OSF Preprint

Boosting model-based control in older adults

Edward H. Patzelt^{1,2}, Wouter Kool¹, Samuel J. Gershman^{1,2}

¹Department of Psychology, Harvard University

²Center for Brain Science, Harvard University

Author Note

Corresponding author:

Edward H. Patzelt

Harvard University

52 Oxford Street, Northwest Building, 295.07, Cambridge, MA, 02138

phone: 651-315-3410, email: patzelt@g.harvard.edu

Keywords: older adults; reinforcement learning; model-based control; incentives; habits and goals

Abstract Words: 238

Article Words: 3715

Figures: 3

Abstract

The tension between habits and plans is reflected in everyday decision-making. Habits are computationally cheap, but fail to flexibly adapt to changes in the environment. Planning is a flexible decision-making strategy, but requires greater resources. Arbitration between habits and plans has been formalized using reinforcement learning algorithms that distinguish between model-free control (habits) and model-based control (plans). Evidence about these two decision-making approaches suggests model-based control follows a developmental trajectory, emerging during adolescence, strengthening during young adulthood, and declining in older adulthood. The normative decline in planning (model-based control) presents the opportunity to develop interventions to increase flexible decision-making. Therefore, we asked if incentives could be used to increase model-based control in older adults. We expected older adults would fail to increase model-based control in response to incentives. This prediction was based upon prior research suggesting older adulthood is associated with deficits in representing and updating the expected value of rewards. Contrary to our expectations, in Experiment 1 we found that incentives could be used to boost model-based control in older adults sampled from an online population. We hypothesized this may be due to previous experience with the task (or with similar tasks). In Experiment 2, a naïve sample of older adults did not boost model-based control in response to incentives. These results suggest that incentives may be a useful intervention to increase model-based planning in older adulthood, but this may require extensive experience with the incentive structure.

Introduction

Everyday decision-making can often rely on habitual behavior, but sometimes demands deliberative planning. Habits provide a computationally cheap problem-solving strategy, but they cannot flexibly adapt to changes in the environment. Planning, on the other hand, provides greater flexibility at a greater computational cost. The tension between habits and plans is particularly useful for understanding the processes underlying development changes in decision-making across the lifespan such as the emergence of flexible planning during adolescence (Decker, Otto, Daw, & Hartley, 2016), and its' subsequent decline in older adulthood (Eppinger, Walter, Heekeren, & Li, 2013).

The distinction between habits and plans has been formalized within a reinforcement learning framework (Daw, Niv, & Dayan, 2005; Dolan & Dayan, 2013; Kool, Cushman, & Gershman, 2017; Sutton & Barto, 1998) which has also been invoked to understand changes in decision-making across the lifespan (Raab & Hartley, 2018). Within this framework, habits are operationalized as a form of “model-free control”, and plans are operationalized as a form of “model-based control”. Model-free control is a modern form of Thorndike's *law of effect* (Thorndike, 1911): actions leading to rewards are more likely to be repeated. This is implemented by updating cached reward predictions based on trial and error. Critically, these cached predictions require interaction with the environment, and therefore cannot quickly adapt to environmental changes, such as decreases in the value of a particular decision. Model-based control uses an internal model of the environment, akin to a “cognitive map” (Tolman, 1948), that specifies the reward and state transition structure. This model allows for prospective simulation (Doll, Duncan, Simon, Shohamy, & Daw, 2015) of sequential decision paths, and permits flexible adaptation in response to environmental changes by updating the model.

Flexible decision-making is a critical skill throughout the lifespan. This is why recent research has focused on methods to boost model-based control within the individual. Efforts to increase model-based control using incentives have been successful (Kool, Gershman, & Cushman,

2017), and this is also true for people with elevated symptoms and traits of psychopathology (Patzelt, Kool, Millner, & Gershman, 2019a). These findings demonstrate the possibility of boosting model-based control in older adults. Indeed, normative aging is associated with model-based learning deficits such as a failure to maintain stable task representations (Hämmerer et al., 2018) and a diminished capacity to integrate uncertainty into decision-making (Eppinger et al., 2013; Nassar et al., 2016).

In the current investigation, we use a sequential decision-making task to ask a critical question about the pliability of model-based control in older adulthood: *can we use incentives to boost model-based control in older adults?* In the next two sections, we review prior research on model-based control across the lifespan and introduce a theoretical framework for understanding the factors governing arbitration between model-free and model-based controllers. We then report two behavioral experiments comparing younger and older adults using a parametric estimate of model-based control under different incentive conditions.

Model-based control across the lifespan

Research into model-based control shows that it follows a developmental arc throughout the lifespan. Decker and colleagues (Decker et al., 2016) demonstrated the apparent absence of model-based control during childhood (ages 8-12), its emergence during adolescence (ages 13-17), and a further bolstering of model-based control during early adulthood (ages 18-25). This study is complemented by the work of Eppinger and colleagues (2013), who investigated model-based control in younger versus older adults using a very similar task. They showed that older adults have model-based deficits which are acutely evident during decisions following unexpected rewards. Whereas younger adults tended to explore the task structure in this state of uncertainty, older adults did not explore the task structure, demonstrating a perseverative response pattern characterized by lower model-based control. The developmental trajectory of model-based control provides an empirical and computational foundation for theoretical accounts of cognitive

processes that are tuned differently across various developmental stages (Hartley & Somerville, 2015; Raab & Hartley, 2018).

Across developmental studies of model-based control, the sequential decision-task used to measure model-based control was developed by Daw and colleagues (Daw, Gershman, Seymour, Dayan, & Dolan, 2011), and we refer to this as the Daw two-step task. The task measured model-based control as a trait within the individual that remained consistent throughout the course of the game. It was then hypothesized that state dynamics may exist that allow model-based control to adaptively fluctuate with changes in the person's situation or environment. This hypothesis was tested in a series of experiments by Kool and colleagues (Kool, Cushman, & Gershman, 2016; Kool, Gershman, et al., 2017) where they demonstrated state-dependent fluctuations in model-based control.

In a novel two-step task design, Kool and colleagues introduced an experimental manipulation to incentivize adaptive boosts in model-based control (Kool, Gershman, et al., 2017). During the course of the game, some trials were “low stakes” (subjects earned the presented reward) and some tasks were “high stakes” (reward was multiplied by 5). This manipulation demonstrated that model-based control could be boosted within the individual. Moreover, recent work using the same task showed model-based deficits could be improved across a range of transdiagnostic symptoms and traits of psychopathology (Patzelt et al., 2019a). This suggests that incentives might also boost model-based control in aging adults and lead to more flexible decision-making strategies.

Metacontrol framework

As mentioned above, model-free and model-based control confer different costs and benefits. The experiments of Kool and colleagues (2017) suggest that arbitration between these two modes of control occurs via a metacontrol cost-benefit analysis. A metacontroller arbitrates between the lower cost (but potentially lower reward) model-free controller, and higher cost (but potentially higher reward) model-based controller (schematized in Figure 1). To illustrate, boosting

model-based control in the Daw two-step task does not lead to greater reward, and therefore the metacontroller will prefer the resource cheap model-free strategy. In contrast, the metacontroller will engage model-based control on the Kool two-step task, because this will lead to greater reward. This metacontrol framework helps explain why it is advantageous to increase model-based control in response to incentives on the Kool two-step task.

The metacontrol cost-benefit analysis also provides a framework for predicting how aging adults may respond to incentives in the two-step task. Aging is associated with declines in cognitive function (Luszcz, 2011), and it has been argued that this is due to selective engagement of cognitive resources due to their greater cost in older adults (Hess, 2014). In turn, this predicts that engagement of model-based control will be associated with greater costs in older adults. Taken together with the work of Eppinger and colleagues (2013), we expect older adults to show model-based deficits, and impairments in their ability to increase model-based control in response to large incentives due to the higher costs of model-based control. If we find that older adults show model-based deficits on the Kool two-step task and/or impairments in responding to incentives, it would suggest older adults may have trouble accurately estimating the costs associated with model-based control. However, if we find the older adults can boost model-based control in response to incentives, it suggests interventions targeting model-based control may help increase goal-directed behavior in older adults.

Experiment 1

Methods

Participants

We used the TurkPrime crowdsourcing platform (Litman, Robinson, & Abberbock, 2017) to recruit participants from Amazon Mechanical Turk (AMT; Seattle, WA). Participants completed

the Kool two-step task (Kool, Gershman, et al., 2017) and a computer adaptive IQ test (CAT; Raven, 2000). They were compensated \$11 (plus a possible bonus between \$0 and \$2.50) and provided informed consent. The study and experimental protocols followed ethical guidelines and were approved by the Committee on the Use of Human Subjects (CUHS), the Institutional Review Board (IRB) for Harvard University. A subset of 377 young adults ($\leq$ age 30) and 47 older adults ($\geq$ age 56) were selected from a larger study conducted on AMT that examined relationships between psychopathology and model-based control (Patzelt et al., 2019a); and separately, psychopathology and mental effort (Patzelt, Kool, Millner, & Gershman, 2019b). All participants were required to have U.S. residency, 90% approval rating, and 100 completed Mechanical Turk Human Intelligence Tasks. We recruited an additional 298 older adults (age $\geq$ 56) from AMT and 16 were excluded for non-response on $>20\%$ of trials. In addition, participants from both younger and older adults were excluded if their computational model parameters were not estimable ($N = 5$), or they did not provide their gender ($N = 2$). Thus, the final sample was 374 younger adults ($M = 26.38$, $SD = 2.85$) and 315 older adults ($M = 61.46$, $SD = 4.91$).

Sequential decision task

Participants performed 200 trials of the two-step task developed by Kool and colleagues (Kool, Gershman, et al., 2017). This allowed us to measure increases in model-based control when comparing high stakes to low stakes (Figure 2A/B). Participants randomly started in one of two first-stage states and chose one of two rocket ships (which were randomly mapped to response keys). Following this choice, the participant deterministically transitioned to either a red or purple alien planet. Upon arriving on the alien planet, the participant received a reward indicated by alien treasure. The amount of reward obtained at each planet changed randomly and gradually over the task, independently for each planet (a Gaussian random walk between 0 and +9 with $SD = 2$). If the participant did not respond, no reward was delivered and the task proceeded to the next trial. Each trial was randomly assigned to one of the two stakes conditions.

In the “high stakes” condition, the participant received 5 times the alien treasure reward (5x multiplier; Figure 1A). In the “low stakes” condition, the participant received the displayed alien treasure reward.

Computational model

We modeled the two-step task using a dual-process reinforcement learning model that captures the mixture of model-based and model-free control exhibited in decision-making on sequential decision tasks (Daw et al., 2011; Kool et al., 2016; Kool, Gershman, et al., 2017). The model free algorithm is SARSA(λ) (Sutton & Barto, 1998), which updates the state-action pair (s, a) at step i and trial t in the following manner:

$$Q_{MF}(s, a) \leftarrow Q_{MF}(s, a) + \alpha \delta_{i,t} e_{i,t}(s, a)$$

where the prediction error is defined as:

$$\delta_{i,t} = r_{i,t} + Q_{MF}(s_{i+1,t}, a_{i+1,t}) - Q_{MF}(s_{i,t}, a_{i,t}).$$

The learning rate is represented by α . This captures the rate at which the participant incorporates reward feedback (i.e. low α leads to slow incorporation and high α leads to quick incorporation) and $e_{i,t}(s, a)$ is an eligibility trace updated in the following way:

$$e_{i,t}(s, a) \leftarrow e_{i-1,t}(s_{i,t}, a_{i,t}) + 1.$$

The eligibility trace allows step 1 values to be partially updated following step 2 outcomes.

The model-based strategy uses a learned transition function to compute values, rather than retrieving state-action values from a cache (as used in the model-free approach):

$$Q_{MB}(s_A, a_j) = P(s_B|a_j) Q_{MF}(s_B) + P(s_C|a_j) Q_{MF}(s_C)$$

where $P(s_B|a_j)$ is the probability of transitioning to state s_B after choosing action a_j , and $Q_{MF}(s_B)$ and $Q_{MF}(s_C)$ are the immediate reward estimates at the second-step states (note that these are not dependent on actions since there is no second-step action).

The mixture of model-free and model-based control is governed by a weighting parameter w :

$$Q_{net}(s_A, a_j) = wQ_{MB}(s_A, a_j) + (1 - w)Q_{MF}(s_A, a_j)$$

The w parameter was fit independently for the low stakes and high stakes conditions, which allowed for comparison of model-based control across stakes.

Finally, the probability of choosing action a given state s on a trial is given by a softmax function:

$$P(s_{i,t}) = \frac{\exp(\beta[Q_{net}(s_{i,t}, a) + \pi \cdot rep(a) + \rho \cdot resp(a)])}{\sum_{a'} \exp(\beta[Q_{net}(s_{i,t}, a') + \pi \cdot rep(a') + \rho \cdot resp(a')])},$$

where we have included two parameters to capture additional variance: motor stickiness π (pressing the same key twice), and stimulus stickiness ρ (choosing the same stimulus twice independent of its value). β represents the inverse temperature parameter that controls choice stochasticity. Lower values of β promote exploration and higher values promoting exploitation. Parameters were optimized using maximum a posteriori estimation with empirical priors (Gershman, 2016; Kool, Gershman, et al., 2017). See Kool et al. (2017) for expanded details

concerning task design and computational model-fitting.

Kool and colleagues (Kool, Gershman, et al., 2017) compared various model formulations of the Kool two-step task and found that the stake-dependent mixture model outperformed other approaches which allowed other parameters to vary (e.g., choice inverse temperature). They showed that the change in model-based control due to stakes cannot be completely explained by other model parameters.

Bayesian analysis of model-based control

We used Bayesian linear regression to estimate the interaction between model-based control across high- and low-stakes and younger- versus older- adults. Whereas classic (frequentist) linear regression analysis typically conducts binary inference testing in the form of rejecting (or accepting) the null hypothesis, Bayesian linear regression (Baldwin & Larson, 2016; Gelman et al., 2013) provides an uncertainty distribution about the strength of the relationship(s) between independent and dependent variables. Specifically, frequentist linear regression provides a point estimate of the regression coefficients, whereas Bayesian regression provides a full posterior distribution, often summarized by the mean or median and the 95% highest posterior density (HPD) credible intervals around the mode, therefore indicating uncertainty about the regression coefficients.

When the posterior is centered around 0, it is both plausible and highly probable (due to the large density over 0) that the IV(s) predict no change in the DV. Therefore, we report the posterior probability that the regression coefficient is less than or greater than 0, to indicate the probability that the interaction of *stakes x age group* is associated with changes in model-based control. We used the brms package (42) with the default prior $\sigma \sim \text{student-}t(3,0,10)$ and regressed *stakes x age group* onto model-based control (*w*) while controlling for gender.

Results

Model-based control, stakes, and reward rate

We replicated previous studies using the Kool two-step task (Kool, Gershman, et al., 2017; Patzelt et al., 2019a), showing that model-based control was greater during high stakes than low stakes, with a posterior probability >0.99 that the regression coefficient was greater than zero. Similarly, increasing reward rate was associated with increases in model-based control, with a posterior probability >0.99 that the regression coefficient was greater than 0.

Model-based control: stakes $\times$ age group

We did not find appreciable evidence for an interaction between stakes and younger versus older adults. Figure 3A shows that both age groups had non-overlapping credible intervals and the posterior probability distribution for this interaction contained zero (76% < 0 , 24% > 0), indicating there is a high probability that *stakes $\times$ age group* is unlikely to be associated with changes in model-based control. Moreover, we did not replicate the Eppinger et al. (2013) finding of model-based deficits in older adults, the main effect of age group regressed onto model-based control produced a posterior probability distribution roughly centered around 0 (40% < 0 , 60% > 0).

. We hypothesized this divergence with Eppinger and colleagues might be due to a lower mean age in the older adults compared to the Eppinger study where older adults were ~69 years old. Thus, we restricted our sample to 77 older adults above the age of 65 ($M = 68.61$, $SD = 3.38$), but still found weak evidence of an association between age group and model-based deficits (*stakes $\times$ age group interaction*; 71% < 0 , 29% > 0).

Discussion

In this study we asked if older adults would increase model-based control in response to incentives. We used a sequential decision-making task (Kool, Gershman, et al., 2017) with

incentives to increase model-based control, and a sample of younger ($M = 26.38$, $SD = 2.85$) and older adults ($M = 61.46$, $SD = 4.91$). Contrary to our expectations, we found that older adults did not show the same model-based deficits found by Eppinger and colleagues (2013). Furthermore, older adults increased model-based control in response to incentives. This divergence could not be explained by differences in ages between our sample and that of Eppinger et al. (2013). We next consider the possibility that older adult MTurk workers might differ in other population characteristics.

A recent review (Chandler & Shapiro, 2016) of studies using MTurk populations suggests that our findings in older adults may be due to increased cognitive aptitude or computer literacy. Moreover, as Chandler and Shapiro note, many MTurk workers have completed several online studies, perhaps increasing their ability to perform on cognitive paradigms. Finally, sequential decision tasks have been increasingly deployed online by our lab (Kool et al., 2016; Kool, Gershman, & Cushman, 2018; Kool, Gershman, et al., 2017; Patzelt et al., 2019a) and others (Gillan, Kosinski, Whelan, Phelps, & Daw, 2016). Thus, we hypothesized older adults on MTurk may *not* show model-based deficits because, 1) they have greater technological skill, or 2) they are not naïve to sequential decision-making tasks.

Experiment 2

In Experiment 2, we sought to recruit a sample of older adults that would be naïve to the sequential decision tasks and the cognitive paradigms commonly used on MTurk. We also added a debriefing question “*Have you ever done a game similar to this?*”. These changes allowed us to increase the probability of recruiting a sample more representative of older adults in the general population. Again, we asked if older adults would show model-based deficits, and if they would increase model-based control in response to incentives. We recruited a naïve sample through the Turkprime crowdsourcing platform of 56 older adults ($M = 66.30$, $SD = 6.42$), where one participant was excluded due to previous experience with a similar task. Of

note, we began with the same age cutoff of 56 years old, but raised this to 65 or older to try and match the mean age of Eppinger et al., (2013).

Results and Discussion

Like experiment 1, we found both stakes and average reward rate were positively associated with increases in model-based control (posterior probability >0.99 that coefficient is greater than 0). The age group factor (younger versus older adults) was not associated with deficits in model-based control, with the posterior probability distribution indicating large uncertainty about the presence of an association ($24\% < 0$, $76\% > 0$). In contrast to experiment 1, Figure 3B demonstrates an interaction whereby older adults did not boost model-based control in response to incentives (relative to younger adults) and the posterior probability distribution did not include zero ($99\% < 0$, $1\% > 0$). These results indicate strong evidence that experimentally naïve older adults fail to increase model-based control in response to incentives. Differences in the results between experiments 1 and 2 could not be attributed to differences in mean age.

General discussion

In two experiments, we asked if older adults would show a boost in model-based control in response to incentives. In Experiment 1, we did not find the expected pattern of model-based deficits in older adults (Eppinger et al., 2013). We recruited younger and older adults from Amazon Mechanical Turk and found that older adults (1) boosted model-based control in response to incentives, and (2) did not show model-based deficits. We next hypothesized that older adults on MTurk had greater technological and/or cognitive aptitude due to experience with sequential decision tasks and/or other cognitive paradigms on the Amazon Mechanical Turk crowdsourcing platform.

Therefore, in experiment 2 we recruited a naïve sample of older adults. Like experiment

1, we found that older adults did not show model-based deficits, but in contrast to experiment 2, they were impaired in boosting model-based control in response to large incentives. The findings across both experiments suggest that incentives can boost model-based control in older adults, but that this may require previous experience with sequential decision-making paradigms, and in particular subjects may need previous experience with the incentive structure in order to learn it adequately.

Prior research on the effect of incentives on physical activity and cognition in older adults has yielded mixed results. For example while some have found incentives increase physical activity (Finkelstein, Brown, Brown, & Buchner, 2008), others have found both peer networks and incentives do not increase physical activity (Kullgren et al., 2014). Kullgren and colleagues (2014) hypothesize their null effect may be due to sample characteristics, such as the older adults in their study being highly educated, primarily white, and in good health. Our study may have faced similar limitations with differences in the populations. For example, MTurk older adults may be more technologically adept, have higher levels of education, or have greater cognitive ability than the naïve sample. However, another study demonstrated older adults can increase memory retrieval speed in response to incentives (Touron, Swaim, & Hertzog, 2007), suggesting that the effect in experiment 1 may not be solely attributable to practice effects or general non-naïveté of cognitive tasks.

There are two findings that suggest the lack of incentive effect in Experiment 2 is a specific deficit in using incentives to dynamically engage model-based control. First, older adults did not show model-based deficits overall, which indicates that they were engaged with the task and the deficit was specific to incentives. Second, model-based control and average reward rate were positively correlated in both younger and older adults. Together, these findings suggest that younger and older adults were both motivated and engaged with the task, but older adults in experiment 2 expressed deficits in their response to incentives.

Limitations

Our study was limited in several ways. First, we are unable to fully disentangle factors that account for differences in model-based control in older adults between experiment 1 and experiment 2. These include general cognitive ability and previous practice with decision-making paradigms. Second, we are unable to determine what component of the metacontroller in older adults might be defective. Their inability to boost model-based control in response to incentives may be due to a biased cost-benefit analysis and/or a deficit in the model-based controller. Third, we are unable to answer if older adults in experiment two would increase model-based control on the Kool two-step task following further practice or experience with other cognitive paradigms.

Conclusions

We studied the possibility of boosting model-based control in older adults using a sequential decision-making task across two experiments. Whereas older adults recruited from MTurk increased model-based control in response to incentives, a naïve sample of older adults did not. Our study provokes an outstanding question: can we use repeated practice with sequential decision tasks to improve goal-directed decision-making in older adults, or are the deficits expressed in experiment two trait-like and unchanging? If older adults who show deficits in their response to incentives can be trained to increase model-based control, then interventions targeting this process may hold promise for increasing cognitive ability or goal-directed decision-making. For the older adults that responded to incentives, such as the experiment 1 sample, providing incentives may be an effective approach to enhance goal-directed behavior.

Acknowledgements

This project was partially supported by a training grant from the NIH Blueprint for Neuroscience Research (T90DA022759/R90DA023427), the Office of Naval Research (N000141712984), and NIH grant CRCNS R01-1207833. Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the NIH. We would like to thank Catherine Hartley for sharing her stimuli, Michal Kosinski for use of the IQ test, Drs. Bruce Rosen and Maria Mody for their work with the AMNTP, Ben Eppinger and Florian Bolenz for their feedback during the analysis, and the members of the Computational Cognitive Neuroscience Laboratory for their support and feedback.

Financial Disclosures

The authors report no biomedical financial interests or potential conflicts of interest.

References

- Baldwin, S. A., & Larson, M. J. (2016). An introduction to using Bayesian linear regression with clinical data. *Behaviour Research and Therapy*, 98, 58–75.
<https://doi.org/10.1016/j.brat.2016.12.016>
- Chandler, J., & Shapiro, D. N. (2016). Conducting Clinical Research Using Crowdsourced Convenience Samples. *Annual Review of Clinical Psychology*, 12, 53–81.
<https://doi.org/10.1146/annurev-clinpsy-021815-093623>
- 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.
<https://doi.org/10.1016/j.neuron.2011.02.027>
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704–1711. <https://doi.org/10.1038/nn1560>
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychological Science*, 27(6), 848–858.
<https://doi.org/10.1177/0956797616639301>
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312–325.
<https://doi.org/10.1016/j.neuron.2013.09.007>
- Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D., & Daw, N. D. (2015). Model-based choices involve prospective neural activity. *Nature Neuroscience*, 18(5), 767–772.
<https://doi.org/10.1038/nn.3981>
- Eppinger, B., Walter, M., Heekeren, H. R., & Li, S.-C. (2013). Of goals and habits: Age-related and individual differences in goal-directed decision-making. *Frontiers in Neuroscience*, 7,

253. <https://doi.org/10.3389/fnins.2013.00253>

Finkelstein, E. A., Brown, D. S., Brown, D. R., & Buchner, D. M. (2008). A randomized study of financial incentives to increase physical activity among sedentary older adults. *Preventive Medicine*, 47(2), 182–187. <https://doi.org/10.1016/j.ypmed.2008.05.002>

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Boca Raton, FL: CRC Press.

Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, 71, 1–6. <https://doi.org/10.1016/j.jmp.2016.01.006>

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *ELife*, 5. <https://doi.org/10.7554/eLife.11305>

Hämmerer, D., Schwartenbeck, P., Gallagher, M., FitzGerald, T. H. B., Düzel, E., & Dolan, R. J. (2018). Older adults fail to form stable task representations during model-based reversal inference. *Neurobiology of Aging*, 74, 90–100. <https://doi.org/10.1016/j.neurobiolaging.2018.10.009>

Hartley, C. A., & Somerville, L. H. (2015). The neuroscience of adolescent decision-making. *Current Opinion in Behavioral Sciences*, 5, 108–115. <https://doi.org/10.1016/j.cobeha.2015.09.004>

Hess, T. M. (2014). Selective Engagement of Cognitive Resources: Motivational Influences on Older Adults' Cognitive Functioning. *Perspectives on Psychological Science*, 9(4), 388–407. <https://doi.org/10.1177/1745691614527465>

Kool, W., Cushman, F. A., & Gershman, S. J. (2016). When Does Model-Based Control Pay Off? *PLOS Computational Biology*, 12(8), 1–34. <https://doi.org/10.1371/journal.pcbi.1005090>

Kool, W., Cushman, F. A., & Gershman, S. J. (2017). *Competition and Cooperation Between Multiple Reinforcement Learning Systems*. (R. W. Morris, A. M. Bornstein, & A. Shenhav,

- Eds.), *Understanding Goal-Directed Decision Making: Computations and Circuits* ((Eds)). New York: Elsevier.
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-benefit arbitration between multiple reinforcement learning systems. *Psychological Science*, 28(9), 1321–1333.
<https://doi.org/10.1177/0956797617708288>
- Kool, W., Gershman, S. J., & Cushman, F. A. (2018). Planning Complexity Registers as a Cost in Metacontrol. https://doi.org/10.1162/jocn_a_01263
- Kullgren, J. T., Harkins, K. A., Bellamy, S. L., Gonzales, A., Tao, Y., Zhu, J., ... Karlawish, J. (2014). A Mixed-Methods Randomized Controlled Trial of Financial Incentives and Peer Networks to Promote Walking Among Older Adults. *Health Education & Behavior*, 41(1_suppl), 43S–50S. <https://doi.org/10.1177/1090198114540464>
- Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442. <https://doi.org/10.3758/s13428-016-0727-z>
- Luszcz, M. (2011). Executive function and cognitive aging. In K. Schaie & S. Willis (Eds.), *Handbook of the psychology of aging* (7th ed., pp. 59–72). San Diego, CA: Elsevier.
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S. C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, 7(May 2015), 1–13. <https://doi.org/10.1038/ncomms11609>
- Patzelt, E. H., Kool, W. J., Millner, A. J., & Gershman, S. J. (2019a). Incentives Boost Model-based Control Across a Range of Severity on Several Psychiatric Constructs. *Biological Psychiatry*, 85(5), 425–433. <https://doi.org/10.1016/j.biopsych.2018.06.018>
- Patzelt, E. H., Kool, W., Millner, A. J., & Gershman, S. J. (2019b). The transdiagnostic structure of mental effort avoidance. *Scientific Reports*, 9(1), 1689. <https://doi.org/10.1038/s41598-018-37802-1>
- Raab, H. A., & Hartley, C. A. (2018). The Development of Goal-Directed Decision-Making. In R.

W. Morris, A. M. Bornstein, & A. Shenhav (Eds.), *Understanding Goal-Directed Decision Making: Computations and Circuits* ((Eds), pp. 1–31). New York: Elsevier.

<https://doi.org/10.1016/B978-0-12-812098-9.00013-9>

Raven, J. (2000). The Raven's Progressive Matrices: Change and Stability over Culture and Time. *Cognitive Psychology*, 41(1), 1–48. <https://doi.org/10.1006/cogp.1999.0735>

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. Abingdon, UK: Routledge.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208.

<https://doi.org/10.1037/h0061626>

Touron, D. R., Swaim, E. T., & Hertzog, C. (2007). Moderation of older adults' retrieval reluctance through task instructions and monetary incentives. *Journals of Gerontology - Series B Psychological Sciences and Social Sciences*, 62(3).

<https://doi.org/10.1093/geronb/62.3.P149>

Figure 1: A. Schematic showing how a metacontroller compares the computational cost against the reward benefit for model-based and model-free controllers. This cost-benefit analysis produces a weighting parameter (w) that expresses the probability of selecting model-based control (MB_{control}).

Figure 2: A. At the start of each trial, participants informed about the stakes, and then made a choice between two rockets. This choice deterministically transitioned the participant to a red or purple planet and they received an alien treasure reward. B. State transition structure of the novel two-step task. Rewards changed gradually over trials on the task according to a Gaussian random walk.

A. Experiment 1

B. Experiment 2

Figure 3: A. The first panel is a coefficient plot containing 95% of the posterior probability density around the mean for the age factor; younger and older adults both increase model-based control in response to incentives. The second panel is the posterior probability distribution of *age group* $\times$ *stakes* interaction, which includes zero and therefore indicates lack of strong evidence for an effect. B. The first panel is a coefficient plot containing 95% of the posterior probability density around the mean for the age factor; younger and older adults both increase model-based control in response to incentives; younger adults increase model-based control in response to incentives but older adults do not (in contrast to experiment 1). The second panel is posterior probability distribution that does not include zero, and this suggests high confidence in the presence of an *age group* $\times$ *stakes* interaction.