit Task,  $De$  and  $Dc$  have exactly the same values. The model can nonetheless be estimated unambiguously in this task since the effect of  $Da$  is linear whereas that of  $Dc$  is hyperbolic. Also note that with hyperbolic formulation, there are dependencies between weights since a null numerator prevents the denominator from impacting the model fit. Thus we discarded models with a null numerator and a non-null weight at the denominator (this is shown with shading in Table 2C). After discarding the models that were not able to produce all the significant results shown in Figure 3, the search space was eventually restricted to 78 models for the Explicit and Dissociation Tasks.

## Figure Legends

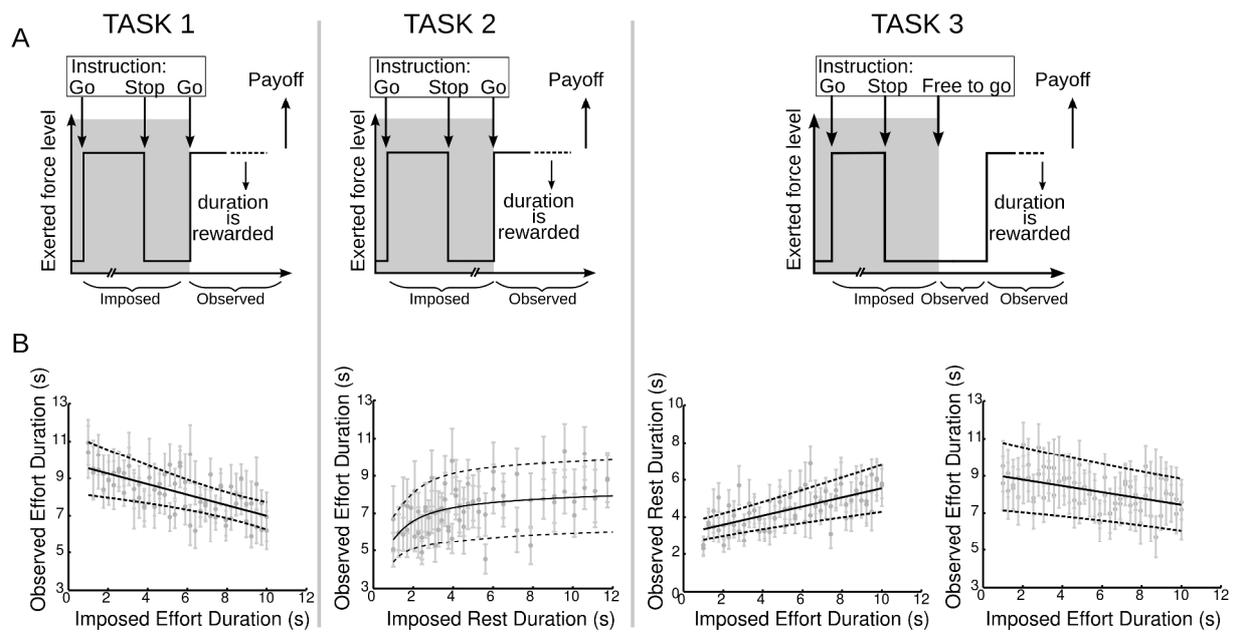


**Figure 1: Introspection of cost evidence**

A: Behavioral tasks. The illustrated screenshots were successively presented every trial. Left: the Effort Allocation Task was exploited in a previous study [10]. When the thermometer image was displayed, participants could squeeze a handgrip to win as much money as possible. Subjects were provided with online feedback on force level and cumulative payoff. The payoff was only increased when force level was above the target bar, at a constant rate proportional to the monetary incentive. The incentive (10, 20 or 50 cents) and the difficulty (i.e. the force required to reach the target bar: 70, 80 or 90% of maximal force) were crossed over trials. The last screen indicated the money won over all preceding trials. Right: the Cost Rating Task was developed to assess introspection of fatigue levels. On each trial, participants were asked to squeeze the hand grip up to the target level (horizontal bar), which corresponded to varying difficulty level (40% to 60% of maximal force), as long as the thermometer was displayed, which could last for varying durations (4 to 7 seconds). After

each effort, participants rated their degree of exhaustion using a visual horizontal analog scale. The last screen of each trial indicated the payoff cumulated over preceding trials.

B: Computational modeling. Bar graph: Bayesian model comparison. For each participant, we estimated several models generating cost evidence from difficulty and duration. Cost evidence was then used to fit the decisions to stop effort exertion (left) or the subjective ratings of exhaustion (right). Models were linear combinations of different possible regressors (main effects, interaction and non-linear effects ), as indicated in the bottom chart (tick: included, cross: not included). Color plot: model simulation. Predictions of winning models (with the highest expected frequency) were calculated using the group-level median value of free parameters and illustrated over the same range of durations and difficulties for visual comparison.



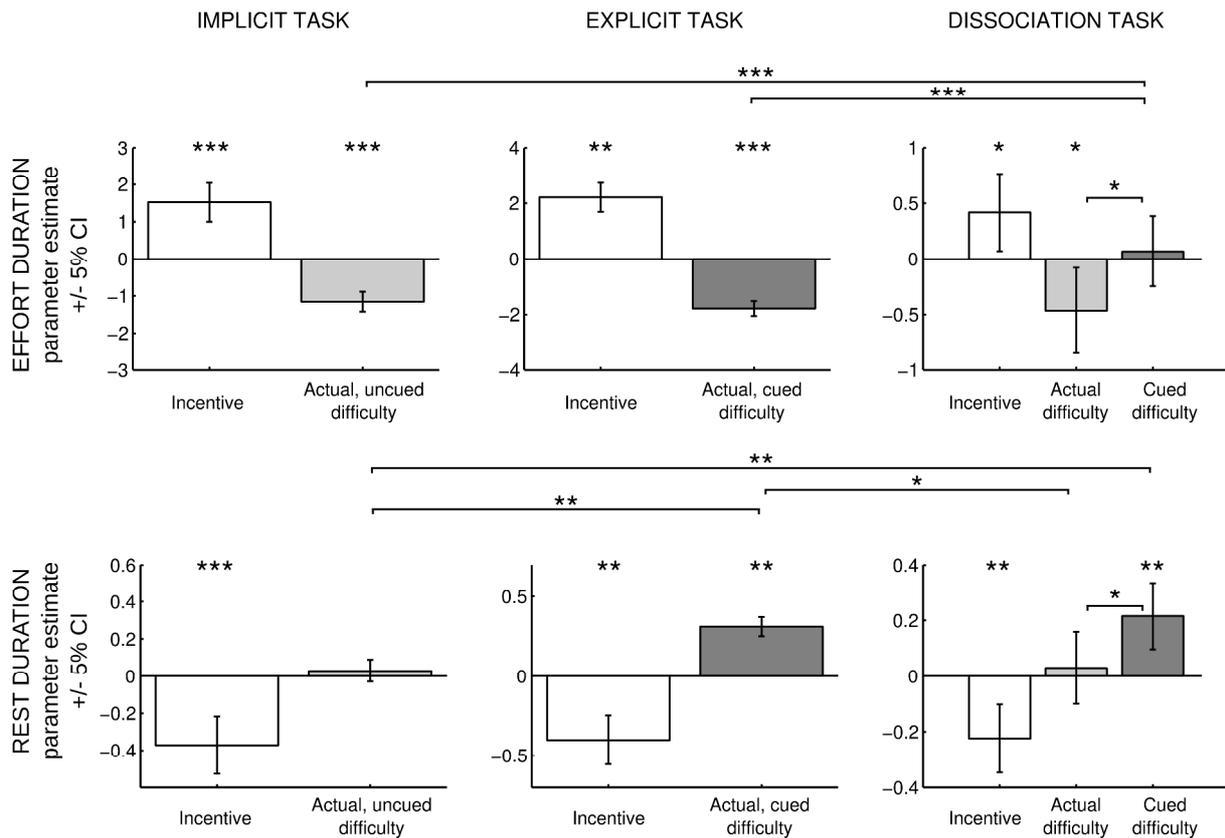
**Figure 2: Behavioral adaptation to cost evidence**

The three columns present three different studies, with results underneath the tasks. Note that there are two sub-columns for the last study (on the right) because there are two dependent variables (rest and effort duration).

A. Behavioral tasks. Each plot sketches the variations of exerted force level within a trial. Gray shading indicates the periods when action was imposed to participants; in the other periods the behavior was let free. The broken line points to periods when durations were systematically varied.

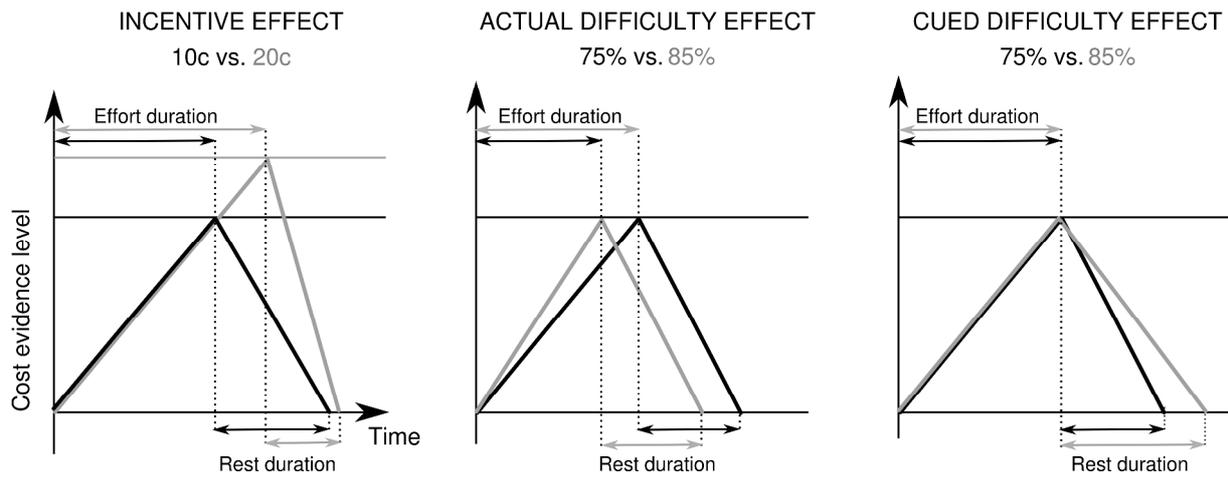
B. Relation between imposed and observed group-level average durations (+/- s.e.m). Two data points are plotted for each imposed duration, corresponding to left and right hands. The black line is the group average of the model fit estimated at the subject level; dash lines size the 5% confidence interval of the average.

C. Significant factors affecting each observed duration. Bars represent the effect sizes (+/- 5% confidence intervals) of the different factors.<sup>5</sup> The variance due to cumulated effort and session number was regressed out of initiation duration prior to estimation. P-value (bilateral test): \*<0.05, \*\*<0.005, \*\*\*<0.0005.



**Figure 3: Dissociation of implicit from explicit cost processing**

Three sets of participants performed three slightly different versions of the Effort Allocation Task. The Implicit Task is sketched in Figure 1A. The only variation introduced in the Explicit Task is that effort difficulty was written on screen (70%, 80%, 90%) along with incentive level, announced as a coin image. The Dissociation Task was visually identical to the Explicit Task, the difference being that the difficulty level announced on screen was not predictive of the actual difficulty level: cued and actual difficulties were crossed into a factorial design. Thus, the factorial combination generated 9 cells for the Implicit and Explicit Tasks (3 incentives x 3 difficulties) and 8 cells for the Dissociation Task (2 incentives x 2 actual x 2 cued difficulties). The effects of the different factors, estimated with linear regression analysis, are illustrated column-wise for each task. Regression coefficients were statistically tested and compared with bilateral t-test, p-values: \*<0.05, \*\*<0.005, \*\*\*<0.0005.



**Figure 4: Computational effects of experimental factors**

The diagrams illustrate how the experimental factors (monetary incentive, actual difficulty, expected difficulty) affect cost-evidence monitoring. Bayesian model comparison dissociated the computational effects of experimental factors: higher incentives increase the amplitude between bounds and the dissipation rate; higher actual difficulty steepens the accumulation rate; higher expected difficulty shallows the dissipation rate.

## Tables

**Table 1 Task-specific information on participants.**

'Exp.' refers to the author who collected the data. The payoff was fractioned into a fix amount ('Fix') for participation and a variable amount ('Var') depending on performance.

Task	Exp.	Period	N after exclusion	N male	N excluded	Mean age +/- s.e.m.	Fix (€)	Var (€)	Var (€) range
Cost Rating Task	FM	02/2013	18	7	0	22.2 +/- 0.5	0	29.8	29 – 30
Adaptation Task 1	FM	03/2012	12	2	0	22.7 +/- 0.8	10	10.1	7 - 13
Adaptation Task 2	FM	03/2012	12	0	0	21.9 +/- 0.4	10	9.6	4 – 15
Adaptation Task 3	FM	03/2012	12	4	2	21.7 +/- 0.7	10	10.3	6 - 15
Implicit Task	FM	04-05/2010	38	16	1	24.2 +/- 0.65	50	31.6	15 – 48
Explicit Task	FM	03/2011	14	10	0	23.7 +/- 0.4	15	13	8 – 19
Dissociation Task	LS	10/2011	15	5	3	25.4 +/- 0.8	10	6	3 – 10

**Table 2: Model space definition.**

Models are characterized by the modulation of parameters (A: amplitude between bounds,  $S_E$ : accumulation slope during effort, and  $S_R$ : dissipation slope during rest) by the experimental factors (incentive, actual difficulty, cued difficulty). Each line is a possible model and each column a potential factor ('1' denotes that the modulation is allowed, '0' that the modulation is absent). The different modulations are linearly combined, except for the cued difficulty in Table 1D, which is integrated as a hyperbolic discounter of monetary incentives. In this case, including cued difficulty (in the denominator) is useless when modulation by incentives (the numerator) is not allowed, which is indicated by gray shadings. Lines appearing in red correspond to models that were not included in the comparison, because they cannot produce the behavioral results (significant effect on effort or rest duration). The winning models for the different tasks appear in blue (note that it is indeed the same model, as there is no cued difficulty, and hence no hyperbolic discounting, in the Implicit Task).

**Table 2A: Implicit task, model space.**

Difficulty			Incentive		
A	$S_E$	$S_R$	A	$S_E$	$S_R$
1	1	1	1	1	1
1	1	0	1	1	0
1	0	1	1	0	1
1	0	0	1	0	0
0	1	1	0	1	1
0	1	0	0	1	0
0	0	1	0	0	1
0	0	0	0	0	0

**Table 2B: Explicit task, model space for linear models**

Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1
1	1	0	1	1	0
1	0	1	1	0	1
1	0	0	1	0	0
0	1	1	0	1	1
0	1	0	0	1	0
0	0	1	0	0	1
0	0	0	0	0	0

**Table 2C: Dissociation task, model space for linear models**

Actual Difficulty			Cued Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1	1	1	1
1	1	0	1	1	0	1	1	0
1	0	1	1	0	1	1	0	1
1	0	0	1	0	0	1	0	0
0	1	1	0	1	1	0	1	1
0	1	0	0	1	0	0	1	0
0	0	1	0	0	1	0	0	1
0	0	0	0	0	0	0	0	0

**Table 2D: Dissociation & explicit task, model space for hyperbolic models**

Actual Difficulty			Cued Difficulty			Incentive		
A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>	A	S <sub>E</sub>	S <sub>R</sub>
1	1	1	1	1	1	1	1	1
1	1	0	1	1	0	1	1	0
1	0	1	1	0	1	1	0	1
1	0	0	1	0	0	1	0	0
0	1	1	0	1	1	0	1	1
0	1	0	0	1	0	0	1	0
0	0	1	0	0	1	0	0	1
0	0	0	0	0	0	0	0	0

## **How the human brain allocates physical effort over time: evidence from behavior, neuroimaging and pharmacology**

Faire le bon choix, c'est trouver le bon compromis entre coût et bénéfice. Dans le cas de la gestion de l'effort physique, ce compromis prend une dimension temporelle. Pour comprendre comment la décision d'arrêter ou reprendre l'effort est prise, nous avons développé un paradigme expérimental chez le sujet humain sain et un modèle computationnel dans lequel le coût estimé augmente à l'effort car la fatigue affecte toute la commande motrice et diminue au repos quand nous récupérons. Le comportement reflète les variations de ce coût estimé et du compromis avec le bénéfice attendu. Grâce à la complémentarité de l'imagerie fonctionnelle par résonance magnétique et de la magnétoencéphalographie (MEG), le coût estimé a été localisé dans les régions proprioceptives du cerveau : l'insula postérieure et le thalamus ventromédian. La MEG a également révélé que la désynchronisation du rythme beta moteur (13-30Hz) permet une reprise plus rapide de l'effort quand les enjeux sont importants. Cette gestion stratégique du repos est liée à l'utilité attendue qui peut être dissociée de l'utilité réelle. Nos résultats montrent que la gestion de l'effort est adaptée en ligne au coût estimé et modulée stratégiquement en fonction des coûts et bénéfices attendus. Les antalgiques (hypnose ou paracétamol) ont un effet limité sur ce processus, à l'inverse de la sérotonine (Escitalopram). Notre contribution, à l'interface entre médecine du sport, théorie de la décision et modèle d'accumulation utilisés en neurosciences, propose un mécanisme pour optimiser la gestion de l'effort physique en maximisant les gains et minimisant les dommages corporels.

Mots clés : effort, accumulation, imagerie fonctionnelle par résonance magnétique, magnétoencéphalographie, prise de décision, motivation

No pain, no gain: optimal decisions involve a tradeoff between cost and benefit. We propose that in physical effort allocation, this tradeoff is unfolded over time. We present a task to investigate this process in the laboratory with healthy humans and we suggest a computational model to account for decisions to stop and resume the effort. Costs increase during exertion, due to fatigue at all stages of the motor command and decrease during rest, due to recovery. We show that this dynamic may be captured by a cost-evidence variable and compared to the expected benefit. Functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG) complementarily showed that cost-evidence may be implemented in proprioceptive regions of the brain: posterior insula and ventromedial thalamus. In addition, MEG showed that motor beta (13-30 Hz) desynchronization mediates the effect of incentives to hasten effort resumption. This strategic invigoration of rest is supported by a behavioral dissociation: the expected utility (not the actual utility) modulates rest durations. Together, our results support that the behavior is adapted on the fly to cost-evidence levels and that this mechanism is modulated strategically according to the expected cost and benefit. This behavior was not affected by pain killers (hypnosis or paracetamol), but by serotonin (Escitalopram). This work bridges a gap between sport medicine, value-based decision-making and accumulation models in neuroscience in showing that accumulation and dissipation of cost-evidence can guide the optimization of effort allocation: this mechanism implements the maximization of benefit while the body costs are minimized.

Key words: effort, accumulation, functional magnetic resonance imaging, magnetoencephalography, decision making, motivation