Need for cognition does not account for individual differences in metacontrol of decision making

Florian Bolenz(1), Maxine F. Profitt(2), Fabian Stechbarth(1), Ben Eppinger(1,2,3) & Alexander Strobel(1)

1: Faculty of Psychology, Technische Universität Dresden, Germany

2: Department of Psychology, Concordia University, Canada

3: PERFORM centre, Concordia University, Canada

Keywords: arbitration, reinforcement learning, cognitive effort

Conflict of interest: The authors declare no competing interests

Acknowledgements: This work was supported by grants from the German Research Foundation (DFG) (SFB 940/2 B6; SFB 940/2 B7) as well as the Natural Sciences and Engineering Research Council of Canada (NSERC) (N01882). Furthermore, this research was undertaken, in part, thanks to funding from the Canada Research Chairs program.

Word count: 2,961

Corresponding author:

Florian Bolenz (florian.bolenz@tu-dresden.de)

Technische Universität Dresden

Fakultät Psychologie

01062 Dresden

Germany

Abstract

Humans show metacontrol of decision-making towards different reward magnitudes. Specifically, when higher rewards are at stake, individuals increase reliance on a more accurate but cognitively effortful strategy. We investigated whether the personality trait Need for Cognition (NFC) explains individual differences in metacontrol. Based on findings of cognitive effort expenditure in executive functions, we expected more metacontrol in individuals low in NFC. In two independent studies, metacontrol was assessed by means of a decision-making task that dissociates different reinforcement-learning strategies. In contrast to our expectations, NFC did not account for individual differences in metacontrol of decision making. These findings suggest a differential role of NFC for the regulation of cognitive effort in decision making and executive functions.

Introduction

Human decision making can be guided by different strategies (Daw, Gershman, Seymour, Dayan, & Dolan, 2011). Individuals show metacontrol of decision-making strategies, that is they adapt their reliance on these strategies to situational demands (Kool, Gershman, & Cushman, 2017, 2018). The framework of Reinforcement Learning (Dayan & Niv, 2008; Sutton & Barto, 1998) discriminates between at least two strategies. First, a *model-free strategy* guides decisions based on previously learnt action-reward associations. This strategy is considered to be computationally simple but can be inaccurate in dynamically changing environments. Second, a *model-based strategy* guides decisions by using a mental model of the environment to predict the consequences of potential decisions. A model-based strategy is usually more accurate, but requires more cognitive effort because it relies at least partly on executive functions (Otto, Gershman, Markman, & Daw, 2013; Otto, Skatova, Madlon-Kay, & Daw, 2015).

Previous studies have shown considerable individual differences in metacontrol of decision-making strategies, in that people differ in how strongly they adapt decision-making strategies toward different reward magnitudes. These individual differences can only partly be accounted for by age or subclinical psychopathological personality traits (Bolenz, Kool, Reiter, & Eppinger, 2019; Patzelt, Kool, Millner, & Gershman, 2018). Here, we aim at extending our understanding of individual differences in metacontrol by asking whether the personality trait Need for Cognition (NFC; Cacioppo & Petty, 1982) explains how much people engage in metacontrol of decision-making strategies.

NFC reflects an individual's intrinsic motivation for cognitively demanding activities (Cacioppo, Petty, Feinstein, Blair, & Jarvis, 1996). Previous research has shown that individuals high in NFC need less monetary incentives to exert cognitive effort than individuals low in NFC (Thompson, Chaiken, & Hazlewood, 1993; Westbrook, Kester, & Braver, 2013). In line with these findings, a study by Sandra and Otto (2018) reported that NFC modulates how rewards affect effort expenditure for executive functions. Specifically, while individuals low in NFC invested more cognitive effort in a high-reward condition compared to a low-reward condition of a task-switching paradigm, this reward-induced effect vanished for individuals high in NFC. This suggests that the adaptation of cognitive effort toward different payoffs is more pronounced when the intrinsic motivation to engage in cognitively effortful behavior is low.

Based on these findings by Sandra and Otto (2018), we expect that NFC explains individual differences in metacontrol of decision-making strategies. Specifically, individuals low in NFC should show more metacontrol of decision making (i.e. a stronger upregulation of model-based decision making when rewards are amplified) than individuals high in NFC. We therefore hypothesize that NFC negatively correlates with reward-related metacontrol of decision

making. We tested this prediction in datasets from two studies employing a decision-making task that has been developed to dissociate model-free from model-based strategies (Kool, Cushman, & Gershman, 2016; Kool et al., 2017).

Study 1

Methods

For Study 1, we report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study (cf. Simmons, Nelson, & Simonsohn, 2012). The dataset and all analysis scripts can be found at osf.io/9wc4u.

Participants

128 participants took part in this study. Sample size was determined based on considerations that it would take around 120 participants to detect a true correlation of $r = .25$ with a power of $1 - \beta = .80$. Data from a subset of this sample ($N = 63$) were already presented as part of a previous study (Bolenz et al., 2019) and for this subsample, we assessed additional measures such as cognitive control and processing speed that are not reported here. We excluded two participants from data analysis due to the following reasons: missing responses in more than 20% of trials in the decision-making task (1) and missing values in the NFC scale (1). Thus, the effective sample consisted of 126 participants (76 female, age range: 18 – 36 years, mean age = 23.4 years). All participants gave informed written consent and received either monetary compensation (5€/hour) or course credit for their participation, as well as an additional monetary compensation related to their performance in the decision-making task (10 cents for every 60 points in the decision-making task). The ethics committee of Technische Universität Dresden approved the study.

Figure 1. The sequential decision-making task. (A) Task transition structure. Each trial offered the choice between one pair of spaceships, both leading deterministically to one of two planets. At the planet, a reward was obtained and the amount of reward slowly drifted over the course of the task. (B) Trial structure. At the beginning of each trial, a stakes condition was cued. Low-stakes trials and high-stakes trials differed in how rewards were converted into points. (C) Transition conditions (only in Study 1). In stable-transitions blocks, the task transition structure remained unchanged throughout the block of 80 trials. In variable-transitions blocks, every 6 to 14 trials, the pair of spaceships in one first-stage state swapped their destination planets. This figure was reprinted from Bolenz et al. (2019) and is licensed under [CC BY 4.0](#).

Decision-making task

Reliance on decision-making strategies was assessed with a sequential decision-making task that dissociates model-free and model-based decision making (Bolenz et al., 2019; Kool et al., 2016; cf. Figure 1). In this task, participants collect rewards (“space treasure”) by choosing between spaceships that travel to different planets. Across trials, we manipulated how rewards were converted into points (cf. Kool et al., 2017). In *low-stakes trials*, participants received one point for every piece of space treasure; in *high-stakes trials*, participants received five points for every piece of space treasure (Figure 1B).

Moreover, we manipulated task complexity by imposing additional demands on structure learning in some trial blocks. In *stable-transitions blocks*, spaceships maintained their

destinations throughout the entire block, while in *variable-transitions blocks*, one pair of spaceships switched their destinations every 6 to 14 trials (Fig. 1C). More details on the decision-making task are reported in the supplementary information.

Reinforcement-learning model

We used an established hybrid reinforcement-learning model (Bolenz et al., 2019; Daw et al., 2011; Gläscher, Daw, Dayan, & O'Doherty, 2010) to determine individual reliance on model-based decision making in the decision-making task. This model integrates model-free reward expectations $Q_{MF}(s,a)$ and model-based reward expectations $Q_{MB}(s,a)$:

$$Q(s, a) = (1 - \omega) \times Q_{MF}(s, a) + \omega \times Q_{MB}(s, a)$$

The model-based weight ω (ranging between 0 and 1) reflects the relative influence of the model-based learner with higher values of ω representing relatively more reliance on model-based decision making. Our model includes four model-based weights, one for low-stakes trials and one for high-stakes trials in both stable-transitions blocks and variable-transitions blocks. We defined metacontrol as the difference between the model-based weight for high-stakes trials and the model-based weight for low-stakes trials, such that a higher difference value reflects a stronger increase of the model-based weight in high-stakes trials. More information on the reinforcement-learning model and the model-fitting procedure are reported in the supplementary information.

Figure 2. Distribution of NFC scores in Study 1. NFC was assessed with the German short scale, consisting of 16 items that are recorded on a 7-point Likert scale (-3 to 3).

NFC scale

We assessed NFC with the German NFC short-scale (Bless, Fellhauer, Bohner, & Schwarz, 1994). This scale consists of 16 items that are recorded on a 7-point Likert scale, ranging from -3 (“totally disagree”) to +3 (“totally agree”). In our sample, NFC scores ranged between -26 and 47 (Figure 2; mean = 14.63, sd = 13.48). Internal consistency was $\alpha = 0.88$.

Data analysis

For quantifying the evidence regarding our hypotheses, we computed Bayes Factors using the BayesFactors package in R (Morey, Rouder, & Jamil, 2015). Bayes Factors reflect how much more likely it is to observe some data under the assumption of the alternative hypothesis than under the assumption of the null hypothesis. Bayes Factors ranging between 3 and 10 are commonly interpreted as providing moderate evidence and Bayes Factors above 10 are interpreted as providing strong evidence for the alternative hypothesis. Conversely, Bayes Factors ranging between 1/3 and 1/10 are interpreted as providing moderate evidence and Bayes Factors below 1/10 are interpreted as providing strong evidence for the null hypothesis (Lee & Wagenmakers, 2013). We report priors and a Bayesian analogue to power analysis in the supplementary information.

Results and discussion

Figure 3. Metacontrol of decision making in Study 1. (A) Mean model-based weights. Error bars represent standard error of the mean. (B-C) Individual differences in model-based weights. Plots show model-based weights in high-stakes trials (y-axis) against model-based weights in low-stakes trials (x-axis) for stable-transitions blocks (B) and variable-transitions blocks (C). Points on the identity line represent individuals that showed no adaptation of model-based weights toward stakes conditions. Points above (below) the identity line represent individuals that showed higher (lower) model-based weights in high-stakes trials compared to low-stakes trials.

Metacontrol of decision making

We found increased model-based weights for high-stakes trials compared to low-stakes trials (Figure 3A), both in stable-transitions blocks ($BF_{10} = 6.9 \times 10^6$) and in variable-transitions blocks ($BF_{10} = 19.3$). In line with results previously reported in a subsample of our participants and in other studies (Bolenz et al., 2019; Kool et al., 2017), this indicates that participants increased reliance on model-based decision making when rewards were amplified.

Figure 4. Relationship between NFC (x-axis) and metacontrol of decision making (y-axis) in Study 1.

Need for Cognition and metacontrol

We found considerable individual differences in how strongly participants adapted decision-making strategies in response to amplified rewards (Figure 3B-C). Therefore, we next analyzed whether NFC explains individual differences in metacontrol of decision-making. We found moderate evidence against a correlation between NFC and metacontrol (Figure 4) for both stable-transitions blocks ($r = 0.00$, $BF_{10} = 0.11$) and variable-transitions blocks ($r = 0.12$, $BF_{10} = 0.28$). Based on findings showing stronger reward-induced increases in executive function performance for individuals low in NFC (Sandra & Otto, 2018), we also tested the hypothesis of a negative correlation between NFC and metacontrol. We found moderate evidence against a negative correlation in stable-transitions block ($BF_{10} = 0.11$) and strong evidence against a negative correlation in variable-transitions blocks ($BF_{10} = 0.05$). These results were largely unaffected by different prior assumptions about the population correlation (see Figure S1 in the supplementary information). These findings speak against the idea that individuals low in NFC show more metacontrol of decision making.

Need for Cognition and model-based decision making

We also investigated whether NFC explains individual differences in model-based decision making. There was moderate evidence against a correlation between NFC and model-based weights in low-stakes trials for both stable-transitions blocks ($r = 0.09$, $BF_{10} = 0.18$) and variable transitions-blocks ($r = 0.07$, $BF_{10} = 0.15$). There was inconclusive evidence regarding a correlation between NFC and model-based weights for high-stakes trials in both stable-transitions blocks ($r = 0.13$, $BF_{10} = 0.34$) and variable-transitions blocks ($r = 0.21$, $BF_{10} = 1.73$). Thus, our results do not show that NFC explains individual differences in model-based decision making.

Study 2

To scrutinize the unexpected findings in Study 1, we tried to replicate these findings in an independent, larger sample. For Study 2, we report how we determined our sample size, all data exclusions and all manipulations (cf. Simmons et al., 2012). The dataset and all analysis scripts can be found at osf.io/9wc4u.

Methods

Participants

214 participants took part in this study. Sample size was determined based on feasibility considerations. Participants completed the decision-making task and the NFC scale as part of a larger task battery, the results of which will be reported elsewhere. We excluded nine participants from data analysis due to the following reasons: missing responses in more than 20% of trials in the decision-making task (1), key repetitions in more than 95% of trials in the decision-making task (1), no or incomplete recording of the decision-making task due to technical difficulties (3), duplicate assessments with the NFC scale (4). Thus, the effective sample consisted of 205 participants (149 female, age range: 18 – 32 years, mean age = 22.0 years). All participants gave informed written consent and received monetary compensation (\$25 or course credit as baseline, 22 cents for every 100 points in the decision-making task). The Human Research Ethics Committee at Concordia University approved the study.

Decision-making task

We used an adapted version of the decision-making task that was employed in Study 1. Similar to Study 1, we manipulated how rewards were converted into points, with each piece of reward being worth one point in low-stakes trials and being worth five points in high-stakes trials. In contrast to Study 1, the spaceships kept their destination planets throughout the entire task, i.e., the complete task corresponded to the stable-transitions condition from Study 1. Details on the task are reported in the supplementary information.

Reinforcement-learning model

Individual reliance on model-based decision making was identified by means of the same reinforcement-learning model as in Study 1. Because there was no change of the transition structure, this model only contained two model-based weights (one for low-stakes trials and one for high-stakes trials).

Figure 5. Distribution of NFC scores in Study 2. NFC was assessed with the English short scale, consisting of 18 items that are recorded on a 5-point Likert scale (1 to 5).

NFC scale

We assessed NFC with the English NFC short-scale (Cacioppo et al., 1996; Cacioppo, Petty, & Kao, 1984). This scale consists of 18 items that are recorded on a 5-point Likert scale, ranging from 1 (“extremely uncharacteristic”) to 5 (“extremely characteristic”). In our sample, NFC scores ranged between 39 and 88 (Figure 5; mean = 64.81, sd = 9.68). Internal consistency was $\alpha = .84$.

Data analysis

We computed Bayes Factors to quantify the evidence regarding our hypotheses. Based on our findings in Study 1, more prior weight was given to correlations closer to $r = 0$. See supplementary information for details, for Bayes Factor robustness checks and a Bayesian analogue to power analysis.

Results and discussion

Figure 6. Metacontrol of decision making in Study 2. (A) Mean model-based weights. Error bars represent standard error of the mean. (B) Individual differences in model-based weights. Plots show model-based weights in high-stakes trials (y-axis) against model-based weights in low-stakes trials (x-axis). Points on the identity line represent individuals that showed no adaptation of model-based weights toward stakes conditions. Points above (below) the identity line represent individuals that showed higher (lower) model-based weights in high-stakes trials compared to low-stakes trials.

Metacontrol of decision making

We found increased model-based weights in high-stakes trials compared to low-stakes trials ($BF_{10} = 2.3 \times 10^7$, Figure 6A). Thus, similar to Study 1, participants showed metacontrol of decision making and relied more on model-based decision making when rewards were amplified.

Figure 7. Relationship between NFC (x-axis) and metacontrol of decision making (y-axis) in Study 1.

Need for cognition and metacontrol

As in Study 1, we observed considerable individual differences in metacontrol of decision making (Figure 6B). We found moderate evidence against a correlation between NFC and metacontrol ($r = .06$, $BF_{10} = 0.24$, Figure 7). Moreover, we found strong evidence against a negative correlation between NFC and metacontrol ($BF_{10} = 0.09$). Consistent with our findings in Study 1, these results speak against a role of NFC in explaining individual differences in metacontrol of decision making. Bayes Factor robustness checks indicated that these findings were largely unaffected by different prior assumptions (see Figure S2 in the supplementary information).

Need for cognition and model-based decision making

We found inconclusive evidence regarding a correlation between NFC and model-based weights in low-stakes trials ($r = .10$, $BF_{10} = 0.47$). In contrast to our findings in Study 1, we found strong evidence for a correlation between NFC and model-based weights in high-stakes trials ($r = .21$, $BF_{10} = 15$).

General discussion

We investigated individual differences in metacontrol of decision making in two independent studies. In contrast to our expectations, NFC did not explain individual differences in metacontrol. That is, individuals low and high in NFC similarly increased their reliance on the more accurate but more effortful model-based decision-making strategy when rewards were amplified. Furthermore, we obtained inconsistent evidence of the relationship between NFC and model-based decision making. While the results from Study 2 suggest that high-NFC individuals show more model-based decision making than low-NFC individuals under high-stakes conditions, this was less evident in Study 1.

Our findings suggest that NFC may have a differential role for the expenditure of cognitive effort in decision-making and executive functions. In a study by Sandra and Otto (2018), participants low in NFC showed a stronger ramp-up of cognitive effort in a task-switching paradigm when rewards were amplified than participants high in NFC. Although model-based decision making relies at least partly on executive functions (Otto et al., 2013; Otto et al., 2015), we did not observe a similar relationship between NFC and cognitive effort expenditure during decision making. This could point to a potential process specificity regarding the role of NFC in the regulation of cognitive effort.

A possible limitation of our studies is that they were based on student samples with a distribution of NFC scores shifted above the scale mean. This restricted range could have attenuated the association between NFC and metacontrol and future studies should rely on more representative samples. However, this sample characteristic equally holds for the study

by Sandra and Otto (2018) and we therefore think it is unlikely that this factor explains the absence of a correlation in our two studies.

The stakes manipulation used in our task (factor 1 versus factor 5) was comparable in size to the stakes manipulation in the study by Sandra and Otto (2018) (1 cent versus 5 cent). In contrast to our experimental design, where the stakes condition was assigned to each trial randomly, Sandra and Otto (2018) employed a blockwise stakes manipulation, keeping the stakes condition constant for several trials in a row. Thus, in our two studies, individual differences in the willingness to exert cognitive effort might have been attenuated by individual differences in the ability to adapt cognitive effort dynamically and immediately. However, since young adults can adapt behavioral performance in cognitive control tasks rapidly to changes in the reward structure (Yee, Adams, Beck, & Braver, 2019), we think that this difference in task design is unlikely to explain the different patterns in our study and the study by Sandra and Otto (2018). Further research is needed to better understand the temporal dynamics of metacontrol.

NFC explained individual differences in model-based decision making under high rewards in Study 2, though we did not observe a similar relationship in Study 1 and in low-stakes trials of Study 2. Thus, our findings remain inconclusive regarding a correlation between NFC and model-based decision making. While NFC generally has been associated with more cognitively effortful modes of information processing (Cacioppo et al., 1996), a recent study points to no relationship between NFC and basic executive functions (Gärtner et al., 2019) which largely mirrors the results from our studies. A potential explanation for these inconsistent findings could be that our decision-making task – like most paradigms in experimental psychology – poses strong situational affordances and thus might reduce interindividual variability, making it more challenging to use in correlational research (Hedge, Powell, & Sumner, 2018).

In conclusion, we found that NFC does not account for individual differences in metacontrol of decision making. That is, individuals low and high in NFC equally adapted their reliance on a more accurate but more effortful strategy when rewards were amplified. While some of our findings suggest that high-NFC individuals exert more cognitive effort during decision making, this relationship was not observed consistently in both studies and encourages more research on the role of personality traits in reinforcement learning.

References

- Bless, H., Fellhauer, R. F., Böhner, G., & Schwarz, N. (1994). Need for cognition: Eine Skala zur Erfassung von Engagement und Freude bei Denkaufgaben [Need for cognition: A scale measuring engagement and happiness in cognitive tasks]. *Zeitschrift für Sozialpsychologie*, 25, 147-154.
- Bolenz, F., Kool, W., Reiter, A., & Eppinger, B. (2019). Metacontrol of decision-making strategies in human aging. *Elife*, 8. doi:10.7554/eLife.49154
- Cacioppo, J. T., & Petty, R. E. (1982). The Need for Cognition. *Journal of Personality and Social Psychology*, 42(1), 116-131. doi:10.1037/0022-3514.42.1.116
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., Blair, W., & Jarvis, G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, 119(2), 197-253. doi:10.1037/0033-2909.119.2.197
- Cacioppo, J. T., Petty, R. E., & Kao, C. F. (1984). The efficient assessment of need for cognition. *J Pers Assess*, 48(3), 306-307. doi:10.1207/s15327752jpa4803_13
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204-1215. doi:10.1016/j.neuron.2011.02.027
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: the good, the bad and the ugly. *Curr Opin Neurobiol*, 18(2), 185-196. doi:10.1016/j.conb.2008.08.003
- Gärtner, A., Grass, J., Wolff, M., Goschke, T., Strobel, A., & Strobel, A. (2019). No relation of Need for Cognition to basic executive functions. *PsyArXiv*. doi:10.31234/osf.io/ts97j
- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66(4), 585-595. doi:10.1016/j.neuron.2010.04.016
- Hedge, C., Powell, G., & Sumner, P. (2018). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behav Res Methods*, 50(3), 1166-1186. doi:10.3758/s13428-017-0935-1
- Kool, W., Cushman, F. A., & Gershman, S. J. (2016). When Does Model-Based Control Pay Off? *PLoS Comput Biol*, 12(8), e1005090. doi:10.1371/journal.pcbi.1005090
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-Benefit Arbitration Between Multiple Reinforcement-Learning Systems. *Psychol Sci*, 28(9), 1321-1333. doi:10.1177/0956797617708288
- Kool, W., Gershman, S. J., & Cushman, F. A. (2018). Planning Complexity Registers as a Cost in Metacontrol. *J Cogn Neurosci*, 1-14. doi:10.1162/jocn_a_01263
- Lee, M. D., & Wagenmakers, E. J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge, UK: Cambridge University Press.
- Morey, R. D., Rouder, J. N., & Jamil, T. (2015). Package "Bayes Factors". Retrieved from <https://cran.r-project.org/web/packages/BayesFactor/BayesFactor.pdf>
- Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychol Sci*, 24(5), 751-761. doi:10.1177/0956797612463080
- Otto, A. R., Skatova, A., Madlon-Kay, S., & Daw, N. D. (2015). Cognitive control predicts use of model-based reinforcement learning. *J Cogn Neurosci*, 27(2), 319-333. doi:10.1162/jocn_a_00709
- Patzelt, E. H., Kool, W., Millner, A. J., & Gershman, S. J. (2018). Incentives Boost Model-Based Control Across a Range of Severity on Several Psychiatric Constructs. *Biological Psychiatry*. doi:10.1016/j.biopsych.2018.06.018
- Sandra, D. A., & Otto, A. R. (2018). Cognitive capacity limitations and Need for Cognition differentially predict reward-induced cognitive effort expenditure. *Cognition*, 172, 101-106. doi:10.1016/j.cognition.2017.12.004
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2012). A 21 Word Solution. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2160588
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: an introduction*. Cambridge, MA: MIT Press.

- Thompson, E. P., Chaiken, S., & Hazlewood, J. D. (1993). Need for Cognition and Desire for Control as Moderators of Extrinsic Reward Effects - a Person X Situation Approach to the Study of Intrinsic Motivation. *Journal of Personality and Social Psychology*, 64(6), 987-999. doi:10.1037/0022-3514.64.6.987
- Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS One*, 8(7), e68210. doi:10.1371/journal.pone.0068210
- Yee, D. M., Adams, S., Beck, A., & Braver, T. S. (2019). Age-Related Differences in Motivational Integration and Cognitive Control. *Cogn Affect Behav Neurosci*. doi:10.3758/s13415-019-00713-3

Supplementary information

Need for cognition does not account for individual differences in metacontrol of decision making

Study 1: Sequential decision-making task

Each trial started with an intertrial interval (black screen, 750 ms) and the presentation of a stakes cue (1000 ms). Both stakes cues (low-stakes cue = “1x” and high-stakes cue = “5x”) were assigned with equal probability to trials. After this, one of two first-stage states was presented (3000 ms) with two spaceships displayed side by side (an orange and a turquoise spaceship in one first-stage state, and a green and a blue spaceship in the other first-stage state; all spaceships were displayed equally often on the left or the right side). Participants selected the left or the right spaceship using the keys F and J on a standard computer keyboard and after a choice was made, the respective spaceship was highlighted for the remaining time of the presentation of the first-stage state. Subsequently, one of two second-stage states (a red or a purple planet with an alien) was presented and the second-stage state was deterministically determined by the choice of the spaceship. For each pair of spaceships, there was always one spaceship leading to the red planet and the other spaceship leading to the purple planet. In stable-transitions blocks, the mapping from spaceships to planets was held constant throughout the block, whereas in variable-transitions blocks, every 6 to 14 trials, one of the two pairs of spaceships switched their destination planets. Participants had 2000 ms to respond to the second-stage state by pressing the space bar. After this response window, the amount of rewards available at this planet and points received in this trial was shown. Rewards available at both planet were based on two independent Gaussian random walks (mean = 0, standard deviation = 2, reflecting boundaries at 0 and 9, values were rounded to integers). During the task, a total point count was displayed in the top-right corner of the screen.

If no response was given during the first-stage state or the second-stage state within the respective response window, the trial was canceled, no reward was given and the task proceeded with the next trial. Trials with missing responses at the first-stage state were not included in the analysis (1% of all trials). Different to previous studies with this task, we included trials with missing responses at the second-stage state because these trials could be potentially informative for updates of the transition structure.

The task consisted of 320 trials, grouped into four blocks of 80 trials. Between blocks, the transition condition (stable vs. variable) was alternating and participants were informed at the beginning of each block about the transition condition for the upcoming trials.

In order to decrease variability between participants due to random variations in the task, we kept the reward trajectories and the sequence of first-stage states identical for all participants. The assignment of stakes conditions to trials and of transition conditions to blocks was counterbalanced across participants.

Before starting with the sequential decision-making task, participants received a detailed instruction about the nature of the reward distribution, the transition structure and the stakes manipulation. To ensure their understanding of the task, participants had to select the spaceship leading to a planet (with 10 consecutive correct choices necessary for each planet to proceed) and to specify the number of points given a number of rewards and a stakes cue (with 10 consecutive correct answers necessary to proceed). Moreover, they performed 20 training trials for the stable-transitions condition and 20 training trials for the variable-transitions condition.

Study 1: Reinforcement-learning model

The model-free learner holds reward expectations $Q_{MF}(s_i, a_i)$ for each action a_i and state s_i at the i -th stage of the task. At the beginning of the task, all reward expectations are set to 4.5 (reflecting the mean of the range of possible rewards). After each choice, all reward expectations are updated according to a temporal-difference learning rule:

$$Q_{MF}(s_i, a_i) \leftarrow Q_{MF}(s_i, a_i) + \alpha \times e(s_i, a_i) \times \delta$$

Here, α is the reward learning rate (bounded between 0 and 1) reflecting how quickly new experiences are integrated into reward expectations. The eligibility trace $e(s_i, a_i)$ is set to 0 for all combinations of s_i and a_i at the beginning of a trial; before updating reward expectations, the eligibility trace for the immediately preceding state-action pair (s_i, a_i) is set to 1 and after the update, all eligibility traces are decayed by the eligibility trace decay parameter λ (bounded between 0 and 1). The reward prediction error δ reflects the discrepancy between experienced and expected reward and is computed as

$$\delta = r + Q_{MF}(s_{i'+1}, a_{i'+1}) - Q_{MF}(s_i, a_i)$$

where r is the immediate reward obtained after a choice (note that r is always 0 after first-stage choices) and $Q_{MF}(s_{i'+1}, a_{i'+1})$ is the reward expectation associated with the subsequent action $a_{i'+1}$ in the new state $s_{i'+1}$ (note that $Q_{MF}(s_{i'+1}, a_{i'+1})$ is always 0 after second-stage choices because the states that offer rewards are terminal).

The model-based learner maintains a model of the task structure represented by a transition matrix $T(s_2 | s_1, a_1)$ that holds probabilities for moving to a second-stage state given an action and a first-stage state. At the beginning of the task, all transition probabilities are 0.5 and after observing a transition to a second-stage state, these probabilities are updated according to

$$T(s_2|s_1, a_1) \leftarrow T(s_2|s_1, a_1) + \eta \times \delta^{SPE}$$

$$T(\neg s_2|s_1, a_1) \leftarrow T(\neg s_2|s_1, a_1) \times (1 - \eta)$$

Here, η is the transition learning rate (bounded between 0 and 1) that reflects how quickly observations of transitions are integrated into the representation of the task structure. To ensure that the sum of transition probabilities stays 1, the probability for transitioning to the alternative second-stage state $\neg s_2$ needs to be adjusted. The state prediction error δ^{SPE} is computed as

$$\delta^{SPE} = 1 - T(s_2|s_1, a_1)$$

It is possible to infer the second-stage state to which the alternative, not-chosen first-stage action would have led because both actions available in a first-stage state always lead to different second-stage states. Thus, the model-based learner also updates transition probabilities for the alternative action in the same way as it does for the actual actions, using a counterfactual transition learning rate η_{CF} (bounded between 0 and 1).

While the model-based reward expectations at the second stage are identical to the model-free reward expectations (because both reflect an estimate of the immediate reward), the model-based reward expectations at the first stage are computed as

$$Q_{MB}(s_1, a_1) = \sum_{s_2} T(s_2|s_1, a_1) Q_{MB}(s_2, a_2)$$

At the first stage, both model-free and model-based reward expectations are combined to an integrated reward expectation $Q(s_1, a_1)$ with the model-based weight ω (bounded between 0 and 1) reflecting the relative influence of the model-based learner.

$$Q(s_1, a_1) = (1 - \omega)Q_{MF}(s_1, a_1) + \omega Q_{MB}(s_1, a_1)$$

Choice probabilities at the first stage are modeled by a softmax function:

$$P(a_1|s_1) = \frac{\exp(\beta[Q(s_1, a_1) + \pi \cdot rep(a_1) + \rho \cdot resp(a_1)])}{\sum_{a'} \exp(\beta[Q(s_1, a') + \pi \cdot rep(a') + \rho \cdot resp(a')])}$$

Here, β is the inverse softmax temperature (left-bounded at 0) that reflects how consistently choices are guided by reward expectations. The choice stickiness π and the response stickiness ρ (both unbounded) capture perseveration (positive values) or switching (negative values) of choices (which stimulus was selected) or responses (which key was pressed) across trials. The indicator variables $rep(a_1)$ and $resp(a_1)$ are set to 1 if the same stimulus or the corresponding response key were selected in the previous trials (and are set to 0 otherwise).

Study 1: Reinforcement-learning model

We obtained individual maximum a posteriori parameter estimates using the mfit toolbox in Matlab (Gershman, 2016) with the following priors: Beta(2,2) priors for α , λ , η , η_{CF} and ω ; Normal(0,1) priors for π and ρ ; a Gamma(3,0.2) prior for β . To avoid local optima, the optimization procedure was started 100 times for each participant and we used the parameters of the run with the highest posterior probability. Model-based weights were estimated separately for low-stakes and high-stakes trials in both stable-transitions and variable-transitions blocks. Transition learning rates were estimated only for variable-transitions blocks and set to 1 during stable-transitions blocks. We fitted different versions of the model where the parameters λ , η , η_{CF} , π and ρ were varied to be free or fixed parameters and we selected the best-fitting model version based on the Akaike Information Criterion (free parameters: λ , π , ρ ; fixed parameters: $\eta = \eta_{CF} = 1$).

Study 1: Data analysis.

For Bayes Factors concerning hypotheses about differences in means, we used JZS priors with scaling parameter $r = \sqrt{2}/2$ (Morey & Rouder, 2011) and for Bayes Factors concerning hypotheses about correlations, we used stretched beta priors with scaling parameter $\kappa = 1$ (Ly, Verhagen, & Wagenmakers, 2016) which assign equal prior probabilities to correlations between -1 and 1. We also conducted Bayes Factor robustness checks for our primary analyses, varying κ between 0.01 and 1 (see Figure S1).

We performed a Bayesian analogue to power analysis for a two-sided correlation test (i.e., comparing evidence for a non-zero correlation and for no correlation) and a one-sided correlation test (i.e., comparing evidence for a negative correlation and for no correlation). For different true population correlations ρ in a population of $N = 1,000$, we computed the proportion of 10,000 random samples of $n = 126$ for which a Bayes Factor would show at least moderate evidence for the null hypothesis or the alternative hypothesis. With our sample of $N = 126$ and in a two-sided correlation test, we would find at least moderate evidence with 80% probability for a non-zero correlation if the true population correlation is $|\rho| \geq .31$ and for no correlation if the true population correlation is $|\rho| \leq .06$. In a one-sided correlation test, we would find at least moderate evidence with 80% probability for a negative correlation if the true population correlation is $\rho \leq -.28$ and for no correlation if the true population correlation is $\rho \geq -0.02$.

Study 2: Sequential decision-making task

We used a variant of the sequential decision-making task from study 1. Trials followed the same structure, but with a different pacing. The intertrial interval was presented for 300 ms, the stakes cue was presented for 800 ms, and both first-stage state and second-stage state were presented for 1500 ms each. The entire task consisted of 280 trials, equally distributed

between low-stakes and high-stakes trials. There were no changes in the transition structure, so all spaceships kept their initial destination planets.

We excluded all trials from analysis in which no response was given during the presentation of either the first-stage state or the second-stage state (3% of all trials).

In order to decrease variability between participants due to random variations in the task, we created two independent trial sequences that determined reward trajectories and first-stage states and that were counterbalanced across participants. Within each trial sequences, the assignment of stakes conditions to trials was counterbalanced.

Before starting the task, participants received a similar task instruction as in study 1, apart from performing 25 training trials, all with a stable task structure.

Study 2: Reinforcement-learning model

We adapted the reinforcement-learning model from study 1. Due to no changes in the transition structure in this version of the task, we set the transition matrix T to reflecting the true transition probabilities at the beginning of the task and did not model any updates of the transition matrix, thus abandoning the transition learning rate and the counterfactual transition learning rate as model parameters. Also, this model only contained two model-based weights (one for low-stakes trials and one for high-stakes trials).

Study 2: Model-fitting procedure:

We used the same model-fitting procedure as in study 1 with the following exceptions: Only two model-based weights were fitted for each participant, one for low-stakes trials and one for high-stakes trials. We fitted different versions of the model where the parameters λ , π and ρ were varied to be free or fixed parameters and we selected the best-fitting model version based on the Akaike Information Criterion (free parameters: π , ρ ; fixed parameter: $\lambda = 0$).

Study 2: Data analysis.

Different from Study 1, we used a scaling parameter $\kappa = 1/3$ for the stretched beta priors for Bayes Factors concerning hypotheses about correlations. Thus, more prior weight was given to correlation coefficients closer to 0, reflecting our findings in Study 1. We also conducted Bayes Factor robustness checks for our primary analyses, varying κ between 0.01 and 1 (see Figure S2).

We performed a Bayesian analogue to power analysis similar to Study 1. With our sample of $N = 205$ and in a two-sided correlation test, we would find at least moderate evidence with 80% probability for a non-zero correlation if the true population correlation is $|\rho| \geq .23$ and for no correlation if the true population correlation is $|\rho| \leq .01$. In a one-sided correlation test, we

would find at least moderate evidence with 80% probability for a negative correlation if the true population correlation is $\rho \leq -.21$ and for no correlation if the true population correlation is $\rho \geq -0.001$.

Figure S1. Bayes Factor robustness checks for the analyses of correlations between NFC and metacontrol in Study 1. Two-sided tests compare evidence for a non-zero correlation and for no correlation; one-sided tests compare evidence for a negative correlation and for no correlation. Plots show Bayes Factors for different values of the scaling parameter κ (ranging between 0.01 and 1). For $\kappa = 1$, all population correlations between -1 and 1 are equally likely a priori and with smaller values of κ , correlations near zero become more likely a priori than correlations close to -1 or 1. Highlighted are the scaling value used in this study ($\kappa = 1$, red) and the scaling value implemented as a default in the BayesFactor package ($\kappa = 1/3$, gray)

Figure S2. Bayes Factor robustness checks for the analyses of correlations between NFC and metacontrol in Study 2. The two-sided test compares evidence for a non-zero correlation and for no correlation; the one-sided test compares evidence for a negative correlation and for no correlation. Plots show Bayes Factors for different values of the scaling parameter κ (ranging between 0.01 and 1). For $\kappa = 1$, all population correlations between -1 and 1 are equally likely a priori and with smaller values of κ , correlations near zero become more likely a priori than correlations close to -1 or 1. Highlighted are the scaling value used in this study ($\kappa = 1/3$, red) and the scaling value used in Study 1 ($\kappa = 1$, gray)

Supplementary references

- Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, 71, 1-6. doi:10.1016/j.jmp.2016.01.006
- Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016). Harold Jeffreys's default Bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72, 19-32. doi:10.1016/j.jmp.2015.06.004
- Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychol Methods*, 16(4), 406-419. doi:10.1037/a0024377