Article

Current Biology

A neuro-metabolic account of why daylong cognitive work alters the control of economic decisions

Highlights

- Cognitive fatigue is explored with magnetic resonance spectroscopy during a workday
- Hard cognitive work leads to glutamate accumulation in the lateral prefrontal cortex
- The need for glutamate regulation reduces the control exerted over decision-making
- Reduced control favors the choice of low-effort actions with short-term rewards

Authors

Antonius Wiehler, Francesca Branzoli, Isaac Adanyeguh, Fanny Mochel, Mathias Pessiglione

Correspondence

antonius.wiehler@gmail.com (A.W.), mathias.pessiglione@gmail.com (M.P.)

In brief

Despite more than a century of research, the origins of cognitive fatigue have remained elusive. Here, Wiehler et al. suggest a neuro-metabolic account: daylong cognitive work leads to glutamate accumulation in the lateral prefrontal cortex, which reduces control over decisions and thereby favors effortless behaviors with immediate gratifications.

Wiehler et al., 2022, Current Biology *32*, 3564–3575
August 22, 2022
https://doi.org/10.1016/j.cub.2022.07.010

CellPress

CellPress

Current Biology

Article

A neuro-metabolic account of why daylong cognitive work alters the control of economic decisions

Antonius Wiehler,[1,2,3,4,6,7,3,*] Francesca Branzoli,[2,3] Isaac Adanyeguh,[3] Fanny Mochel,[3,5] and Mathias Pessiglione[1,2,3,4,*]
[1]Motivation, Brain and Behavior Lab, Paris Brain Institute (ICM), Pitié-Salpêtrière Hospital, Paris, France
[2]Center for NeuroImaging Research (CENIR), Paris Brain Institute (ICM), Pitié-Salpêtrière Hospital, Paris, France
[3]Sorbonne Universités, Inserm U1127, CNRS U7225, Paris, France
[4]Department of Psychiatry, Service Hospitalo-Universitaire, Groupe Hospitalier Universitaire Paris Psychiatrie & Neurosciences, Paris, France
[5]Assistance Publique - hôpitaux de Paris, Pitié-Salpêtrière Hospital, Department of Genetics, Paris, France
[6]Twitter: @ntonius_w
[7]Lead contact
*Correspondence: antonius.wiehler@gmail.com (A.W.), mathias.pessiglione@gmail.com (M.P.)
https://doi.org/10.1016/j.cub.2022.07.010

SUMMARY

Behavioral activities that require control over automatic routines typically feel effortful and result in cognitive fatigue. Beyond subjective report, cognitive fatigue has been conceived as an inflated cost of cognitive control, objectified by more impulsive decisions. However, the origins of such control cost inflation with cognitive work are heavily debated. Here, we suggest a neuro-metabolic account: the cost would relate to the necessity of recycling potentially toxic substances accumulated during cognitive control exertion. We validated this account using magnetic resonance spectroscopy (MRS) to monitor brain metabolites throughout an approximate workday, during which two groups of participants performed either high-demand or low-demand cognitive control tasks, interleaved with economic decisions. Choice-related fatigue markers were only present in the high-demand group, with a reduction of pupil dilation during decision-making and a preference shift toward short-delay and little-effort options (a low-cost bias captured using computational modeling). At the end of the day, high-demand cognitive work resulted in higher glutamate concentration and glutamate/glutamine diffusion in a cognitive control brain region (lateral prefrontal cortex [lPFC]), relative to low-demand cognitive work and to a reference brain region (primary visual cortex [V1]). Taken together with previous fMRI data, these results support a neuro-metabolic model in which glutamate accumulation triggers a regulation mechanism that makes lPFC activation more costly, explaining why cognitive control is harder to mobilize after a strenuous workday.

INTRODUCTION

Even professional chess players start making mistakes, typically after 4–5 h in the game, which they would not make when well rested. A consensual explanation of why chess playing induces cognitive fatigue is that planning moves cannot rely on learned effortless routines (except at the beginning of the game) because the space of possibilities is way too large. Winning the game, therefore, requires the capacity to monitor new context-action mappings, a capacity that is known as cognitive or executive control.[1,2] In neuroscience, behavioral tasks have been developed to vary the demand in cognitive control, by imposing frequent remapping of stimulus-response associations based on contextual information, as in working memory and task-switching paradigms. Using these paradigms, functional neuroimaging studies have identified a lateral prefrontal-parietal system recruited when more cognitive control must be engaged.[3,4]

However, the reason why exerting cognitive control is exhausting remains unclear.[5] Explanations have been proposed in the behavioral economics and social psychology literatures, which have investigated a related notion of self-control or self-regulation. This self-control capacity is notably involved when resisting an impulse (e.g., to eat tasty junk food or to scream with pain) for the sake of long-term goals (e.g., to stay in good health or to remain socially acceptable). Resource depletion theories[6,7] have suggested that exerting such control may tap on global energetic supply (such as blood glucose). However, evidence in favor of these theories has been reconsidered, such that empirical ground is still lacking.[8,9] In any case, these theories fail to explain what is special with self-control and why other cognitive processes such as vision would not induce (and suffer from) global resource depletion. Moreover, a global resource depletion account would contradict the well-shared idea that energy consumption by the brain is constant and globally unaffected by cognitive processing.[10,11]

Noting the absence of biological grounds for fatigue related to cognitive control, other authors have suggested functional explanations. The general idea is that cognitive fatigue would be a sensation generated by the brain, whose purpose is to stop performing the current demanding task for the benefit of a more rewarding activity.[12–14] In that framework, fatigue would stem from a cost-benefit calculation that adjusts the behavior to enjoy available pleasures or avoid opportunity costs. An argument in favor of cognitive fatigue as a functional adaptation is that increasing the payoff of the ongoing task tends to improve performance, even in a state of exhaustion,[15] suggesting that there is no true loss of capacity. However, a difficulty with a purely functional account is that fatigue appears unnecessary: if there is a good reason to stop working on a task and turn to a more gratifying activity, the brain could figure it out without generating an illusion of fatigue. Also, it becomes implausible to maintain that cognitive fatigue is just a functional adaptation when considering the numerous pathological conditions, such as burnout and depression, in which fatigue precisely prevents the patient from enjoying life opportunities.[16,17]

To articulate the functional and biological accounts of cognitive fatigue, we propose (1) that such fatigue stems from an increase in the cost of exerting cognitive control, (2) which in turn stems from metabolic alterations in the brain system underpinning cognitive control. Rather than performance decrement with time-on-task, which can be confounded with boredom, counteracted by training, or compensated by motivation,[18] we reasoned that cognitive fatigue might be better captured by economic choices where monetary benefits are discounted by effort or delay costs. A choice bias for low-cost (LC) options would thus represent an objective marker of cognitive fatigue, even in the absence of a conscious fatigue sensation that could be reported on a psychometric scale. Indeed, subjective fatigue reports are notoriously unreliable, due to limitation of insight, social desirability bias, and variability in the mapping from sensations to rating scales.[19–23]

In a previous study,[24] we developed a daylong protocol mixing cognitive control tasks meant to induce cognitive fatigue and intertemporal choices to reveal cognitive fatigue. To avoid confounding cognitive control exertion with boredom or simply time, we compared groups of participants performing easy and hard versions of the same tasks for the same duration. Only with hard versions did we observe an increase in the preference for immediate rewards over larger later rewards. This increased choice impulsivity after a day of hard work was associated with decreased activity in a specific component of the cognitive control system: the left lateral prefrontal cortex (lPFC). The same behavioral and neural signatures of cognitive fatigue were also observed in endurance athletes suffering from a mild form of burnout due to training overload.[25] Previous results, therefore, support our hypothesis that cognitive fatigue arises from an increase in the cost of recruiting the cognitive control brain system.

However, these previous studies using functional MRI (fMRI) could not address the question of why the cost of cognitive control increases when used for a prolonged duration. This study aimed to fill this gap, using *in-vivo* ^{1}H magnetic resonance spectroscopy (MRS) to quantify metabolites in neural tissues, while participants followed a similar daylong protocol. It has been suggested before that the cost could arise from the need to clear waste products that would accumulate during cognitive control exertion and would potentially be toxic for neural functions.[26] Another possibility would be that the cost arises from the need to restore some energetic resource or metabolic precursor which, unlike blood glucose, would be specifically consumed in cognitive control brain regions. The case of glutamate is particularly interesting, as when released in high quantity, it may be both missing inside the cell (for the neuron to maintain its activity) and disturbing synaptic transmission (to other neurons) outside the cell. Thus, glutamate may be considered a substance that could be either depleted or accumulated with neural activity. To monitor metabolites of the glutamate family, we used an optimized ^{1}H MRS sequence.[27,28] We then looked for a three-way interaction between groups (easy versus hard cognitive work), brain regions (lPFC versus visual cortex), and time on task (session number).

In addition, we used diffusion-weighted ^{1}H MRS[29] to monitor the diffusion of glutamate-related substances,[30] because we anticipated that their concentration might remain constant, even if the metabolic account for cognitive fatigue is correct. The reason is that spectroscopic measures of metabolic concentrations cannot distinguish between cell types nor between cell compartments. A depletion inside the cell could therefore be compensated by an accumulation outside the cell. Diffusion measures are useful to detect this sort of phenomenon because they are differentially sensitive to compartments, as diffusion is more limited inside cells or vesicles than outside cells.[31,32] The same three-way interaction between group, region, and session was therefore tested on diffusion MRS data.

Finally, compared with our previous protocol, we added behavioral tests and measures to better specify the link between cognitive fatigue and the cost of cognitive control. In particular, we included other domains of economic choice with options that trade monetary rewards against costs that were either related to cognitive control (effort discounting) or not (probability discounting). Relatedly, we extended our computational model of economic choice with additional bias parameters favoring LC options in all domains. In addition, we measured pupil dilation during decision-making as an index of cognitive effort invested in deliberation, as we had no fMRI measurement to document the reduction in cognitive control exertion. Last, to better dissociate these choice-related markers of cognitive fatigue from experienced fatigue sensation, we collected subjective self-reports on a visual rating scale.

RESULTS

Behavioral measures of cognitive fatigue

To induce cognitive fatigue, participants followed a previously validated experimental protocol[24] (see Figure 1). For 6.25 h, participants performed cognitive tasks (N-switch and N-back) that require cognitive control. In both tasks, letters were displayed on a computer screen every 1.60 s. In the N-switch task, participants had to perform either a vowel versus consonant or an upper versus lower case discrimination task, depending on the color of the letter. The difficulty in this task depends on the rate of switches (i.e., color changes), which was 1 versus 12 in 24 trials for the easy versus hard versions. In the N-back task,

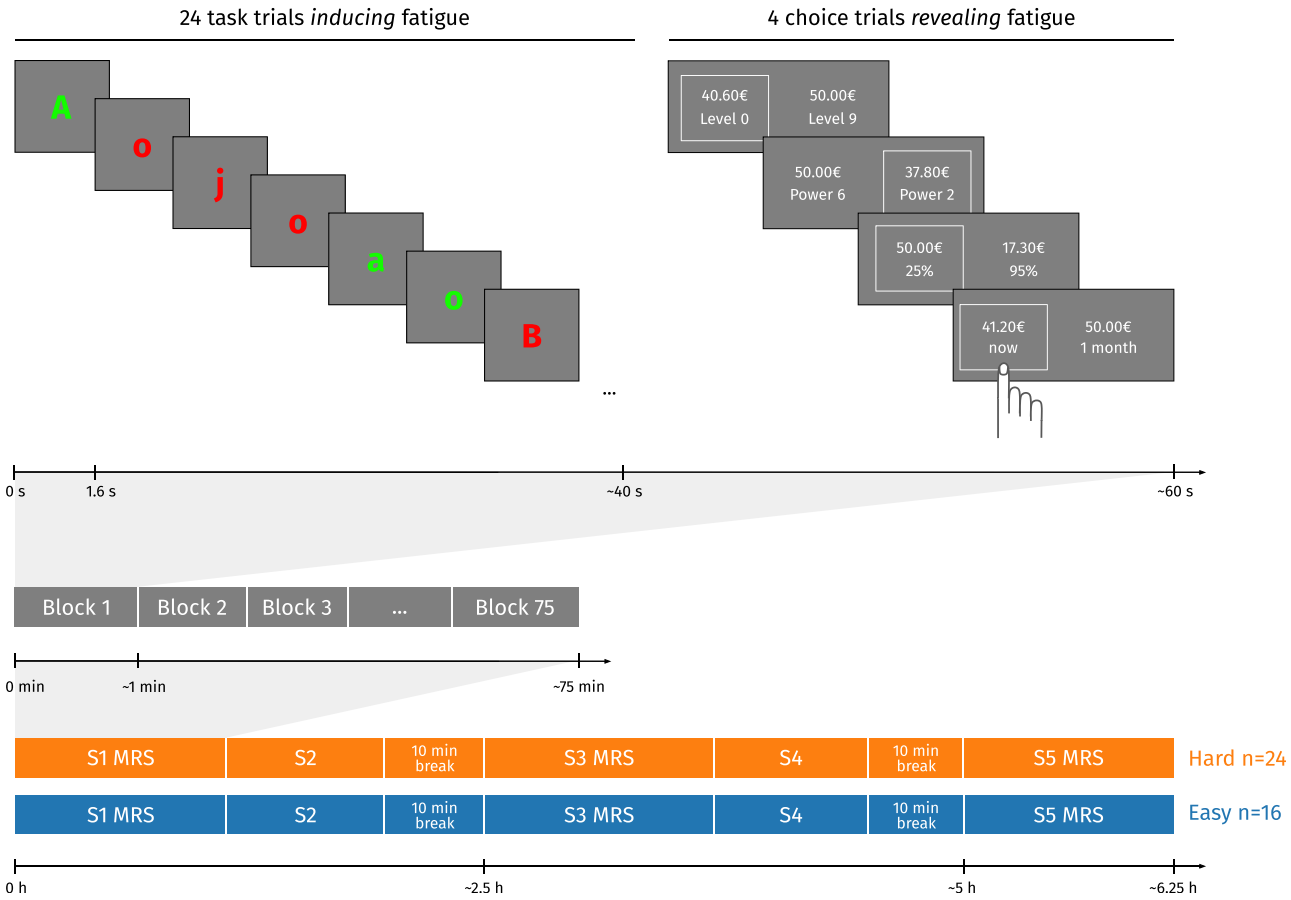

Figure 1. Experimental design

From top to bottom, the protocol is shown with diminishing time resolution, from single-trial to daylong experiment. In two training sessions (not shown) preceding the testing day, participants learned to perform the cognitive tasks with a correct response rate higher than 90% and practiced the economic choices to reveal their indifference points. During the main experiment, participants alternated between cognitive tasks and economic choices. One group of participants (n = 16) was assigned the easy version and the other group (n = 24) the hard version of cognitive tasks. In every block, cognitive tasks included either 24 N-switch trials (1 versus 12 switches in the easy versus hard condition) or 24 N-back trials (1 versus 3-back in the easy versus hard condition). The task to do was announced at the beginning of the block and was changed once per session. Economic choices included one trial per cost domain (delay, probability, and physical and cognitive effort). In all choices, the two options were a variable reward at a low cost versus 50€ at a variable cost. Participants performed 5 sessions (S1–S5) of 75 blocks, for a total duration of 6.25 h. Three of these sessions were performed in the scanner to simultaneously collect magnetic resonance spectroscopy (MRS) data.

participants had to state whether the letter on-screen was identical or different from the letter presented in N trials before. The difficulty in this task depends on the distance between trials (i.e., the load of information to keep in working memory), which was 1 versus 3 for the easy versus hard versions. Participants were split into two groups: the test group (n = 24) performed the hard version (12-switch and 3-back), whereas the control group (n = 16) performed the easy version of cognitive tasks.

In our search for behavioral signatures of mental fatigue, we started with traditional measures: performance decrement and self-report. Both groups were trained on the day before and maintained a high level of performance throughout the test day (correct response rate >80%, see Figure 2A). A three-factor (group × session × task) linear mixed model fit on performance showed a significant difference between groups ($\beta = 0.06$, $p < 0.01$) and a significant decrease with session number ($\beta = -0.01$, $p = 0.001$), but no significant difference between tasks ($\beta = -0.009$, $p = 0.68$) and no significant group-by-session interaction ($\beta = -0.002$, $p = 0.61$). In log-transformed response times (RTs), the same model showed a significant difference between groups ($\beta = -0.11$, $p < 0.001$), a significant difference between tasks ($\beta = -0.14$, $p < 0.001$), but no significant effect of session ($\beta = 0.001$, $p = 0.63$) and no significant group-by-session interaction ($\beta = -0.004$, $p = 0.20$). Thus, performance measures confirmed that low-demand tasks were easier (with higher accuracy and shorter RT) than high-demand tasks but provided no evidence for a fatigue effect (across sessions) that would differ between groups.

In addition to objective performance, participants were also asked to rate their subjective level of fatigue on a rating scale. A two-factor (group × session) linear mixed model analysis showed a significant increase with session number ($\beta = 9.30$, $p < 0.001$) but no significant difference between groups ($\beta = -8.14$, $p = 0.19$) and no significant interaction ($\beta = -0.32$, $p = 0.85$). Thus, the subjective report could not capture a fatigue that would specifically relate to task difficulty, hence to the load

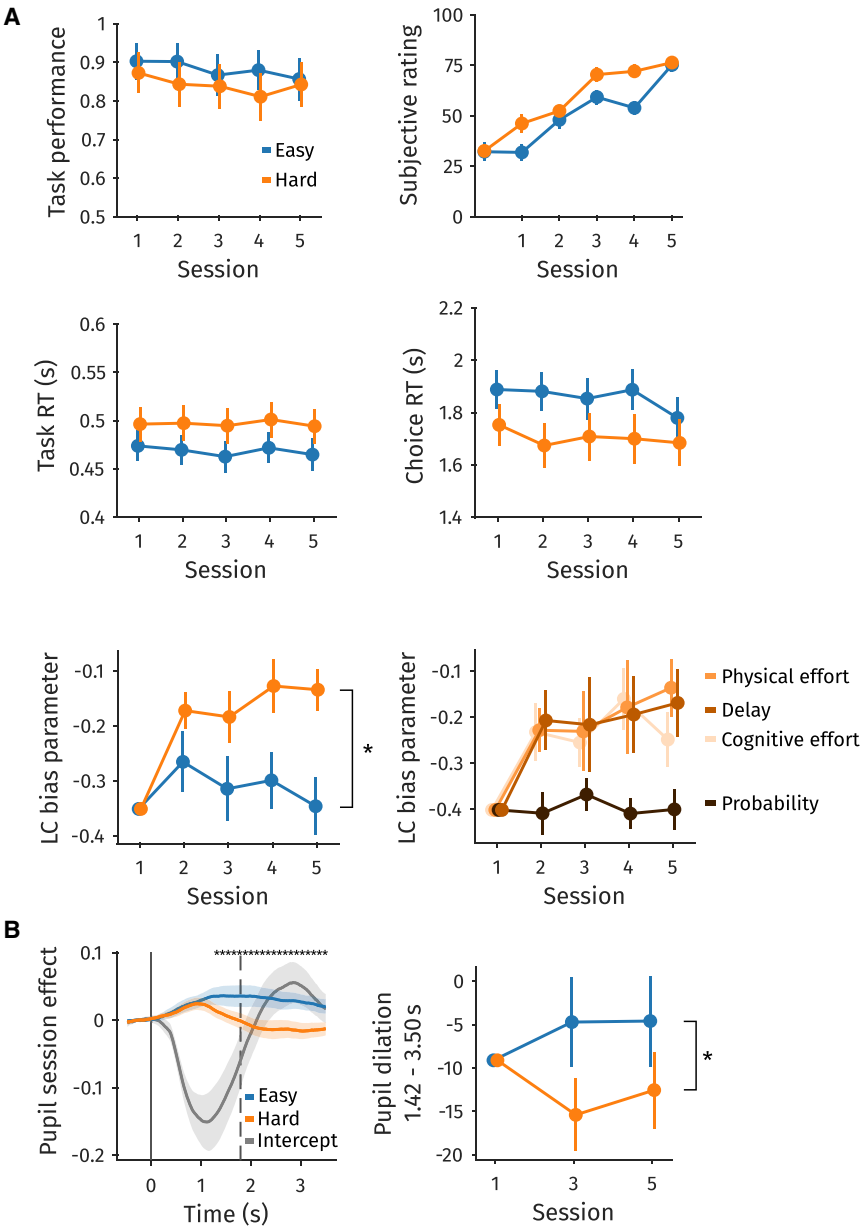

Figure 2. Behavioral results

(A) Behavioral measures of cognitive fatigue. Top graphs: average correct response rate (left) in cognitive tasks (pooling N-switch and N-back trials) and subjective rating of experienced fatigue (right), between 0 ("I'm in top form") and 100 ("I'm totally exhausted"), separately for the groups performing easy and hard versions. Middle graphs: median response time (RT) during N-switch and N-back tasks (left) and median RT during economic choice tasks except probability discounting (right). Bottom graphs: average bias parameter in economic choices (pooling all cost domains but probability) across experimental sessions (left) and for the different cost domains in the hard condition (right) shown across experimental sessions. The bias parameter is an additive bonus for the low-cost (LC) option in the choice model. Bias data are normalized to the grand mean in the first session. See Figure S1 for model-free behavioral results and Table S1 for results of the statistical analysis.

(B) Pupillary measures of cognitive fatigue. The left graph shows estimates (beta weights) from a linear regression model meant to explain pupil size, at each time point within an economic choice trial. The model included intercept, group, and session as factors of interest and various factors of no interest (including choice and response time, see STAR Methods). Before linear regression, pupil recordings were preprocessed, aligned to choice onset (i.e., option display), and baseline corrected (by subtracting the mean over the −500–0 ms time window). The vertical dashed line indicates the median RT. The intercept shows the time course of pupil dilation, irrespective of group and session. The impact of cognitive fatigue corresponds to the difference in session effects between the easy and hard conditions (stars indicate time points when the difference is significant). For visual comparison with the other behavioral results, the right graph shows non-*Z* scored, but baseline corrected, pupil size during the time window when the group difference was significant, per group and session (normalized to the grand mean of the first session). In all graphs, error bars indicate inter-participant standard errors of the mean (SEM). The error bars for subjective ratings of fatigue are plotted but are too small to be visible in the graph. Stars on brackets indicate significant group-by-session interactions.

on cognitive control. To assess other potential psychological states that may have affected choices, we also asked participants to rate their level of hunger and stress. Hunger ratings increased before lunch (sandwich and fruit eaten during the 10-min break) but trivially decreased afterward, whereas stress ratings were stable throughout the day (Figure S1A). These results suggest that self-reported fatigue (like hunger and stress) should be distinguished from the actual fatigue related to the exertion of cognitive control.

To reveal cognitive fatigue associated with demand on cognitive control, we turned to economic choices. Participants made four choices after each 24-trial block of cognitive task trials. Each choice opposed a small-reward/low-cost (LC) option to a big-reward/high-cost (HC) option. The four choices corresponded to the four possible cost domains: (1) delay in reward delivery (bank transfer, as in previous protocol), (2) cognitive effort (difficulty level of a 30-min block of N-switch task to be performed after the experiment), (3) physical effort (power of a 30-min cycling session on a home bike after the experiment), and (4) probability of reward delivery (in a lottery played after the experiment). Big rewards were always 50€, and HCs were randomly varied across five predefined levels. LCs were either zero or a close-to-zero fixed level (such as 3 days for delay discounting). Small rewards (associated with LC options) were adjusted to individual indifference points, in a calibration session before the experiment, such that each participant started the protocol in the morning with a rate of LC choices of around 50%.

To analyze economic choices, we used a computational modeling approach that was previously validated.[23,25] Models included a domain-specific discounting function (with a free discount parameter k) that integrated rewards and costs to generate option values, and a softmax choice function (with a free inverse temperature β) that generated selection probabilities from the difference between option values. In addition to these standard parameters (k for the steepness of discounting and β for the consistency of choices), we included an additive *bias* parameter that shifted the softmax function toward the selection of the LC option. All free parameters were estimated per participant and session.

To determine whether a computational parameter would capture cognitive fatigue related to task difficulty, we looked for group-by-session interactions in a two-factor linear mixed model analysis performed across cost domains. The interaction was neither significant for discounting steepness k ($p = 0.24$), nor for choice consistency β ($p = 0.63$), but was significant for the LC bias ($p = 0.02$, see Figure 2A and Table S1). The same analysis performed separately in each cost domain showed that the interaction was significant for delay, cognitive, and physical effort discounting (all $p < 0.05$) but not for probability discounting ($p = 0.98$). Thus, the signature of cognitive fatigue was captured by the same LC bias parameter as in our previous studies,[23,25] now extended to other cost domains that involve cognitive control (not only to wait for gratification but also to exert effort), but not to factors that do not require cognitive control (such as the outcome probability in a lottery).

Note that a model-free analysis would lead to similar conclusions. Indeed, the LC choice rate followed the same pattern as the LC bias parameter: no significant change in the low-demand group, but a highly significant increase ($p < 0.001$) in the high-demand group. Among domains, the LC choice rate significantly increased for delay, physical effort, and cognitive effort discounting, but not probability discounting (see Figure S1), again mirroring the LC bias parameter. Thus, the computational decomposition into parameters capturing different sources of variance only contributed to clarifying the impact of cognitive fatigue on economic choice, with a significant group-by-session interaction. Note also that including performance as a covariate in the linear mixed model analysis did not change the main result (additional performance regressor: $p = 0.85$, group-by-session interaction: $p = 0.02$), ruling out a potential effect of frustration related to the error rate (which participants could not monitor anyway because they were provided no feedback about their performance in cognitive tasks).

The computational signature of cognitive fatigue suggests a shift of preference favoring the options involving low control costs. An alternative would be that fatigue changes the way choices are made, favoring heuristics that avoid the pain of deliberation. From this perspective, participants might choose the LC option because it is easier to valuate. However, this would not explain why we observed a similar shift toward options with non-zero LC, which are harder to valuate than zero-cost options. When adding the zero/non-zero-cost factor to the linear model, we found no evidence for a different effect of session on the two trial types (interaction trial-type × session: $\beta = 0.003$, $p = 0.87$). In addition, a time-saving heuristic would predict that choices would be made faster and faster along the day, which was not the case (see Figure 2A). Indeed, a three-factor (group × session × choice) linear mixed model fitted on choice RT across cost domains (except probability) showed borderline effect of task difficulty (hard versus easy: $\beta = 0.11$, $p = 0.06$), no main effect of session number (from 1 to 5: $\beta = -0.01$, $p = 0.34$), no effect of choice type (LC versus HC option: $\beta = -0.01$, $p = 0.58$), and, critically, no significant interaction (all $p > 0.14$). Thus, there was no evidence for cognitive fatigue resulting in the use of time-saving heuristics.

Pupillary measures of cognitive fatigue

Although fatigue induced a systematic bias favoring options that require less cognitive control, we could not check whether this bias was related to reduced lPFC activity during choice, as we observed in previous studies,[24,25] because there was no fMRI scanning here. However, as an index of cognitive effort exerted during economic choice, we measured pupil dilation, which has been associated to the demand in cognitive control and the level of effort invested.[33]

Within-trial pupil size time series were submitted to a sample-wise multiple regression against session number and all relevant nuisance variables (see STAR Methods). The time course of the intercept (see Figure 2B, left plot) shows, irrespective of experimental variables, an initial constriction for the first second, followed by a dilation ending with the choice. The initial constriction is commonly seen as a response to low-level features, like a surprise signal evoked by a change in stimulus display, independent of luminance,[34] whereas the late dilation likely reflects the effort invested in the deliberation.[33] When comparing the regression estimates of session number between groups, we observed a significant difference (after correcting for multiple comparisons using random field theory) in a late time cluster (from 1.42 to 3.50 s). The difference was driven by the late dilation plummeting with session number in the hard group (see Figure 2B, right plot), suggesting that less effort was invested in choice deliberation with cognitive fatigue.

Metabolic measures of cognitive fatigue

Behavioral results suggest that the fatigue-induced preference shift toward LC options is associated with less effort invested in decision-making. This pattern is consistent with our proposal that cognitive fatigue can be conceived as an increase in the cost of recruiting cognitive control. To explain why control cost is increased by the performance of hard cognitive tasks, we turned to brain metabolism. Our model (see Figure 3A) assumes that recruiting cognitive control regions either exhaust some metabolic resource or accumulate some toxic by-product. This metabolic alteration may somehow be sensed by a meta-controller that would adjust the intensity of control exertion depending on expected costs and benefits, as previously suggested.[35] Thus, the diminution of control exertion would be the outcome of a regulation loop aiming at maintaining the concentration of metabolites within acceptable limits. Increasing the expected benefits could naturally counteract the impact of increased metabolic costs and maintain the intensity of control exertion. This is why the reduction of control exertion is more salient in economic choices, where the expected benefit cannot be precisely estimated, compared with cognitive tasks, in which correct performance results in a precise payoff.

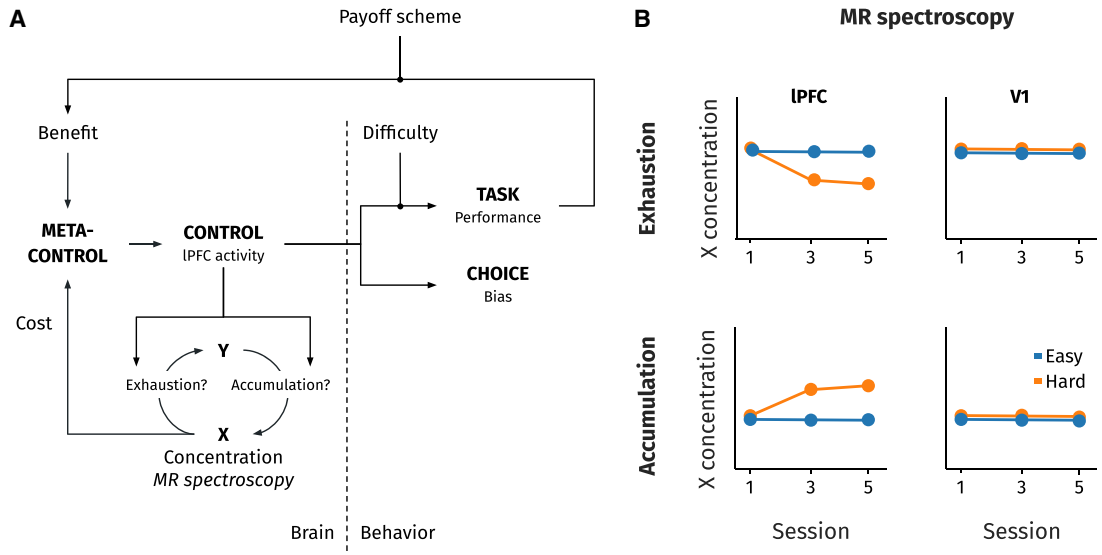

Figure 3. Neuro-metabolic predictions

(A) Model. Our conceptual model assumes that the amount of cognitive control to be invested over the behavior is submitted to a cost/benefit trade-off. The benefit of control is to improve performance (depending on task difficulty) and hence the associated reward (depending on payoff schedule), or to make a sound economic choice. The cost of control relates to the necessity of either restoring an exhausted resource or clearing out an accumulated by-product of neural activity. The model, therefore, specifies the relationships between behavioral observables (accuracy in cognitive tasks and bias in economic choice) and brain measures (lPFC activity with fMRI, and X concentration with MRS).

(B) Predictions. Regarding the metabolite X (be it an exhausted resource or an accumulated by-product), the model predicts a three-way interaction between region (lPFC versus V1), group (easy versus hard conditions), and session (growing fatigue).

To keep the predictions simple (see Figure 3B). We assumed that metabolic alteration (either exhaustion or accumulation) (1) would only occur when performing hard cognitive tasks (not easy ones), (2) would only affect the lPFC (the region identified as related to cognitive fatigue in previous studies) and not the primary visual cortex (V1) (although it is constantly stimulated in this experiment), and (3) would progressively develop with session number, until it reaches a plateau due to meta-regulation. Thus, the two metabolic hypotheses predicted a (group × region × session) three-way interaction.

To test for this three-way interaction, we measured *in-vivo* metabolic concentrations with MRS, using an optimized semi-LASER ^{1}H sequence,[27,28] in both the left lPFC and V1 voxels of interest (VOIs), during sessions 1, 3, and 5 of the experiment (Figure 4A). The only metabolite showing a significant three-way interaction was glutamate, irrespective of the ratio computed for normalization (Figure 4; see Figure S2 for other metabolites): glutamate/myo-Inositol (Glu/Ins, $\beta = 0.06$, $p = 0.02$), glutamate/total creatine (Glu/tCr, $\beta = 0.04$, $p < 0.01$, see Table S2), and glutamate/total N-acetyl aspartate (Glu/tNAA, $\beta = -0.02$, $p = 0.03$). A log-likelihood ratio test confirmed a significant improvement of the linear mixed model fit when adding interaction terms (Glu/Ins: $p < 0.001$, Glu/tCr: $p < 0.001$, Glu/tNAA: $p < 0.001$).

The observed pattern was consistent with the accumulation hypothesis, in the sense that the group performing hard cognitive tasks ended the experiment with more glutamate in the lPFC than the control group, (Glu/Ins: $t = 2.23$, $p = 0.03$, Glu/tCr: $t = 2.27$, $p = 0.03$, Glu/tNAA: $t = 1.68$, $p = 0.10$, between-group t test in session 5), no difference being observed in the visual cortex (Glu/Ins: $t = -1.10$, $p = 0.28$, Glu/tCr: $t = -0.95$, $p = 0.35$, Glu/tNAA: $t = -0.44$, $p = 0.66$; between-group t test in session 5). However, contrary to our expectations, the interaction was not driven by an increase in the hard condition but by a decrease in the easy condition.

We also measured metabolic diffusion in the same VOIs using a second semi-LASER sequence added with diffusion gradients[29] and again found a significant three-way interaction involving glutamate (Figure 4B, lower part). More specifically, the interaction was observed for the apparent diffusion coefficient (ADC) of Glx (glutamate and glutamine quantified together: $\beta = 6.95e{-}06$, $p = 0.02$, see Table S3). This interaction was driven by Glx ADC increasing across sessions in the lPFC (but not in V1) of participants performing the hard version of cognitive tasks (but not the easy version). This is consistent with the hypothesis that glutamate concentration is regulated during cognitive control exertion, potentially leading to higher concentrations in the extracellular space where diffusion is faster. The pattern was in line with the prediction of the metabolic accumulation model, as the interaction was driven by an increase in the hard condition versus a steady state in the easy condition.

Finally, we assessed the correlation between behavioral and neuro-metabolic measures of cognitive fatigue, as was done in our previous studies between behavioral and fMRI measures.[24,25] After regressing out variables of no interest (see STAR Methods), we tested the correlation across participants between the increase in LC bias (from session 1 to session 5) and the increase in Glu concentration and Glx diffusion, separately (see Figure 5). The correlation was positive in both cases, but significant for Glx diffusion only ($r = 0.43$, $p = 0.039$), not Glu concentration ($r = 0.30$, $p = 0.24$). The significant correlation suggests that the choice bias toward LC options was linked to the level of Glx diffusion and hence to a need for slowing down glutamate accumulation.

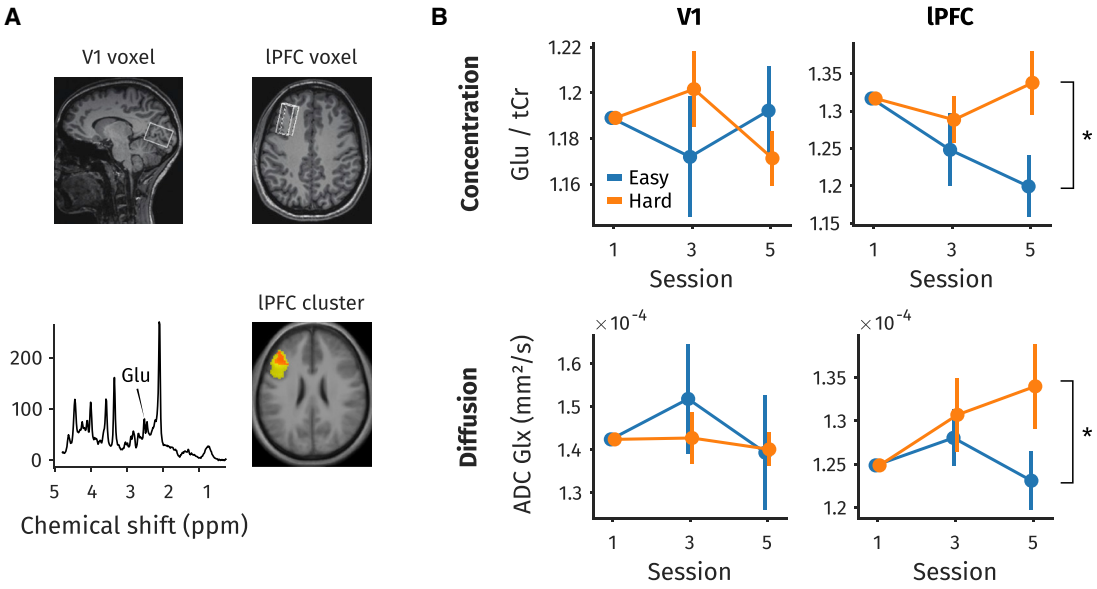

Figure 4. MR spectroscopy results

Glutamate was the only metabolite to show the expected three-way interaction (group × region × session).

(A) Data collection. The left lateral prefrontal region (axial slice at the bottom) was defined as the cluster showing reduced activity with cognitive fatigue in a previous fMRI study.[24] On top slices are shown the locations of the MRS voxels of interest (VOIs, 35 × 25 × 15 mm) in the primary visual cortex (V1, left) and in the lateral prefrontal cortex (lPFC, right). The lPFC VOI was individually adjusted to cover most of the cluster identified with fMRI, whereas the V1 VOI was placed to cover the medial part of the occipital cortex. The glutamate peak (Glu) is indicated within an example individual spectrum acquired in the lPFC.

(B) Main results from ^{1}H MRS. Top: glutamate concentration levels were normalized over concentrations of total creatine (tCr, shown here), myoinositol (Ins, shown in Figure S2), and total N-acetyl aspartate (tNAA, shown in Figure S2). The three-way interaction (group × region × session) was significant in the three normalized measures. For visual comparison with behavioral measures and model predictions, data were also normalized to the grand mean of the first session. See Figure S3 for additional metabolites. Bottom: main results from diffusion-weighted MRS. Glx (glutamate plus glutamine) were again the only metabolites for which diffusion measures (apparent diffusion coefficients [ADCs], see STAR Methods) showed the expected three-way interaction. Error bars indicate inter-participant standard errors of the mean (SEM). Stars on brackets indicate significant group-by-session interactions.

See Tables S2 and S3 for results of the statistical analysis.

Modeling the link between neural activity and metabolic measures

To examine whether the downregulation of cognitive control, resulting in decreased lPFC activity, as observed using fMRI in our previous study,[24] could explain the pattern of lPFC glutamate measures, as observed using MRS in the present study, we developed a Markov chain model. This model (see Figure 6) predicts the evolution of a metabolic measure *X* across time, depending on two opposite flows: an accumulation due to control exertion (indexed by fMRI activity) and a clearance proportional to *X* (measured by MRS).

Results show that MRS measures were consistent with the principle of a metabolic regulation mechanism, reducing lPFC activity to maintain glutamate accumulation within acceptable levels (i.e., manageable by clearance processes). However, the concentration and diffusion measures suggest different parameterizations of the model, notably regarding baseline lPFC glutamate levels in the morning (relative to glutamate accumulation rate).

DISCUSSION

In this study, we investigated the impact of performing hard cognitive control tasks for several hours, compared with performing easy versions of the same tasks for the same duration. We observed (1) a shift in preferences toward LC options, (2) a reduction of pupil dilation during economic choice, (3) a glutamate concentration maintained at a high level in the lPFC, and (4) an increase in glutamate/glutamine diffusion within the lPFC. This pattern of results is compatible with the assumption of an increase in the cost of cognitive control, related to the necessity of maintaining glutamate levels within acceptable boundaries. The elevated cost would both limit the recruitment of cognitive control during choice and bias decisions away from costly options.

The present behavioral results replicate and extend previously published observations that exerting intense cognitive control, either for intellectual work[24] or endurance sport,[25] induces a form of cognitive fatigue that manifests as an increased preference for immediate options. The key results are significant group-by-session interactions, showing a specific effect of task difficulty, ruling out time or boredom as potential explanations. We did not find such an interaction in self-reported fatigue, which similarly increased across sessions in the two groups. This may be due to participants mapping the range of their subjective sensations on the same portion of the rating scale, as they were unaware of the other condition (imposed on the other group). It could also reflect a dissociation between actual fatigue of the cognitive control brain system and conscious perception of fatigue. This dissociation is common in everyday life; for instance, when people go on working or driving and start making errors because they failed to detect their true fatigue state. In any case, it shows

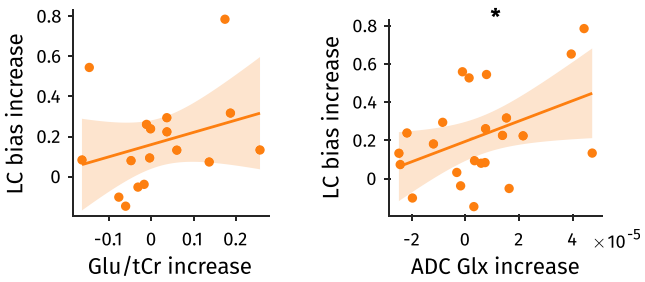

Figure 5. Correlation between behavioral and MR spectroscopy measures

For each participant of the high-demand group, the increase (from first to last session) was calculated for both the behavioral measure (choice bias toward low-cost option [LC] after regressing out the effect of cost domain) and the MR spectroscopy measures (Glu/tCr ratio or Glx apparent diffusion coefficient [ADC] after regressing out nuisance variables such as quality estimates of metabolic spectra, gray matter concentration in the scanned voxel, and head movement parameters). Graphs show the correlation across participants (dots), with regression lines and confidence intervals (shadow areas). The star indicates that the correlation with LC bias was only significant for ADC Glx.

that subjective ratings cannot be taken as absolute measures and that cognitive fatigue might be better evidenced by preference shifts toward LC options in economic decisions. Conversely, choice-related markers of cognitive fatigue might not account for the subjective perception of intense fatigue that represents a frequent clinical symptom in many neuropsychiatric conditions.[23] As in previous studies, this fatigue did not affect performance in cognitive control tasks (N-back and N-switch), suggesting that cognitive control exertion is sustained in these tasks, despite an increase in its cost, due to the high benefit attached to correct performance. Indeed, making correct responses in cognitive control tasks entailed an objective monetary benefit, whereas the benefit of making a sound decision in economic choice tasks was more difficult to estimate. The new observations provide further support for an interpretation of the preference shift as stemming from an elevation of cognitive control cost.

First, we have introduced other types of discount factors in our economic choices (probability, cognitive effort, and physical effort), on top of delay. The shift in preference was specifically observed for choices that involve cognitive control (waiting longer or exerting more effort to obtain better rewards), not for choices involving risk (a cost associated to the lottery, not imposed on the participant). These results are consistent with the idea that mental and physical effort both involve cognitive control because they both require to over-rule automatic routines.[25] In all cases but probability discounting, the shift in preference was captured in the choice model by an additive bias favoring LC options, not by the choice consistency (inverse temperature) parameter. Note that a decrease in choice consistency could not mimic a preference shift with our design because options were tailored around individual indifference points, such that participants started the experiment at chance level (with a 50/50 preference between LC and HC options). Also, choices were not more impulsive in the sense that they were made faster with cognitive fatigue. As in previous experiments, choice RT did not vary with the number of sessions completed along the day and did not depend on which option was selected (costly or uncostly).

Second, we used eye-tracking during scanning sessions and observed that cognitive fatigue was accompanied by lesser pupil dilation when making a choice. This can be taken as evidence for lower cognitive effort invested in economic choice, consistent with the reduced lPFC activity that was found in our previous study using fMRI. Indeed, pupil dilation has been validated as an index of cognitive effort.[33,36–40] Pupil dilation has also been associated with the activation of noradrenaline neurons in the locus coeruleus and hence activation of the anterior cingulate cortex.[41–44] It is tempting to interpret these results in the light of the theoretical framework assuming that the intensity of cognitive control is adjusted by the anterior cingulate cortex, depending on expected costs and benefits.[35] Applied to our case, the decrease in pupil dilation would suggest that the reduction of cognitive control during economic choice, due to its elevated cost, is mediated by the anterior cingulate cortex downregulating activity in the lPFC, which we observed with fMRI.

The next question in this general account of cognitive fatigue was about the cost of cognitive control: why is it increasing with the performance of high-demand tasks? To articulate the cost-benefit arbitration framework with a neuro-metabolic account of cognitive fatigue, we imagined two scenarios: cognitive control could be reduced to prevent some resource from dangerous exhaustion or to prevent some by-product from dangerous accumulation. These two scenarios predicted a three-way interaction between group, region, and session, which we only found in glutamate levels, whatever the normalization procedure. At the end of the day, lPFC glutamate concentration and glutamate/glutamine diffusion were significantly higher in the group performing high-demand tasks relative to the low-demand group, although there was no difference in the visual cortex. These observations are consistent with higher-demand cognitive control tasks being associated with greater glutamate release,[45,46] which would result in steeper glutamate accumulation with time on task across a workday.

However, the interaction observed in glutamate concentration was mainly driven by a decrease in lower-demand conditions, which we did not expect. However, we verified that such a pattern was still consistent with our dynamic model including clearance proportional to glutamate concentration and accumulation related to the intensity of cognitive control (hence to lPFC activity). The fitted parameters indicated that what we had not anticipated was a high level of lPFC glutamate at the beginning of the day. Facing the new scanning environment and implementing new instructions (e.g., to provide a manual response without moving the head) might already be cognitive control demanding,[47] possibly explaining the elevated glutamate level. In this scenario, the gradual elimination of glutamate would be observable in the low-demand conditions, although it would be compensated by gradual accumulation related to task performance in the high-demand conditions. Such a scenario was corroborated by results showing that glutamate/glutamine diffusion was higher in the lPFC after high-demand cognitive control exertion, compared with the control group and region. Diffusion measures displayed the expected pattern, with an increase related to intense cognitive control, and no change in low-demand conditions. The interpretation can hence be refined, as a change in glutamate/glutamine diffusion might signal a relative accumulation in the extracellular compartment (where diffusion

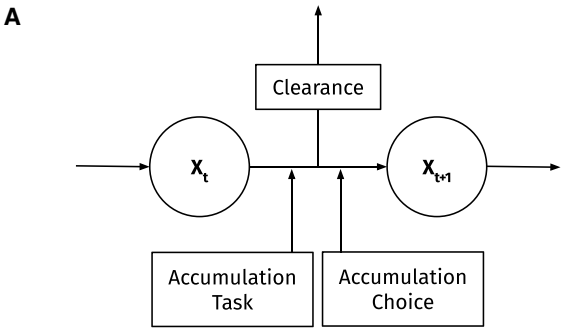

$$Glu_t = Glu_{t-1} - \beta_{clear}Glu_{t-1} + \beta_{task}BOLD_{task,t} + \beta_{choice}BOLD_{choice,t}$$

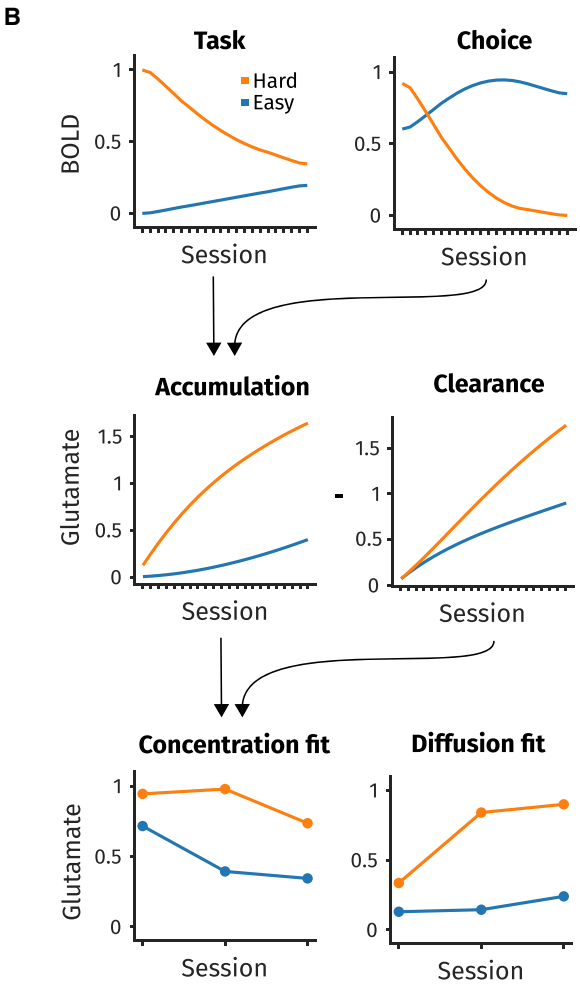

Figure 6. Model linking neural activity to metabolic measures

(A) Formalization of the dynamic model of metabolite concentration. In our model, the evolution of metabolite X across time follows a Markov chain. X_{t+1} is determined by X_t, minus a passive clearance depending on X level, plus accumulation of X due to exerting control at time t (for cognitive tasks or economic choices). We illustrate here the accumulation model because it fits with glutamate measurements. Note that a resource exhaustion model would follow the same logic, except that exerting control would deplete X, which would need to be restored instead of cleared.

(B) Predictions of glutamate measurements. Top plots represent the input to the model, i.e., lPFC fMRI activity extracted from a previous study.[24] Data were normalized between 0 and 1 across conditions, but separately for the cognitive tasks and economic choice, upsampled to 22 sessions and smoothed with a moving average. Middle plots show the two components driving the dynamics of glutamate measure: accumulation due to lPFC activity (during both cognitive task and economic choice) and clearance proportional to glutamate level. The bottom plots show two simulations with different sets of parameters that correspond to glutamate concentration measures (β_0 = 0.85, β_{clear} = 0.09, β_{task} = 0.12, β_{choice} = 0.01) and glutamate/glutamine diffusion measures (β_0 = 0.14, β_{clear} = 0.04, β_{task} = 0.11, β_{choice} = 0.001). In both cases, the interaction between group and session is driven by the reduction of lPFC activity with fatigue in the hard group.

of molecules is faster), presumably related to strong spiking activity.[31,32] Note however that although release in the extracellular space is a standard interpretation of the diffusion measure, it could in principle reflect other phenomena, such as release from vesicles to the intracellular space, or change in the glutamate/glutamine ratio (because the two molecules may diffuse in different compartments).

Obviously, our results are only correlational and cannot be taken as proof that what limits cognitive control exertion is the need to prevent glutamate accumulation. Causal manipulations would be required to validate this assumption. In addition, the metabolic spectrum was narrow and constrained by technical limitations, as there are metabolites that cannot be quantified with *in-vivo* MRS methods or at least in a 3T MRI scanner, using a semi-LASER sequence, with the echo-time optimized for glutamate. For instance, GABA would also have been a possible candidate but could not be reliably quantified with our MRS data acquisition sequence. In any case, it should be noted that the target substance requiring regulation may not be glutamate itself but any substance whose concentration is linked to glutamate accumulation. Nevertheless, glutamate regulation has been pointed out as an essential component in the brain energy budget and discussed as a potential source of cognitive fatigue.[48–50] Indeed, there are good reasons for which glutamate accumulation may need to be regulated.

Glutamate is well known as the main excitatory neurotransmitter in the brain, which must be maintained in tight balance with inhibitory neurotransmission for regular cortical functioning.[51–53] Glutamate is present in the cells at high concentrations, as it is involved in the detoxification of ammonia and also serves as a precursor for the synthesis of proteins.[54] It is therefore important to limit glutamate release, both because it is a useful resource in the intracellular compartment and because it is a potentially toxic by-product in the extracellular compartment. In line with our results, extracellular glutamate tends to accumulate in stressful conditions or with increasing task demands such as working memory load.[46,55] The issue with too high concentrations of extracellular glutamate is not only the disruption of excitation/inhibitory balance but also the induction of activation bursts, which might impair the transmission of information and cause excitotoxicity in the most severe cases.[51,52,54] A known regulation mechanism at the synaptic level is glutamate reuptake through transportation into surrounding glial cells[56,57] or axons[58,59] and conversion into glutamine. Unfortunately, our measurement technique was not sensitive enough to explore molecular/cellular mechanisms. Although we could distinguish

concentrations of glutamate and glutamine, diffusion measures pooled the two metabolites. Additional limitations relate to the low spatial and temporal resolution of MRS scanning. As our VOI was about 40 cm^3, it is impossible to draw precise conclusions about anatomical locations, and because data acquisition takes about 10 min, it is impossible to know which particular task events most contributed to glutamate accumulation.

In the dynamic model, fMRI and MRS measures were integrated to verify that our data were compatible with glutamate accumulation being the trigger of cognitive control regulation. Nevertheless, there are some gaps in this demonstration, as we took some shortcuts. An obvious one is that we did not collect fMRI and MRS data in the same participants; hence, we could only make predictions at the group level and could not test inter-individual correlations. However, the assumption that the BOLD signal is linearly related to glutamate release is corroborated by simultaneous fMRI and MRS recordings in the visual cortex.[60] On a related note, the very principle of the model, postulating that glutamate is regulated, may have weakened inter-individual correlations between glutamate levels and behavioral signatures of fatigue. This is because glutamate level is supposed to be maintained at a given boundary, by reducing lPFC activity during choice; hence, the choice bias inferred from the behavior should be strongly correlated with lPFC activity, as we indeed observed in our previous studies,[24,25] but only weakly with glutamate level, as we observed here. Note that our MRS findings are not just mirroring the fMRI findings: lPFC BOLD activity was reduced across sessions, whereas glutamate concentration was steady or slightly increasing, as was glutamate/glutamine diffusion. This is evidence of some accumulation taking place: if glutamate was purely reflecting the momentary activity of the brain region, it should have plummeted throughout the day.

Even if our model provides proof of concept that a metabolic account of cognitive fatigue can be combined with a cost-benefit mechanistic framework, several aspects remain speculative at this stage. Notably, how glutamate levels would be monitored to estimate the costs of cognitive control is unclear. It remains possible that the brain may not monitor glutamate itself but any phenomenon linked to glutamate accumulation (e.g., GABA synthesis). Also, an explanation is still missing for why cognitive control regions would accumulate glutamate more than other regions like the visual cortex. On a different note, research is needed to explore the recovery of glutamate levels at rest or during sleep. Interestingly, the cognitive control network is deactivated in rest conditions that activate the default mode network,[10,61] which could favor the clearance of extracellular glutamate. Moreover, it has been shown that glutamate concentrations decrease during sleep, in relation to EEG slow-wave activity.[62] Glutamate could therefore belong to the potentially toxic substances that are eliminated during sleep, which could mediate recovery from cognitive fatigue.[63] Finally, how cognitive fatigue due to excessive use of cognitive control relates to other forms of fatigue remains to be specified. As it was also observed in a mild form of burnout syndrome[25] and patients with low-grade glioma,[23] we tend to believe that an elevated cost of cognitive control is key to several clinical manifestations of fatigue,[64,65] but this speculation still requires empirical confirmation. It would also require a theoretical articulation between the objective fatigue of the cognitive control brain system documented here with choice-related markers and the subjective fatigue sensation that might represent the main complaint of patients in the clinics.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
  - Lead contact
  - Materials availability
  - Data and code availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
  - Participants
- METHOD DETAILS
  - Tasks
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Computational modeling
  - Pupil dilation
  - MRS data acquisition
  - MRS data analysis
  - Diffusion-weighted magnetic resonance spectroscopy acquisition
  - DW-MRS spectral processing
  - Metabolic accumulation model
  - Statistical analysis

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2022.07.010.

ACKNOWLEDGMENTS

The work was supported by a Labex "Bio-Psy" (http://www.biopsy.fr/index.php/fr/) research grant and by the "Big Brain Theory" program of the Paris Brain Institute. The MRS package was developed by Gülin Öz, Ivan Tkac, Dinesh Deelchand (semi-LASER sequence), Edward J. Auerbach, and Małgorzata Marjańska (diffusion-weighted semi-LASER sequence) and was provided by the University of Minnesota under a C2P agreement. We would like to thank Jean Daunizeau and Gerhard Jocham for their insightful comments.

AUTHOR CONTRIBUTIONS

A.W., F.M., and M.P. designed research. A.W. conducted experiments and collected data. A.W., F.B., I.A., and M.P. conceived behavioral and neural analyses. A.W. analyzed data. A.W. and M.P. wrote the paper and all authors provided feedback on the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: October 8, 2021
Revised: May 27, 2022
Accepted: July 6, 2022
Published: August 11, 2022

REFERENCES

1. Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. *24*, 167–202.

2. Koechlin, E., and Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. Trends Cogn. Sci. *11*, 229–235.
3. Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. Science *302*, 1181–1185.
4. Owen, A.M., McMillan, K.M., Laird, A.R., and Bullmore, E. (2005). N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. Hum. Brain Mapp. *25*, 46–59.
5. Hockey, G.R.J. (2011). A motivational control theory of cognitive fatigue. In Cognitive Fatigue: Multidisciplinary Perspectives on Current Research and Future Applications, P.L. Ackerman, ed. (American Psychological Association), pp. 167–187.
6. Baumeister, R.F., Bratslavsky, E., Muraven, M., and Tice, D.M. (1998). Ego depletion: is the active self a limited resource? J. Pers. Soc. Psychol. *74*, 1252–1265.
7. Kahneman, D. (1973). Attention and Effort (Prentice-Hall).
8. Molden, D.C., Hui, C.M., Scholer, A.A., Meier, B.P., Noreen, E.E., D'Agostino, P.R., and Martin, V. (2012). Motivational Versus metabolic effects of carbohydrates on self-control. Psychol. Sci. *23*, 1137–1144.
9. Hagger, M.S., Chatzisarantis, N.L.D., Alberts, H., Anggono, C.O., Batailler, C., Birt, A.R., Brand, R., Brandt, M.J., Brewer, G., Bruyneel, S., et al. (2016). A multilab preregistered replication of the ego-depletion effect. Perspect. Psychol. Sci. *11*, 546–573.
10. Raichle, M.E. (2015). The Brain's default mode network. Annu. Rev. Neurosci. *38*, 433–447.
11. Sokoloff, L. (1977). Relation between physiological function and energy metabolism in the central nervous system. J. Neurochem. *29*, 13–26.
12. Kurzban, R., Duckworth, A., Kable, J.W., and Myers, J. (2013). An opportunity cost model of subjective effort and task performance. Behav. Brain Sci. *36*, 661–679.
13. Inzlicht, M., Schmeichel, B.J., and Macrae, C.N. (2014). Why self-control seems (but may not be) limited. Trends Cogn. Sci. *18*, 127–133.
14. Kool, W., and Botvinick, M. (2018). Mental labour. Nat. Hum. Behav. *2*, 899–908.
15. Boksem, M.A.S., and Tops, M. (2008). Mental fatigue: costs and benefits. Brain Res. Rev. *59*, 125–139.
16. Fava, M., Ball, S., Nelson, J.C., Sparks, J., Konechnik, T., Classi, P., Dube, S., and Thase, M.E. (2014). Clinical relevance of fatigue as a residual symptom in major depressive disorder. Depress. Anxiety *31*, 250–257.
17. Grahek, I., Shenhav, A., Musslick, S., Krebs, R.M., and Koster, E.H.W. (2019). Motivation and cognitive control in depression. Neurosci. Biobehav. Rev. *102*, 371–381.
18. Ackerman, P.L. (2011). Cognitive Fatigue: Multidisciplinary Perspectives on Current Research and Future Applications (American Psychological Association).
19. Dittner, A.J., Wessely, S.C., and Brown, R.G. (2004). The assessment of fatigue: a practical guide for clinicians and researchers. J. Psychosom. Res. *56*, 157–170.
20. Prue, G., Rankin, J., Allen, J., Gracey, J., and Cramp, F. (2006). Cancer-related fatigue: a critical appraisal. Eur. J. Cancer *42*, 846–863.
21. Burke, S.E., Babu Henry Samuel, I., Zhao, Q., Cagle, J., Cohen, R.A., Kluger, B., and Ding, M. (2018). Task-based cognitive fatigability for older adults and validation of mental fatigability subscore of Pittsburgh fatigability scale. Front. Aging Neurosci. *10*, 327.
22. Kim, I., Hacker, E., Ferrans, C.E., Horswill, C., Park, C., and Kapella, M. (2018). Evaluation of fatigability measurement: integrative review. Geriatr. Nurs. *39*, 39–47.
23. Facque, V., Wiehler, A., Volle, E., Mandonnet, E., and Pessiglione, M. (2022). Present bias in economic choice demonstrates increased cognitive fatigability of glioma patients. Cortex *151*, 281–293.
24. Blain, B., Hollard, G., and Pessiglione, M. (2016). Neural mechanisms underlying the impact of daylong cognitive work on economic decisions. Proc. Natl. Acad. Sci. USA *113*, 6967–6972.
25. Blain, B., Schmit, C., Aubry, A., Hausswirth, C., Le Meur, Y., and Pessiglione, M. (2019). Neuro-computational impact of physical training overload on economic decision-making. Curr. Biol. *29*, 3289–3297.e4.
26. Holroyd, C.B. (2015). The waste disposal problem of effortful control. In Motivation and Cognitive Control, T. Braver, ed. (Routledge), pp. 235–260.
27. Öz, G., and Tkáč, I. (2011). Short-echo, single-shot, full-intensity proton magnetic resonance spectroscopy for neurochemical profiling at 4 T: validation in the cerebellum and brainstem. Magn. Reson. Med. *65*, 901–910.
28. Deelchand, D.K., Adanyeguh, I.M., Emir, U.E., Nguyen, T.M., Valabregue, R., Henry, P.G., Mochel, F., and Öz, G. (2015). Two-site reproducibility of cerebellar and brainstem neurochemical profiles with short-echo, single-voxel MRS at 3T. Magn. Reson. Med. *73*, 1718–1725.
29. Genovese, G., Marjańska, M., Auerbach, E.J., Cherif, L.Y., Ronen, I., Lehéricy, S., and Branzoli, F. (2020). In vivo diffusion-weighted MRS using semi-LASER in the human brain at 3 T: methodological aspects and clinical feasibility. NMR in. Biomedicine, e4206.
30. Palombo, M., Shemesh, N., Ronen, I., and Valette, J. (2018). Insights into brain microstructure from in vivo DW-MRS. NeuroImage *182*, 97–116.
31. Vincent, M., Gaudin, M., Lucas-Torres, C., Wong, A., Escartin, C., and Valette, J. (2021). Characterizing extracellular diffusion properties using diffusion-weighted MRS of sucrose injected in mouse brain. NMR Biomed. *34*, e4478.
32. Pfeuffer, J., Tkáč, I., and Gruetter, R. (2000). Extracellular–intracellular distribution of glucose and lactate in the rat brain assessed noninvasively by diffusion-weighted 1 H nuclear magnetic resonance spectroscopy in vivo. J. Cereb. Blood Flow Metab. *20*, 736–746.
33. van der Wel, P., and van Steenbergen, H. (2018). Pupil dilation as an index of effort in cognitive control tasks: a review. Psychon. Bull. Rev. *25*, 2005–2015.
34. Korn, C.W., and Bach, D.R. (2016). A solid frame for the window on cognition: modeling event-related pupil responses. J. Vision *16*, 28.
35. Shenhav, A., Botvinick, M.M., and Cohen, J.D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron *79*, 217–240.
36. Tursky, B., Shapiro, D., Crider, A., and Kahneman, D. (1969). Pupillary, heart rate, and skin resistance changes during a mental task. J. Exp. Psychol. *79*, 164–167.
37. Borderies, N., Bornert, P., Gilardeau, S., and Bouret, S. (2020). Pharmacological evidence for the implication of noradrenaline in effort. PLoS Biol. *18*, e3000793.
38. Piquado, T., Isaacowitz, D., and Wingfield, A. (2010). Pupillometry as a measure of cognitive effort in younger and older adults. Psychophysiology *47*, 560–569.
39. Hess, E.H., and Polt, J.M. (1964). Pupil size in relation to mental activity during simple problem-solving. Science *143*, 1190–1192.
40. Zénon, A., Sidibé, M., and Olivier, E. (2014). Pupil size variations correlate with physical effort perception. Front. Behav. Neurosci. *8*, 286.
41. Gilzenrat, M.S., Nieuwenhuis, S., Jepma, M., and Cohen, J.D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. Cogn. Affect. Behav. Neurosci. *10*, 252–269.
42. de Gee, J.W., Colizoli, O., Kloosterman, N.A., Knapen, T., Nieuwenhuis, S., and Donner, T.H. (2017). Dynamic modulation of decision biases by brainstem arousal systems. eLife *6*, 1–36.
43. Murphy, P.R., O'Connell, R.G., O'Sullivan, M., Robertson, I.H., and Balsters, J.H. (2014). Pupil diameter covaries with BOLD activity in human locus coeruleus. Hum. Brain Mapp. *35*, 4140–4154.
44. Joshi, S., Li, Y., Kalwani, R.M., and Gold, J.I. (2016). Relationships between pupil diameter and neuronal activity in the locus coeruleus, colliculi, and cingulate cortex. Neuron *89*, 221–234.
45. van Vugt, B., van Kerkoerle, T., Vartak, D., and Roelfsema, P.R. (2020). The contribution of AMPA and NMDA receptors to persistent firing in the dorsolateral prefrontal cortex in working memory. J. Neurosci. *40*, 2458–2470.

46. Woodcock, E.A., Anand, C., Khatib, D., Diwadkar, V.A., and Stanley, J.A. (2018). Working memory modulates glutamate levels in the dorsolateral prefrontal cortex during 1H fMRS. Front. Psychiatry *9*, 1–10.
47. Hommel, B., Fischer, R., Colzato, L.S., van den Wildenberg, W.P.M., and Cellini, C. (2012). The effect of fMRI (noise) on cognitive control. J. Exp. Psychol. Hum. Percept. Perform. *38*, 290–301.
48. Bélanger, M., Allaman, I., and Magistretti, P.J. (2011). Brain energy metabolism: focus on astrocyte-neuron metabolic cooperation. Cell Metab. *14*, 724–738.
49. Shulman, R.G., Rothman, D.L., Behar, K.L., and Hyder, F. (2004). Energetic basis of brain activity: implications for neuroimaging. Trends Neurosci *27*, 489–495.
50. Rönnbäck, L., and Hansson, E. (2004). On the potential role of glutamate transport in metal fatigue. J. Neuroinflammation *1*, 1.
51. Petroff, O.A.C. (2002). GABA and glutamate in the human brain. Neuroscientist *8*, 562–573.
52. Zhou, Y., and Danbolt, N.C. (2014). Glutamate as a neurotransmitter in the healthy brain. J. Neural Transm. *121*, 799–817.
53. Tatti, R., Haley, M.S., Swanson, O.K., Tselha, T., and Maffei, A. (2017). Neurophysiology and regulation of the balance Between excitation and inhibition in neocortical circuits. Biol. Psychiatry *81*, 821–831.
54. Popoli, M., Yan, Z., McEwen, B.S., and Sanacora, G. (2011). The stressed synapse: the impact of stress and glucocorticoids on glutamate transmission. Nat. Rev. Neurosci. *13*, 22–37.
55. Moghaddam, B. (1993). Stress preferentially increases extraneuronal levels of excitatory amino acids in the prefrontal cortex: comparison to hippocampus and basal ganglia. J. Neurochem. *60*, 1650–1657.
56. Danbolt, N.C. (2001). Glutamate uptake. Prog. Neurobiol. *65*, 1–105.
57. Bak, L.K., Schousboe, A., and Waagepetersen, H.S. (2006). The glutamate/GABA-glutamine cycle: aspects of transport, neurotransmitter homeostasis and ammonia transfer. J. Neurochem. *98*, 641–653.
58. Tzingounis, A.V., and Wadiche, J.I. (2007). Glutamate transporters: confining runaway excitation by shaping synaptic transmission. Nat. Rev. Neurosci. *8*, 935–947.
59. Vandenberg, R.J., and Ryan, R.M. (2013). Mechanisms of glutamate transport. Physiol. Rev. *93*, 1621–1657.
60. Ip, I.B., Berrington, A., Hess, A.T., Parker, A.J., Emir, U.E., and Bridge, H. (2017). Combined fMRI-MRS acquires simultaneous glutamate and BOLD-fMRI signals in the human brain. NeuroImage *155*, 113–119.
61. Pochon, J.B., Levy, R., Fossati, P., Lehericy, S., Poline, J.B., Pillon, B., Le Bihan, D., and Dubois, B. (2002). The neural system that bridges reward and cognition in humans: an fMRI study. Proc. Natl. Acad. Sci. USA *99*, 5669–5674.
62. Volk, C., Jaramillo, V., Merki, R., O'Gorman Tuura, R., and Huber, R. (2018). Diurnal changes in glutamate + glutamine levels of healthy young adults assessed by proton magnetic resonance spectroscopy. Hum. Brain Mapp. *39*, 3984–3992.
63. Xie, L., Kang, H., Xu, Q., Chen, M.J., Liao, Y., Thiyagarajan, M., O'Donnell, J., Christensen, D.J., Nicholson, C., Iliff, J.J., et al. (2013). Sleep drives metabolite clearance from the adult brain. Science *342*, 373–377.
64. Leone, S.S., Wessely, S., Huibers, M.J.H., Knottnerus, J.A., and Kant, I. (2011). Two sides of the same coin? On the history and phenomenology of chronic fatigue and burnout. Psychol. Health *26*, 449–464.
65. Bianchi, R., Schonfeld, I.S., and Laurent, E. (2015). Burnout–depression overlap: a review. Clin. Psychol. Rev. *36*, 28–41.
66. Samuelson, P.A. (1937). A note on measurement of utility. Rev. Econ. Stud. *4*, 155–161.
67. Peters, J., and Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. J. Neurosci. *29*, 15727–15734.
68. Chong, T.T.J., Apps, M.A.J., Giehl, K., Hall, S., Clifton, C.H., and Husain, M. (2018). Computational modelling reveals distinct patterns of cognitive and physical motivation in elite athletes. Sci. Rep. *8*, 11888.
69. Soutschek, A., and Tobler, P.N. (2020). Causal role of lateral prefrontal cortex in mental effort and fatigue. Hum. Brain Mapp. *41*, 4630–4640.
70. Daunizeau, J., Adam, V., and Rigoux, L. (2014). VBA: A probabilistic treatment of nonlinear models for neurobiological and behavioural data. PLoS Comp. Biol. *10*, e1003441.
71. Gruetter, R., and Tkáč, I. (2000). Field mapping without reference scan using asymmetric echo-planar techniques. Magn. Reson. Med. *43*, 319–323.
72. Provencher, S.W. (1993). Estimation of metabolite concentrations from localized in vivo proton NMR spectra. Magn. Reson. Med. *30*, 672–679.
73. K.J. Friston, ed. (2007). Statistical Parametric Mapping: the Analysis of Funtional Brain Images, First Edition (Elsevier/Academic Press).
74. Tkác, I., Starcuk, Z., Choi, I.Y., and Gruetter, R. (1999). In vivo 1H NMR spectroscopy of rat brain at 1 ms echo time. Magn. Reson. Med. *41*, 649–656.
75. de Brouwer, H. (2009). Evaluation of algorithms for automated phase correction of NMR spectra. J. Magn. Reson. *201*, 230–238.
76. Henry, P.-G., Marjanska, M., Walls, J.D., Valette, J., Gruetter, R., and Ugurbil, K. (2006). Proton-observed carbon-edited NMR spectroscopy in strongly coupled second-order spin systems. Magn. Reson. Med. *55*, 250–257.
77. Govindaraju, V., Young, K., and Maudsley, A.A. (2000). Proton NMR chemical shifts and coupling constants for brain metabolites. NMR Biomed. *13*, 129–153.
78. Kaiser, L.G., Marjańska, M., Matson, G.B., Iltis, I., Bush, S.D., Soher, B.J., Mueller, S., and Young, K. (2010). 1H MRS detection of glycine residue of reduced glutathione in vivo. J. Magn. Reson. *202*, 259–266.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and algorithms		
Matlab R2019a	Mathworks	SCR_001622
Psychophysics Toolbox	http://psychtoolbox.org/	SCR_002881
SPM12	FIL, London	SCR_007037
Lcmodel	http://s-provencher.com/lcmodel.shtml	SCR_014455
Deposited data	https://zenodo.org/record/6795446	https://doi.org/10.5281/zenodo.6795446

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Antonius Wiehler (antonius.wiehler@gmail.com).

Materials availability

No new materials have been generated.

Data and code availability

The data are available at https://zenodo.org/record/6795446.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Participants

We included a total of n = 40 participants, n = 24 in the high-demand condition (age mean m = 22.21 y, standard deviation SD = 6.91, 13 female), n = 16 in low-demand condition (age m = 24.56 y, SD = 6.12, 10 female). The two groups were matched regarding impulsivity trait (as measured by Barratt's Impulsiveness Scale BIS-11) and baseline fatigue state (as measured by Brief and Multidimensional Fatigue Inventory questionnaires, BFI and MFI-20). All participants gave informed consent before participating in the study. All participants were screened for exclusion criteria: left-handedness, age under 20 or above 39 y, regular use of drugs or medication, any history of psychiatric or neurological diagnosis, and contraindications to MRI scanning (pregnancy, claustrophobia, metallic implants). In both conditions, 2 additional participants stopped the experiment before completion and were therefore excluded from the data analysis. All participants ate a sandwich and a fruit during the first break. Water was the only allowed drink during the day and was available without restrictions. Participants received 50€ as financial compensation for the two training sessions. For their performance on the experiment day, they received 5€ plus another 3€ for each percent above 75% in their average performance in the cognitive tasks, which would result in 50€ for an average of 90% correct responses maintained throughout the day. Additionally, one choice trial in each domain (probability, delay, effort) was pseudo-randomly chosen and implemented. The study was approved by the local ethics committee of the Pitié-Salpêtrière Hospital (CPP no 113-15, ID RCB: 2015-A01445-44).

METHOD DETAILS

Tasks

Two cognitive control tasks were used to induce cognitive fatigue: N-switch and N-back. In each trial of both tasks, a letter appeared on the screen, colored either red or green. Participants had to give their response within a 0.8s time window, followed by a 0.8 s inter-trial interval. During N-switch blocks, participants had to perform a discrimination task that depended on the color of the letter: upper case vs. lower case for one color, vowel vs. consonant for the other color (red or green, counterbalanced across participants). The task was switched 12 times per block of 24 trials in the hard condition, whereas it was switched only once in the easy condition. During N-back blocks, participants had to indicate whether the letter on the screen was the same as the letter presented in three trials (hard condition) or one trial (easy condition) before.

To reveal cognitive fatigue, we presented four choice trials after each block of 24 task trials. The time out for choice trials was 3.25 s, followed by a jittered inter-trial interval (mean 1.25 s, SD = 0.33). Each choice trial opposed a small-reward/low-discount

option with a big-reward/high-discount option. Rewards ranged from 0.1€ to 50€. All rewards were presented with 2 digits precision but rounded to the first digit: for example, 41.23€ had been rounded to 41.20€. Discount factors were of four different types: delay, probability, cognitive effort, and physical effort. Delays ranged from zero (reward received in cash immediately after the experiment) to one year (reward received by bank transfer), with intermediate levels of three days, one week, one month, and three months). The probability of winning the lottery (vs. nothing) ranged from 5% to 100% with intermediate levels of 25%, 50%, 75%, and 95%. The cognitive effort consisted of performing the N-switch task for 30 minutes after the main experiment. Effort levels corresponded to the number of switches in a 24-trial block: 0, 2, 4, 6, 9, or 12. The physical effort consisted of pedaling on a stationary bike for 30 minutes after the main experiment. Effort levels corresponded to the resistance of the bike, expressed in percentage of the maximum power that the participant could develop during calibration: 0%, 12.5%, 25%, 37.5%, 56.25%, 75%. Participants were instructed that one randomly selected choice of each type would be realized, meaning that money would be given to the participant but only after the chosen delay / playing the chosen lottery / exerting the chosen effort level.

All participants completed two sessions of training on the day before the main experiment. Each training session started with the easy version of the cognitive control tasks and gradually increased in difficulty until performance reached 90% correct responses at the highest difficulty level. During the second training session, participants were also instructed about the choice tasks, and they practiced with a test set of choices to get familiarized.

On the day of the main experiment, participants first rehearsed the cognitive control tasks to ensure their performance was still above 90% correct responses. Then they underwent a choice calibration procedure. For all choice trials, the big reward was fixed to 50€, while the small reward was associated with either zero cost (e.g., 0 days for delay) or the lowest cost (e.g., 3 days). For each discounting domain and cost level separately, the size of the small reward was adjusted with a staircase procedure depending on the choice of the participant. Had the participant chosen the low-cost option twice in a row, the small reward was reduced to the mean between its current value and that of the last rejected low-cost option (0€ for the first trial). Had the participant rejected the low-cost option twice in a row, the small reward was increased to the mean between its current value and that of the last accepted low-cost option (50€ for the first trial). If the low-cost option was accepted/rejected only once, the small reward value was reduced/increased by 10%. The staircase procedure stopped when the difference between accepted and rejected small rewards was smaller than 4€. The mean between the last-rejected and last-accepted small reward was taken as the indifference point for each cost level of a given domain. On average, participants made 15.47 [7-108] choices to reach the indifference point.

To reduce noise in these estimates, the calibration procedure was repeated three times and indifference points were averaged, separately for every cost level. In total, 36 indifference points were estimated, corresponding to four choice domains times nine cost levels (five opposed to zero discount, e.g., one year vs. 0 days, and four opposed to lowest discount, e.g., one year vs. 3 days).

In the main experiment, for every cost level, we tailored choice options around participant-specific indifference points, with five trials presenting small rewards at the indifference point (drawn from a normal distribution centered at the indifference point with SD=1), one trial with the small reward 30% below and one trial with the small reward 30% above the indifference point. The choice trials close to indifference were meant to maximize the sensitivity for detection of a preference shift, while distant trials were meant to ensure that computational models could be fitted with recoverable parameters. The consequence of tailoring choice options was that participants started the experiment with an average choice rate close to 50%, leaving room for either decrease or increase with fatigue. Under these constraints, small rewards were randomly drawn such that each of the 5 sessions presented novel choices. This resulted in a total of 5 sessions × 4 choice types × (4+5 cost levels) × (5+1+1 trials) = 1260 choice trials per participant.

At the beginning of each session, participants rated their subjective fatigue by positioning a cursor on a visual analog scale between 0 (“I’m in top form”) and 100 (“I’m totally exhausted”).

QUANTIFICATION AND STATISTICAL ANALYSIS

Computational modeling

To analyze choice behavior, we used a computational modeling approach. A priori, we favored multiplicative forms for delay and probability discounting, because a delayed or probabilistic reward must be positive (always better than nothing) and subtractive forms for effort discounting because the reward may not be worth the cost (so the value can be negative). We nonetheless used Bayesian model comparison to identify the discounting functions that provided the best account of choices made during calibration for each cost domain. The most plausible discounting functions were indeed exponential for delay (as in our previous studies), hyperbolic for probability, and parabolic for effort (see equations below).

In all models, rewards were first discounted with different functions depending on the considered factor (D: delay in days, P: winning probability, E: effort level) to generate subjective values (SV):

Delay discounting:[66]

$$SV = Re^{-kD} \quad \text{(Equation 1)}$$

Probability discounting:[67]

$$SV = \frac{R}{1 + k\frac{1-P}{P}} \quad \text{(Equation 2)}$$

Cognitive and physical effort discounting:[68,69]

$$SV = R - kE^2 \quad \text{(Equation 3)}$$

with the free parameter k controlling the steepness of discounting in each function. Then subjective values of the low-cost (LC) and high-cost (HC) options are compared and transformed into the probability of choosing the low-cost option (p_{LC}) through a softmax function:

$$p_{LC} = \frac{1}{1 + e^{-(\beta(SV_{LC} - SV_{HC}) + bias)}} \quad \text{(Equation 4)}$$

With the inverse temperature parameter β controlling choice consistency with subjective values and the *bias* parameter shifts the choice probability towards the low-cost option.

The different models were inverted using a variational Bayes approach under the Laplace approximation, implemented in the VBA toolbox[70] (available at https://mbb-team.github.io/VBA-toolbox) programmed in Matlab R2019a (MathWorks, Natick, MA, USA). We first modeled the calibration choices in each participant to build informed priors for the main experiment. Then, using the individual priors, we estimated all posterior parameters independently for each participant, session, and choice type.

Pupil dilation

During sessions one, three, and five, we recorded pupil size with an Eyelink 1000 eye-tracker (SR Research, right eye recorded, 1000 Hz sampling rate). Time series were pre-processed by removing blink periods, fixations outside the screen, and samples outside three times the median signal. Removed samples were replaced by linear interpolation. Time series were then band-pass filtered between 1/128 Hz and 1 Hz, down-sampled to 60 Hz, z-scored across all sessions per participant, epoched around choice trials, and corrected for baseline differences across trials (by subtraction of the mean pupil size within the 500 ms before choice onset). Within-trial pupil size time course was analyzed using a sample-wise regression approach: for each time sample, pupil size was regressed against session number across trials. To remove potential confounds with low-level variables, the linear regression model also included low-cost choice (1 or 0), low-cost choice in the trial before (1 or 0), the choice type (probability, delay, cognitive effort, physical effort), type of low-cost option (zero or non-zero), the inter-trial interval before choice onset, the cognitive task performance for the current block, the block number and the trial number within the block. The regression was run at the individual level and regression estimates were tested against 0 at the group level with a 1-D random field theory implementation in the VBA toolbox.[70]

MRS data acquisition

Magnetic Resonance Spectroscopy (MRS) was performed on a 3 T Siemens MAGNETOM Prisma Fit MRI scanner (Siemens Medical Solutions, Erlangen, Germany), equipped with gradient coils capable of reaching 80 mT/m on each of the three axes. The standard RF body coil was used for excitation and a 64-channel receive-only head coil for reception.

Participants alternated between performing the tasks inside the scanner (sessions 1, 3, and 5) and outside the scanner (sessions 2 and 4). Before every new scanning session, the spectroscopic volumes of interest (VOI) were manually placed to maximize the overlap with those of the previous session. To precisely position the VOI and to perform tissue segmentation, the MRS protocol was preceded by a 3D T1-weighted magnetization-prepared rapid gradient echo sequence [field of view = 256 x 256; isotropic resolution = 1 mm; TR/TE = 2300/4.18 ms; total acquisition time = 4 min 44 s]. A VOI of 35×25×15 mm^3 was positioned in the lateral prefrontal cortex. The size of the VOI was adapted to cover most of the activation cluster observed at the group level in our fMRI study[24] while respecting the block shape imposed by the MRS sequence and a minimum volume which was needed to reach an acceptable signal-to-noise ratio. First, on the axial slice, the VOI, with an anterior-posterior expansion of 35 mm, was placed on the middle frontal gyrus. We used the posterior border of the triangular part of the inferior frontal gyrus as the posterior reference. The VOI was then rotated to follow the cortex orientation and to include as little CSF as possible. Second, on the coronal slice, the VOI, which spanned 25 mm, was again placed to cover the middle frontal gyrus and was rotated to follow the cortex orientation and to avoid including CSF. A control VOI of the same size was positioned over the primary visual cortex.

For the MRS data acquisition, we used a modified single-voxel semi-LASER sequence[27,28] (TR/TE = 5000/28 ms; number of complex points = 2048; averages = 64; total acquisition time = 5 min 30 s). B0 shimming in the VOIs was performed using a fast automatic shimming technique with echo-planar signal trains utilizing mapping along projections, FAST(EST)MAP.[71] Before the MRS acquisition, the RF power for the asymmetric slice-selective 90° pulse (duration, 2 ms) of the semi-LASER sequence was optimized to produce the maximum signal. This was in turn used to automatically adjust the power of the 180° hyperbolic secant adiabatic full passage pulses (duration, 4 ms). Water suppression was performed using variable power with optimized relaxation delays (VAPOR) and suppression of signal contamination from other brain regions was achieved with outer volume suppression (OVS). In addition, unsuppressed water spectra were acquired for eddy current corrections. The contribution of cerebrospinal fluid (CSF) to the VOI was corrected by segmenting the brain and estimate the proportion of CSF in the VOI.

MRS data analysis

All spectra were processed in MATLAB R2019a (MathWorks, Natick, MA, USA). Eddy currents and shot-to-shot phase and frequency correction were performed as described previously.[28] LCModel[72] (Version 6.3.0-G) was used for the quantification of metabolite

concentrations. The basis set was simulated using the density matrix formalism and included alanine, ascorbate, aspartate, creatine, γ-aminobutyric acid, glycerophosphorylcholine, phosphocholine, phosphocreatine, glucose, glutamine, glutamate, glutathione, myo-inositol, scyllo-inositol, lactate, N-acetylaspartate, N-acetylaspartylglutamate, phosphorylethanolamine, and taurine, as well as macromolecule spectra that were acquired in healthy volunteers for a previous study.[28] Metabolite quantification was considered reliable for Cramér Rao Lower Bounds (CRLB) < 20%, so we excluded data for 10% of the participants. According to this criterion, the following metabolites were reliably quantified: total creatine (tCr), myo-inositol (Ins), total N-acetylaspartate (tNAA), total choline (tCho), glutathione (GSH), glycerophosphocholine (GPC), glutamine (Gln) and glutamate (Glu).

To account for local inhomogeneities in the VOIs, as well as reduce the high inter-participant variability in the absolute concentration due to the variable amounts of CSF, gray and white matter in the VOIs, we report metabolic concentration as a ratio to the baseline peaks of Ins, tNAA, and tCr. To control for potential confounds, we included the following measures as covariates in the statistical analysis: age, linewidth (as estimated by LCModel), signal-to-noise ratio (as estimated by LCModel), grey matter concentration in the voxel (estimated by a Freesurfer segmentation of the VOI), and movement regressors which were estimated in SPM12[73] based on pre/post-measurement EPI images.

Diffusion-weighted magnetic resonance spectroscopy acquisition

Diffusion-weighted magnetic resonance spectroscopy (DW-MRS) was performed in the same VOIs as for the MRS acquisitions using a single-voxel semi-LASER sequence with diffusion gradients added in a bipolar configuration (TE = 120 ms, spectral width = 3 kHz, number of complex points = 2048).[29] All resonances were excited using a slice-selective 90° pulse (pulse length of 2.52 ms) followed by two pairs of slice-selective adiabatic refocusing pulses in the other two dimensions (HS1, R = 20, pulse length 7.5 ms). All acquisitions were synchronized with the cardiac cycle using a pulse-oximeter device, to start each acquisition every three heartbeats, while maintaining a minimum TR of 2500ms. Diffusion-weighting was applied in three orthogonal directions ([1, 1, -0.5], [1, -0.5, 1], [-0.5, 1, 1] in the VOI coordinate system) with diffusion gradient duration = 22 ms, diffusion time = 60 ms and four increasing gradient strengths g = 0, 19, 39, 58 mT/m, resulting in the b-values b0 = 0, b1 = 1080, b2 = 4300 and b3 = 9770 s/mm^2.

Sixteen averages were collected for each diffusion-weighting condition and saved as individual free induction decays for further post-processing. Water suppression was performed using VAPOR and OVS suppression.[74] For eddy current corrections, unsuppressed water reference scans were acquired from the same VOIs using the same parameters as water suppressed spectra.

DW-MRS spectral processing

DW-MRS data were corrected as described previously.[29] Eddy current corrections were performed for each DW condition using water reference scans. Zero-order phase fluctuations and frequency drifts were corrected on single averages before summation using an area minimization and penalty algorithm and a cross-correlation algorithm, respectively.[75] A peak-thresholding procedure was applied, for each DW condition, to discard the single spectra with artifactually low SNR caused by non-translational tissue motion.[29] A threshold of 70% for the proportion of spectra rejected per DW condition was set. No datasets were excluded using this criterion. The remaining spectra, for each condition, were then averaged.

The averaged spectra were analyzed with LCModel for metabolite quantification. The basis set was simulated based on the density matrix formalism[76] and using chemical shifts and J-couplings reported previously.[77,78] The basis set included alanine, ascorbate, aspartate, creatine, γ-aminobutyric acid, glucose, glutamate, glutamine, glutathione, glycerophosphorylcholine, mIns, lactate, N-acetylaspartate, N-acetylaspartylglutamate, phosphocreatine, phosphorylcholine, phosphorylethanolamine, scyllo-inositol, and taurine. Independent spectra for the CH3 and CH2 groups of NAA, Cr, and PCr were simulated and included in the basis set. Metabolite quantification was considered reliable for CRLB < 20% at b0, and no data were excluded due to this criterion. Signal/noise was > 10 at all b-values.

ADCs were calculated for the metabolites fitted by LCModel. Metabolite ADCs for tNAA, tCr, tCho, Glx, and Ins were computed in each VOI and session by fitting a stretched exponential to the logarithmic metabolite signal decay:

$$\log\left(\frac{S_{met,b}}{S_{met,0}}\right) = -\alpha + (-ADC * b)^{\gamma} \quad \text{(Equation 5)}$$

Were $S_{met,b}$ is the metabolite signal at a given b-value, $S_{met,0}$ is the metabolite signal at b_0, *ADC* is the apparent diffusion coefficient (scaling factor) and γ is the stretching factor. To control for potential confounds in the statistical analysis of *ADC*, we added the following measures as covariates: age, linewidth and signal/noise as estimated by LCModel, grey matter concentration in the voxel, and the proportion of spectra rejected during preprocessing.

Metabolic accumulation model

To test the link between neural activity measured using fMRI in our previous study using the same behavioral protocol[24] and the present spectroscopy measures, we developed a metabolic accumulation model (see Figure 6), following a Markov chain:

$$Glu_t = Glu_{t-1} - \beta_{clear} Glu_{t-1} + \beta_{task} BOLD_{task,t} + \beta_{choice} BOLD_{choice,t} \quad \text{(Equation 6)}$$

The model assumes that glutamate accumulates at a rate (Glu_t - Glu_{t-1}) that is proportional to instantaneous neural activity (proxied by BOLD signal measured with fMRI during the cognitive tasks and economic choice) and dissipates at a rate that is proportional to its current level. To generate the inputs to the dynamic model, we averaged BOLD activity measured for each session and group in our

previous study.[24] The BOLD activity was then upsampled to 22 sessions, normalized between 0 and 1, and smoothed with a moving average, before entering into the dynamic model. We then fitted both Glu concentration (mean of Glu/tCr, Glu/Ins, and Glu/tNAA) and Glx diffusion (ADC) measures, separately, as if they had been measured during sessions 2, 12, and 22, to match the timing of data acquisition in the two studies. The model was inverted to estimate posterior parameters of the three scaling factors (β_{clear}, β_{task}, and β_{choice}) using the VBA toolbox.[70]

Statistical analysis

All statistical analyses were performed in MATLAB R2019a (MathWorks, Natick, MA, USA) with linear mixed models (function *fitglme*). Intercepts and all within-participant factors (e.g., session) were estimated on the participant level. All between-participant factors (e.g., condition) were estimated at the group level. Response times (RTs) during choice were log-transformed to correct a skewed distribution and we excluded trials with short RTs of <0.1 s.