Valence bias in metacontrol of decision making in adolescents and young adults

Florian Bolenz (1) & Ben Eppinger (1,2,3)

1: Faculty of Psychology, Technische Universität Dresden, Dresden, Germany

2: Department of Psychology, Concordia University, Montreal, Canada

3: PERFORM centre, Concordia University, Montreal, Canada

Keywords: arbitration, development, reinforcement learning, effort, framing

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 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. We thank Andrea Reiter and Alexa Ruel for helpful feedback on this manuscript. We are grateful to Annkathrin Gomon, Annekathrin Rätsch, David Reindel, Johanna Dicke, Jan Hubert, Maria Schießl and Theresa Kolb for assistance during data acquisition.

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

Abstract

Metacontrol refers to the human ability to dynamically adapt decision-making strategies to changes in internal and external demands. In this study, we investigated the development of metacontrol from adolescence into young adulthood as well as developmental differences in the sensitivity of metacontrol to framing effects. Adolescents and young adults were assessed with a decision-making task that dissociates model-free and model-based decision-making strategies. In this task, we manipulated outcome magnitude and outcome valence, i.e. the framing of outcomes. With increasing age, we found a greater adaptation of model-based decision making to outcome magnitudes. Model-based decision making was more pronounced for loss compared to gain frames but this framing effect did not differ with age. Our findings suggest that metacontrol continues to develop into young adulthood. While losses generally increase the motivation to invest cognitive resources into an effortful decision-making strategy, the development of metacontrol is not sensitive to framing effects.

Introduction

During adolescence, the ability to engage in more complex decision-making strategies increases (Raab & Hartley, 2019). However, the successful use of a given decision-making strategy does not only depend on the mere ability to engage in it - it also depends on how flexible individuals are in adjusting their reliance on decision-making strategies to changes in internal and external demands. In this study, we ask how the ability for metacontrol of decision making (i.e. the dynamic adaptation of decision-making strategies) develops from adolescence into young adulthood and whether framing effects differentially affect the flexible usage of decision-making strategies in adolescents as compared to young adults.

To study metacontrol we draw on previous work that dissociates two major decision-making strategies: model-based and model-free decision making (Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Dayan & Niv, 2008). Model-based decision making represents a deliberative, prospective strategy that evaluates different choice options by means of forward planning based on knowledge about the structure of the environment (a cognitive model). In contrast, model-free decision making represents a more reflexive, retrospective strategy that relies on previously experienced action-reward contingencies. Previous developmental research shows that the reliance on model-based decision making (but not model-free decision making) becomes more pronounced from childhood to adulthood (Decker, Otto, Daw, & Hartley, 2016; Potter, Bryce, & Hartley, 2017) and these findings mirror the development of executive functions (Munakata, Snyder, & Chatham, 2012) which model-based decision making is thought to rely on (Otto, Gershman, Markman, & Daw, 2013; Otto, Skatova, Madlon-Kay, & Daw, 2015).

Due to this reliance on executive functions, model-based decision making is effortful and carries an intrinsic cost (Kool, McGuire, Rosen, & Botvinick, 2010). In return, it provides a higher behavioral flexibility than model-free decision making because dynamic changes in the environment can be accounted for more quickly (Dayan & Niv, 2008). Recent theories of metacontrol have proposed that humans weigh the cost and benefits associated with a decision-making strategy against each other when controlling their reliance on these strategies (Kool, Cushman, & Gershman, 2019; Lieder & Griffiths, 2017). According to this view, the cognitive effort associated with employing model-based decision making therefore needs to be matched by

sufficiently high benefits. This is in line with findings demonstrating that adults show increased reliance on model-based decision making when stakes are high; thus obtaining a higher monetary pay-off when engaging in a model-based decision-making strategy. (Kool, Gershman, & Cushman, 2017). Conversely, the engagement in model-based decision making is reduced when the task at hand gets more complex and thus this strategy becomes cognitively more effortful (Bolenz, Kool, Reiter, & Eppinger, 2019; Eppinger, Walter, & Li, 2017; Kool, Gershman, & Cushman, 2018). Since the strategic adaptation of physical and cognitive effort improves from adolescence into adulthood (Insel, Kastman, Glenn, & Somerville, 2017; Rodman et al., 2019), metacontrol of decision making may improve with increasing age.

So far, research on the interplay of model-free and model-based decision-making strategies has almost exclusively focused on decisions in the gain domain. The initial work on the dissociation of model-free and model-based decision making used rewards as outcomes to reinforce behavior (e.g., Daw et al., 2011; Gläscher, Daw, Dayan, & O'Doherty, 2010) and the studies on metacontrol applied rewards of varying magnitudes (e.g., Kool et al., 2017; Kool et al., 2018). This focus on decisions in the gain domain is surprising given behavioral and neuroscientific studies emphasizing the role of outcome valence in decision making, i.e. the framing of outcomes as gains or losses. Losses carry a higher subjective weight than gains in decisions under risk (Tversky & Kahneman, 1981) and losses have been associated with a greater allocation of cognitive resources (Yechiam & Hochman, 2013). Evidence from neuroscience studies supports the idea of dissociable mechanisms involved in decisions in the gain domain compared to the loss domain. For example, electrophysiological and neuroimaging work point to a unique role of the dorsal anterior cingulate cortex in the processing of monetary loss as well as in mediating effortful behavioral adjustments in response to negative feedback (Fischer & Ullsperger, 2013; Holroyd et al., 2004; Ullsperger, Danielmeier, & Jocham, 2014). Taken together, these findings suggest that losses may bias the cost-benefit evaluations underlying metacontrol and point to partially dissociable cognitive and neural processes underlying decisions in the gain and loss domain.

The potential impact of losses on decision making is interesting from a developmental perspective because of several findings that suggest that the neural systems involved in the processing of loss outcomes and the associated behavioral adaptations continue to develop into adolescence

(Crone & Steinbeis, 2017; Kelly et al., 2009; Rubia et al., 2006). The question whether adolescents and adults differ in how they attribute subjective weight to losses as compared to gains is not fully answered. Results of a developmental neuroimaging study on loss aversion during descriptive decision making show differences in fronto-striatal activation during decision making between adolescents and adults but no differences in behavioral preferences (Barkley-Levenson, Van Leijenhorst, & Galvan, 2013). In contrast, results from studies using the Iowa Gambling Task suggest that performance deficits in children and adolescents may result from a disproportionate tendency to shift behavior after receiving loss feedback when compared to young adults (Cassotti, Houde, & Moutier, 2011). Moreover, several studies have reported different developmental trajectories for gains as compared to losses with respect to effects on behavior and neural processing. For example, the neural sensitivity for differences in monetary outcomes follows independent developmental curves during adolescence in gains versus loss contexts (Insel & Somerville, 2018) and adolescents show a stronger asymmetry between gains and loss than adults during probabilistic reinforcement learning (Palminteri, Kilford, Coricelli, & Blakemore, 2016).

Here, we investigated metacontrol and its susceptibility to framing effect across development from adolescence to young adulthood by means of a decision-making task that dissociates model-free and model-based decision making. In this task, we manipulated outcome magnitude (low vs. high) and outcome valence (gains vs. losses). We expected that with increasing age, participants would show more metacontrol with respect to different outcome magnitudes and that the development of metacontrol differs between gain and loss contexts.

Results

We used a sequential decision-making task (adapted from Kool, Cushman, and Gershman (2016)) to determine reliance on model-free and model-based decision making in a sample of 97 adolescents (12-17 years) and 104 adults (18-25 years; see Figure 1 for age distribution). In this task, participants make repeated decisions between two spaceships - either between an orange and a turquoise spaceship or between a blue and a green spaceship. Each spaceship deterministically leads to one of two planets where the participants obtain a number of outcomes (see Figure 2A). The number of outcomes drifts over the course of the experiment (see Figure 2A).

Figure 1. Sample distribution of age and sex.

To understand how model-free and model-based decision making differ in this task, consider that a model-free decision-maker makes choices based on an independent reward expectation for each of the four spaceships. In contrast, a model-based decision-maker predicts the planet to which a spaceship will lead and uses the reward expectation associated with this planet to guide choices. Crucially, because multiple spaceships lead to one planet, the model-based decision-maker can integrate over experiences with different spaceships.

To investigate how the framing of outcomes affects model-based decision making, we manipulated the outcome valence across blocks of trials: In some blocks, outcomes were framed as gains (“space treasure”) and participants were instructed to maximize reward. In other blocks, outcomes were framed as losses (“antimatter”) and participants were instructed to avoid losing. Moreover, similar to previous studies (e.g., Kool et al., 2017), we manipulated how many points were at stake in each trial: In low-stakes trials, participants gained/lost one point for each outcome they had collected and in high-stakes trials participants gained/lost five points for each outcome (Figure 2B).

Figure 2. Decision making task. (A) Task transition structure. In every trial, one of two pairs of spaceships is presented. Each spaceship deterministically leads to one of two planets where a number of outcomes is obtained. The number of outcomes available at the two planets drifts over the course of the task. (B) Trial structure. At the beginning of a trial, the stakes condition is cued which determines how the outcomes of this trial are converted into points. During some blocks, outcomes are framed as gains and during the other blocks, outcomes are framed as losses.

To analyze how reliance on model-based decision making differed between experimental conditions and across development, we used a hierarchical Bayesian version of an established reinforcement-learning model (Daw et al., 2011; Kool et al., 2016). In our analyses, we focused on the model-based weight parameter that reflects the degree to which choices are guided by model-based decision making (for results regarding model-free weights, see Table S1-S2 and Figure S1 in the supplementary information). We used two variants of the reinforcement-learning model, a non-developmental model and a developmental model. In the non-developmental model, model-based weights were regressed on stakes condition and valence condition; thus,

regression weights reflect the effects of the experimental manipulation across all participants regardless of age. In the developmental model, model-based weights were regressed on stakes condition, valence condition and participants' age, which allows to test how model-based weights and the effects of the experimental conditions change across development. We will report results from the developmental model for all regression weights that relate to age and from the non-developmental model otherwise.

We will first report analyses on developmental changes in model-based decision making. As a second step, we will turn to developmental changes in stakes-based metacontrol, that is in the adaptation of model-based decision making to differences in outcome magnitudes. Third, we will present the effects of outcome valence on model-based decision making and metacontrol and how these effects changed with development.

Figure 3. Metacontrol of model-based decision making and its development across adolescence. (A) Mean model-based weights across all participants. Error bars represent Bayesian 95% credible intervals. (B) Model-based weights in all experimental conditions as a function of age. The shaded areas represent 95% credible intervals. Points represent means of the participant-level distributions (jittered along the x-axis).

Development of model-based decision making

Across conditions, model-based weights increased with age ($b_{\text{age}} = 0.21$, credible interval (CI) = [0.12, 0.30]). Separate analyses showed age-related increases in model-based weights for all

experimental conditions (Figure 3; $b_{\text{age}(\text{low}, \text{gain})} = 0.19$, $\text{CI} = [0.09, 0.30]$; $b_{\text{age}(\text{high}, \text{gain})} = 0.25$, $\text{CI} = [0.15, 0.35]$; $b_{\text{age}(\text{low}, \text{loss})} = 0.16$, $\text{CI} = [0.06, 0.26]$; $b_{\text{age}(\text{high}, \text{loss})} = 0.24$, $\text{CI} = [0.13, 0.34]$). These findings suggest that model-based decision making continues to develop across adolescence.

Development of stakes-based metacontrol

Across all participants, we found increased model-based weights for high-stakes trials compared to low-stakes trials ($b = 0.09$, $\text{CI} = [0.07, 0.12]$; Figure 3). Consistent with previous non-developmental studies (e.g., Kool et al., 2017), this indicates that participants showed stakes-based metacontrol, i.e. they adapted their reliance on model-based decision making to differences in outcome magnitudes. Separate analyses for adolescents and adults showed stakes-based metacontrol in both age groups (see Table S3 and Figure S2 in the supplementary information).

To investigate whether stakes-based metacontrol changes across adolescence, we analyzed stakes effects as a function of age. We found that the stakes effect increased with age ($b_{\text{age} \times \text{stakes}} = 0.07$, $\text{CI} = [-0.01, 0.15]$, 94.9% of the posterior mass above 0). Furthermore, separate analyses for gain and loss blocks showed at least moderate evidence for age-related increases in stakes effects for both valence conditions ($b_{\text{age} \times \text{stakes}(\text{gain})} = 0.06$, $\text{CI} = [-0.03, 0.15]$, 88.4% of the posterior mass above 0; $b_{\text{age} \times \text{stakes}(\text{loss})} = 0.08$, $\text{CI} = [-0.02, 0.17]$, 94.2% of the posterior mass above 0). These results show that metacontrol of decision making becomes more pronounced with age across adolescence.

As metacontrol could be related to the ability to engage in model-based decision-making (i.e., younger participants could show less metacontrol because they are generally more constrained in their ability to engage in model-based decision making), we analyzed if higher overall reliance on model-based decision making correlated with metacontrol. Our results show that there was no or at most only a small positive correlation between mean model-based weights and stakes effects ($r = -0.03$, $\text{CI} = [-0.20, 0.13]$, 35.8% of the posterior mass above 0; Figure 4). This suggests that increases in metacontrol do not necessarily result from a greater ability to engage in model-based decision making.

Figure 4. Relationship between average model-based weight and stakes effect. Dashed lines indicate 95 % credible interval. Points represent means of participant-level distributions.

Valence bias in metacontrol

Across all participants, we found higher model-based weights in loss blocks compared to gain blocks ($b_{\text{valence}} = -0.03$, $\text{CI} = [-0.05, -0.01]$; Figure 3). This indicates that participants relied more on model-based decision making when outcomes were framed as losses than when outcomes were framed as gains. Moreover, we found an interaction of stakes condition and valence condition ($b_{\text{stakes} \times \text{valence}} = 0.06$, $\text{CI} = [0.03, 0.09]$), indicating that the effect of outcome valence differed between low-stakes trials and high-stakes trials. While we found increased model-based weights in loss blocks compared to gain blocks for low-stakes trials ($b_{\text{valence}(\text{low})} = -0.06$, $\text{CI} = [-0.08, -0.03]$), we did not find evidence for a valence effect for high-stakes trials ($b_{\text{valence}(\text{high})} = 0.00$, $\text{CI} = [-0.02, 0.02]$). Thus, loss-induced increases of model-based decision making were particularly pronounced in trials with low outcome magnitudes. Separate analyses for adolescents and adults indicated interactions of stakes condition and valence condition in both age groups (see Table S3 and Figure S2 in the supplementary information).

In order to examine developmental differences in how gains and losses affect model-based decision making, we analyzed the effect of outcome valence as well as the interaction of stakes condition and valence condition as a function of age. We did not find evidence for an age-related change in the effect of outcome valence on model-based weights ($b_{\text{age} \times \text{valence}} = 0.02$, $\text{CI} = [-0.03, 0.07]$, 82.0% of the posterior mass above 0). Moreover, there was no evidence for an age-related change in the interaction of stakes condition and valence condition ($b_{\text{age} \times \text{stakes} \times \text{valence}} = -0.02$, $\text{CI} = [-$

0.11, 0.08], 34.3% of the posterior mass below 0). Thus, we did not find developmental differences in the asymmetry of gains and loss with respect to metacontrol of decision making.

Discussion

In this study, we investigated the development of metacontrol of decision making across adolescence. Specifically, we asked how the adaptation of model-based decision making toward different outcome magnitudes changes from adolescence into young adulthood and whether metacontrol in adolescents and adults is differentially sensitive to framing effects. We show that model-based decision making and its adaptation toward different outcome magnitudes (stakes-based metacontrol) improves with age. We also show that metacontrol of decision making is sensitive to outcome valence, i.e. the framing of outcomes as gains or losses. However, we do not find evidence for developmental differences in this valence bias.

Development of model-based decision making

The reliance on model-based decision making increased across adolescence. This finding is consistent with results of previous work that has shown an age-related increase in model-based decision making from childhood (age 8 years) to young adulthood (Decker et al., 2016; Potter et al., 2017). Going beyond these findings, we observe developmental differences in model-based decision making in a more constrained age range. Thus, the current data suggest that model-based decision making is a truly late-developing capacity that continues to develop into early adulthood. This could be linked to the protracted maturation of prefrontal brain regions (Casey, Tottenham, Liston, & Durston, 2005), which are implicated in model-based decision making (Gläscher et al., 2010; Huang, Yapple, & Yu, 2020; Smittenaar, FitzGerald, Romei, Wright, & Dolan, 2013).

Development of stakes-based metacontrol

Both adolescents and adults showed stakes-based metacontrol, that is they adapted their reliance on model-based decision making toward the different outcome magnitudes. However, stakes-based metacontrol became more pronounced with age across adolescence: Older participants showed stronger increases in model-based decision making when outcomes were amplified. This finding extends the results from a recent study reporting that stakes-based metacontrol increased

with age in five- to eleven-year-old children (Smid, Kool, Hauser, & Steinbeis, 2020), an age range which is adjacent to our study. Furthermore, the increased adaptation of model-based decision making parallels findings that the strategic adaptation of physical and cognitive effort toward different reward magnitudes improves from adolescence into adulthood (Insel et al., 2017; Rodman et al., 2019).

We interpret this result as showing that the ability for metacontrol continues to develop during adolescence. Together with studies in other age groups (Bolenz et al., 2019; Smid et al., 2020), this suggests that metacontrol of decision making is a cognitive process which is subject to dynamic developmental changes across the lifespan. However, one could also think of potential alternative interpretations for these developmental differences. First, less adaptation of model-based decision making toward different outcome magnitudes might be explained by an attenuated sensitivity to differences in outcomes; that is, the difference in subjective values for low-stakes and high-stakes outcomes might be reduced for adolescents. However, findings from neuroimaging studies speak against this interpretation by showing an increased sensitivity for differences in rewards in adolescents (Insel & Somerville, 2018). Second, the reduced capacity for model-based decision making in younger participants might have constrained the degree to which adaptation is possible. This would mirror suggestions that higher incentives only boost performance when the task demands match an individual's cognitive capacities (Davidow, Insel, & Somerville, 2018). However, we did not find that participants who showed more model-based decision making also showed more stakes-based metacontrol. Thus, these findings suggest that the development of model-based decision making is independent from the development of stakes-based metacontrol. This is in line with findings that varying the difficulty of cognitive-control tasks does not necessarily affect the adaptation of cognitive effort (Devine et al., in prep.), suggesting that the adaptation of effort is to some degree independent from one's capacity for the task. A dissociated development of model-based decision making and metacontrol is also consistent with the view that different brain regions underlie these two cognitive functions: While prefrontal cortex and hippocampus are assumed to be implicated in model-based decision making (Gläscher et al., 2010; Huang et al., 2020; Smittenaar et al., 2013; Vikbladh et al., 2019), dorsal anterior cingulate cortex has been proposed to be involved in the implementation of cost-benefit evaluations (Shenhav, Botvinick, & Cohen, 2013) that are considered to underlie metacontrol

(Kool et al., 2019; Kool et al., 2017). We therefore think that our finding is best explained by a process-specific, age-related improvement of metacontrol. Future research should shed more light on the neurodevelopmental trajectories that underlie age-related changes in metacontrol.

Valence bias in metacontrol

When avoiding losses, participants showed more reliance on model-based decision making than when seeking gains. It is commonly assumed that losses carry a higher subjective weight than gains (Tversky & Kahneman, 1981). This higher subjective weight of losses might increase the expected benefit of investing in an effortful decision-making strategy and thus affect cost-benefit evaluations that have been hypothesized to underlie metacontrol of decision making (Kool et al., 2019; Kool et al., 2017). In short, the motivation to avoid losses makes people allocate more cognitive resources to a more accurate but also more demanding decision-making strategy.

Previous studies have reported valence effects on model-based decision making in psychiatric patients or under pharmacological interventions but not in healthy or placebo controls (Voon et al., 2015; Worbe et al., 2016). Importantly, these findings were based on a different decision-making task (Daw et al., 2011) that has been questioned to elicit metacontrol (Kool et al., 2016; Kool et al., 2017). Therefore, the valence effects reported in these studies might not be reflective of an actual modulation of metacontrol but might represent more general effects associated with altered psychiatric or physiological conditions. In contrast, by using a task that previously has been shown to elicit metacontrol (e.g., Kool et al., 2017) and by testing a sample of healthy participants, our study provides evidence that outcome valence typically modulates model-based decision making and its metacontrol.

We observed a valence bias when outcome magnitude was low but not when outcomes were amplified. One possible explanation for this differentiation could be that participants show ceiling performance already for high-stakes trials during gain blocks and thus might not be able to further increase model-based decision making when outcomes are framed as losses. In this case, a valence bias should also arise in high-stakes trials when making the task more demanding because higher costs would lead to a more selective use of model-based decision making.

Contrary to our expectations, we did not find developmental differences in this valence bias. While previous studies have reported independent developmental trajectories for processing gains and losses (Insel & Somerville, 2018) and stronger valence-related asymmetries for learning in adolescents (Palminteri et al., 2016), our results might suggest that these developmental effects are not necessarily reflected in higher-order decision-making processes. However, our task design also differed in some aspects from previous studies (Insel & Somerville, 2018; Palminteri et al., 2016): For instance, we employed a blockwise manipulation of outcome valence while in these other studies, the framing of outcomes as gains or losses changed between single trials. It remains a question for future research how these differences in task design might affect behavior in adolescents and adults.

Conclusions

In this study, we found that metacontrol continues to develop from adolescence into young adulthood. With increasing age, participants showed more adaptation of model-based decision making toward different outcome magnitudes. Moreover, metacontrol was sensitive to outcome valence: When outcomes were framed as losses, participants showed a greater willingness to engage in the cognitively more demanding, model-based decision-making strategy. However, contrary to our predictions, we did not find evidence for developmental differences in this valence bias, indicating that framing effects had a similar impact on the adjustment of decision strategies in adolescents and young adults.

Methods

Participants

97 adolescents (50 female, mean age: 14.5 years, age range: 12 to 17 years) and 104 adults (61 female, mean age: 21.3 years, age range: 18 to 25 years) took part in this study. We aimed for a larger sample size than previous developmental studies on model-based decision making because of our intention to investigate interaction effects and because we expected generally weaker effects due to the more restricted age range. The target sample size of around 200 participants was determined based on feasibility considerations. No participant was excluded from data analysis. Participants' age followed an approximately uniform distribution within the age range

(Figure 1). We compensated participants with 5 € per hour or course credit as a baseline compensation and with an additional performance-dependent bonus payment of 7 cents per 100 points in the sequential decision-making task (range: 3.99 € to 5.46 €). Beside the sequential decision-making task, participants also completed a task battery including a cognitive-control task and a risk-preference task as well as questionnaires on impulsivity, cognitive effort investment and real-world risk taking, but we do not report analyses on these data here. All participants and the parents of all underage participants provided written informed consent. The ethics committee of Technische Universität Dresden approved the study.

Sequential decision-making task

We adapted a sequential decision-making task used in previous studies (Kool et al., 2016; Kool et al., 2017) that dissociates model-free and model-based decision making. Every trial started with an inter-trial interval (black screen, 300 ms) after which a stakes cue was presented (“1x” or “5x”, randomly chosen with equal probability, 800 ms) that indicated the stakes condition of the current trial. Then, one of two first-stage states was presented showing two spaceships side by side (an orange and a turquoise spaceship or a blue and a green spaceship, randomly chosen with equal probability, 1500 ms). Participants selected one of the two spaceships using the F and J keys on a standard computer keyboard for the left and the right spaceship, respectively, with the side on which the spaceship appeared being balanced across the task in both first-stage states. Contingent on their choice, one of two second-stage states was presented subsequently (a red or a purple planet with a corresponding alien, 1500 ms). In each first-stage state, one spaceship deterministically led to the red planet and the other spaceship deterministically led to the purple planet with a fixed mapping between spaceships and planets throughout the task. Participants responded in the second-stage state by pressing the space bar. After this, the outcomes received in this trial were presented above the alien (“space treasure” or “antimatter”, depending on the valence condition, 800 ms). The number of outcomes available at the two planets was determined by independent Gaussian random walks (mean = 0, standard deviation = 2, reflecting boundaries at 0 and 9, values rounded to integers) for each block of trials. At the end of the trial, the presentation of received outcomes was replaced with a presentation of points obtained in this

trial (depending on the number of outcomes, valence condition and stakes condition, 800 ms). A point count was displayed in the top-right corner of the screen during each block.

The task consisted of 480 trials divided into eight blocks of 60 trials. We varied outcome valence between blocks in alternating order and participants were informed about the upcoming valence condition before the start of each block. In gain blocks, the outcomes available at the two planets were presented as “space treasure” and participants gained points with every piece of outcome. In loss blocks, outcomes were presented as “antimatter” and participants lost points with every piece of outcome. Participants were instructed to seek as many outcomes as possible during gain blocks but to avoid as many outcomes as possible during loss blocks. To keep the overall payoff comparable between gain blocks and loss blocks, participants started each loss block with 1600 points and each gain block with 0 points.

Moreover, we manipulated between trials how outcomes were converted into points. In low-stakes trials, participants gained one point for each piece of space treasure and lost one point for each piece of antimatter. In high-stakes trials, the value of outcomes was amplified: Participants gained five points for each piece of space treasure and lost five points for each piece of antimatter. This stakes manipulation has previously been shown to induce metacontrol of model-based decision making (Kool et al., 2017).

The timing of trials was constant irrespective of how fast participants responded during the first-stage state or the second-stage state. Thus, participants could not increase their reward rate over time by responding faster. As soon as a response was given, the selected option was highlighted for the remaining time of presentation. If no response was given during the presentation of the first-stage state or the second-stage state, the trial was canceled and the task proceeded with the next trial. In trials with missing responses, points were calculated based on the stakes condition of the trial and the worst possible number of outcomes given the respective valence condition (0 pieces of space treasure in gain blocks or 9 pieces of antimatter in loss blocks). Therefore, it was not possible to avoid losing points simply by not responding during loss blocks. We excluded all trials with missing responses from analysis (2.1% of all trials).

To reduce variability between participants due to differences in the task, we created four different trial sequences that determined the sequence of first-stage states, number of outcomes available

at the two planets, sequence of stakes conditions and sequence of valence conditions. The sequence of first-stage states and the number of available outcomes was identical for all four trial sequences. Stakes conditions and valence conditions were counterbalanced across trial sequences. We assigned these four trial sequences in a balanced way to participants across the age range.

Prior to starting the task, participants received extensive training on the outcome distribution, the mapping between spaceships and planets as well as the stakes manipulation. To ensure that participants understood of the mapping between spaceships and planets, they needed to correctly select the spaceship leading to a given planet ten times in a row before being able to proceed. To ensure that participants understood the stakes manipulation, they needed to correctly report the number of points given a stakes cue and a number of outcomes ten times in a row before being able to proceed. Furthermore, participants completed 15 practice trials in the gain condition and 15 practice trials in the loss condition.

Reinforcement-learning model

To identify reliance on different decision-making strategies, we used an established reinforcement-learning model (Daw et al., 2011; Kool et al., 2017; Otto, Raio, Chiang, Phelps, & Daw, 2013). This model maintains model-free and model-based reward expectations, $Q_{MF}(s_i, a_i)$ and $Q_{MB}(s_i, a_i)$ respectively, for each available action a_i in every state s_i at the i -th stage of the task and both reward expectations can guide choice behavior. Note that while first-stage states have two available actions (selecting one of the two spaceships), second-stage states have only one action available (accepting to receive the outcome).

Model-free reward expectations. At the beginning of a block, we set all model-free reward expectations to $Q_{MF}(s_i, a_i) = 4.5$ for gain blocks and to $Q_{MF}(s_i, a_i) = -4.5$ for loss blocks thereby assuming an average model-free reward expectations before any information on the outcomes had been observed. After experiencing the consequences of a choice (i.e., obtaining outcomes or transitioning to a new state), all model-free reward expectations were updated according to:

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

Here, α is the reward learning rate (ranging between 0 and 1) that determines how strongly the model-free reward expectation is shifted toward the most recent experiences. The eligibility trace $e(s_i, a_i)$ grades which state-action pairs are eligible for updating.

At the beginning of a trial, $e(s_i, a_i)$ was set to 0 for all state-action pairs; before updating model-free reward expectations, the eligibility trace for the immediately preceding state-action pair (s'_i, a'_i) was set to 1 and after updating, all eligibility traces were decayed by the eligibility trace parameter λ (bounded between 0 and 1). By means of this, model-free reward expectations in the first-stage state could be updated twice in a trial: first, directly after moving from the first-stage state to the second-stage state and second, after obtaining the outcome at the end of the trial.

The reward prediction error δ is computed as the difference between the expected reward (given by the reward expectation for the immediately preceding state-action pair) and the actually experienced reward (which is the sum of the immediately obtained reward r and the predicted future reward in the new state, given by the reward expectation for the following state-action pair):

$$\delta = r + Q_{MF}(s_{i+1}', a_{i+1}') - Q_{MF}(s'_i, a'_i)$$

Note that $r = 0$ after first-stage choice because there is no immediate reward in second-stage states and $Q_{MF}(s_{i+1}', a_{i+1}') = 0$ after second-stage choices because reward states are considered terminal.

Model-based reward expectations. Model-based reward expectations at the first stage are based on the reward expectations associated with the second-stage state to which a first-stage choice will lead. Formally, this is computed as follows:

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

Here, T reflects transition probabilities according to the task structure with $T(s_2 | s_1, a_1) = 1$ if selection a_1 in s_1 will lead to s_2 and $T(s_2 | s_1, a_1) = 0$ otherwise. By means of these transition probabilities, model-based decision making incorporates knowledge about the task structure (i.e., a model of the task).

Model-based reward expectations at the second stage are identical to model-free reward expectations because both reflect an estimate of the immediate outcome.

Choice rule. The probability for selecting a particular action in a given first-stage state $P(a_1'|s_1')$ was modeled with a softmax function that takes into account model-free and model-based reward expectations as well as perseverative tendencies.

$$P(a_1'|s_1') = \frac{\exp(\beta_{MF}Q_{MF}(s_1', a_1') + \beta_{MB}Q_{MB}(s_1', a_1') + \pi \times rep(a_1') + \rho \times resp(a_1'))}{\sum_{a_1} \exp(\beta_{MF}Q_{MF}(s_1', a_1) + \beta_{MB}Q_{MB}(s_1', a_1) + \pi \times rep(a_1) + \rho \times resp(a_1))}$$

The choice stickiness parameter π (unbounded) and the response stickiness parameter ρ (unbounded) capture repeating or switching behavior (positive or negative values, respectively) of choices (selection of a spaceship) and responses (key presses). The indicator variable $rep(a_1')$ is set to 1 when a_1' was the action selected in the previous trial (and 0 otherwise). The indicator variable $resp(a_1')$ is set to 1 when selecting a_1' requires the same key press as in the previous trial (and 0 otherwise).

The model-free weight and the model-based weight, β_{MF} and β_{MB} respectively, are unbounded parameters that reflect the degree to which decision making is guided by either strategy. Note that while more conventional formulations of the reinforcement-learning model (Daw et al., 2011; Kool et al., 2017) use a single weighting parameter that reflects the relative influence of model-based versus model-free decision making, the formulation we use here is algebraically equivalent (Otto, Raio, et al., 2013) and comes with less bounded parameters which facilitates hierarchical model-fitting.

We fitted two variants of the reinforcement-learning model to the behavioral data. First, we fitted a non-developmental model in order to investigate differences in model-free and model-based decision making between the experimental conditions across all participants regardless of age. In this model variant, we regressed model-free and model-based weights on stakes condition, valence condition and their interaction (here illustrated for the model-based weight, the same procedure applies to the model-free weight):

$$\beta_{MB} = b_0 + b_{stakes}x_{stakes} + b_{valence}x_{valence} + b_{stakes \times valence}x_{stakes}x_{valence}$$

We used effect coding for the categorical variables ($x_{stakes} = -0.5$ for low-stakes trials and $x_{stakes} = 0.5$ for high-stakes trials; $x_{valence} = 0.5$ for gain blocks and $x_{valence} = -0.5$ for loss blocks). By means of this coding style, we can interpret b_0 as average model-based weight across all conditions and b_{stakes} , $b_{valence}$ and $b_{stakes \times valence}$ as main and interaction effects similar to an analysis of variance.

To analyze developmental effects, we also used an extended version of this model. In this developmental model, model-free and model-based weights were regressed on stakes condition, valence condition, age (scaled to the range [0, 1]) and all potential two-way and three-way interactions:

$$\beta_{MB} = b_0 + b_{stakes}x_{stakes} + b_{valence}x_{valence} + b_{stakes \times valence}x_{stakes}x_{valence} + age(b_{age} + b_{stakes \times age}x_{stakes} + b_{valence \times age}x_{valence} + b_{stakes \times valence \times age}x_{stakes}x_{valence})$$

In the developmental model, b_0 , b_{stakes} , $b_{valence}$ and $b_{stakes \times valence}$ can be interpreted as the mean, the main effect and the interaction effect for participants at the lower bound of the age range while b_{age} , $b_{stakes \times age}$, $b_{valence \times age}$ and $b_{stakes \times valence \times age}$ represent how the mean, the main effects and the interaction effect change when moving from the lower bound to the upper bound of the age range. We report results from the developmental model for all effects that include age (b_{age} , $b_{stakes \times age}$, $b_{valence \times age}$ and $b_{stakes \times valence \times age}$) and from the non-developmental model otherwise (b_0 , b_{stakes} , $b_{valence}$ and $b_{stakes \times valence}$).

Model-fitting procedure. We fitted a hierarchical Bayesian version of the reinforcement-learning model to the behavioral data using the probabilistic programming language Stan via the RStan package (Stan Development Team, 2018). All participant-level parameters were modeled as being drawn from a group-level distribution. Reward learning rate, the choice stickiness parameter and the response stickiness parameter were estimated separately for gain and loss blocks and we transformed bounded parameters (reward learning rate) to logit space. Participant-level parameters for these model parameters were modeled as being drawn from independent, univariate normal distribution with a group-level mean and a group-level standard deviation (e.g., $\pi_{gain} \sim N(\mu_{\pi_{gain}}, \sigma_{\pi_{gain}})$ or $\text{logit}(\alpha_{gain}) \sim N(\mu_{\text{logit}(\alpha_{gain})}, \sigma_{\text{logit}(\alpha_{gain})})$). To reduce model complexity, the eligibility trace parameter was set to $\lambda = 1$ and our results remain qualitatively unaffected when assuming other values for this parameter (see Table S4-S5 in the supplementary information). The linear combination weights $\vec{b} = (b_0, b_{stakes}, b_{valence}, b_{stakes \times valence})$ that determine model-based

weights were modeled as being drawn from a multivariate normal distribution with a vector of group-level means and a group level covariance matrix ($\vec{b} \sim N(\vec{\mu}_b, \Sigma_b)$) with $\Sigma_b = \text{DPD}$, where P is a correlation matrix and D is a diagonal matrix with $\text{diag}(D) = \vec{\sigma}_b$). The identical procedure applied to linear combination weights for the model-free weights.

For all group-level means, we used Normal(0,1) priors. For all group-level standard deviations, we used Normal(0,2) priors. For correlations, we used LKJ(2) priors. For the developmental parameters (b_{age} , $b_{stakes \times age}$, $b_{valence \times age}$ and $b_{stakes \times valence \times age}$), we used Normal(0,1) priors.

We ran four independent chains of the Markov Chain Monte Carlo procedure with 4000 iterations each and discarded the first 2000 iterations of each chain as warm-up. We assessed the convergence of all chains via the Gelman-Rubin statistic ($\hat{R} < 1.1$) for all parameters and ensured that bulk and tail effective sample size was sufficiently high (> 400) for all parameters.

When presenting group-level parameters, we report means and 95% credible intervals of the marginal distributions. Credible intervals are computed as the [.025, .975] percentile interval and can be interpreted as including the true value of the parameter of interest with a probability of 95%. When presenting participant-level parameters, we report means of the marginal participant-level distributions.

Data availability

All data and analysis scripts can be found at osf.io/a7ek9.

References

- Barkley-Levenson, E. E., Van Leijenhorst, L., & Galvan, A. (2013). Behavioral and neural correlates of loss aversion and risk avoidance in adolescents and adults. *Dev Cogn Neurosci*, 3, 72-83. doi:10.1016/j.dcn.2012.09.007
- Bolenz, F., Kool, W., Reiter, A., & Eppinger, B. (2019). Metacontrol of decision-making strategies in human aging. *Elife*, 8. doi:10.7554/eLife.49154
- Casey, B. J., Tottenham, N., Liston, C., & Durston, S. (2005). Imaging the developing brain: what have we learned about cognitive development? *Trends in Cognitive Sciences*, 9(3), 104-110. doi:10.1016/j.tics.2005.01.011
- Cassotti, M., Houde, O., & Moutier, S. (2011). Developmental changes of win-stay and loss-shift strategies in decision making. *Child Neuropsychol*, 17(4), 400-411. doi:10.1080/09297049.2010.547463
- Crone, E. A., & Steinbeis, N. (2017). Neural Perspectives on Cognitive Control Development during Childhood and Adolescence. *Trends Cogn Sci*, 21(3), 205-215. doi:10.1016/j.tics.2017.01.003
- Davidow, J. Y., Insel, C., & Somerville, L. H. (2018). Adolescent Development of Value-Guided Goal Pursuit. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2018.05.003
- 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
- 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. *Psychol Sci*, 27(6), 848-858. doi:10.1177/0956797616639301
- Eppinger, B., Walter, M., & Li, S. C. (2017). Electrophysiological correlates reflect the integration of model-based and model-free decision information. *Cogn Affect Behav Neurosci*, 17(2), 406-421. doi:10.3758/s13415-016-0487-3
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, 79(6), 1243-1255. doi:10.1016/j.neuron.2013.07.006

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

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G., & Cohen, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nat Neurosci*, 7(5), 497-498. doi:10.1038/nn1238

Huang, Y., Yaple, Z. A., & Yu, R. (2020). Goal-oriented and habitual decisions: Neural signatures of model-based and model-free learning. *Neuroimage*, 215, 116834. doi:10.1016/j.neuroimage.2020.116834

Insel, C., Kastman, E. K., Glenn, C. R., & Somerville, L. H. (2017). Development of corticostriatal connectivity constrains goal-directed behavior during adolescence. *Nat Commun*, 8(1), 1605. doi:10.1038/s41467-017-01369-8

Insel, C., & Somerville, L. H. (2018). Asymmetric neural tracking of gain and loss magnitude during adolescence. *Soc Cogn Affect Neurosci*, 13(8), 785-796. doi:10.1093/scan/nsy058

Kelly, A. M., Di Martino, A., Uddin, L. Q., Shehzad, Z., Gee, D. G., Reiss, P. T., . . . Milham, M. P. (2009). Development of anterior cingulate functional connectivity from late childhood to early adulthood. *Cereb Cortex*, 19(3), 640-657. doi:10.1093/cercor/bhn117

Kool, W., Cushman, F., & Gershman, S. J. (2019). Competition and Cooperation Between Multiple Reinforcement Learning Systems. In R. Morris, A. Bornstein, & A. Shenhav (Eds.), *Goal-Directed Decision Making: Computations and Neural Circuits* (pp. 153-178). New York: Academic Press.

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

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *J Exp Psychol Gen*, 139(4), 665-682. doi:10.1037/a0020198

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychol Rev*, 124(6), 762-794. doi:10.1037/rev0000075

Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing Cognitive Control: Three Key Transitions. *Curr Dir Psychol Sci*, 21(2), 71-77. doi:10.1177/0963721412436807

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., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proc Natl Acad Sci U S A*, 110(52), 20941-20946. doi:10.1073/pnas.1312011110

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

Palminteri, S., Kilford, E. J., Coricelli, G., & Blakemore, S. J. (2016). The Computational Development of Reinforcement Learning during Adolescence. *PLoS Comput Biol*, 12(6), e1004953. doi:10.1371/journal.pcbi.1004953

Potter, T. C. S., Bryce, N. V., & Hartley, C. A. (2017). Cognitive components underpinning the development of model-based learning. *Dev Cogn Neurosci*, 25, 272-280. doi:10.1016/j.dcn.2016.10.005

Raab, H. A., & Hartley, C. A. (2019). The Development of Goal-Directed Decision-Making. In R. Morris, A. Bornstein, & A. Shenhav (Eds.), *Goal-Directed Decision Making: Computations and Neural Circuits* (pp. 279-308). New York: Academic Press.

Rodman, A. M., Powers, K. E., Insel, C., Kastman, E. K., Kabotyanski, K. E., Stark, A. M., . . . Somerville, L. H. (2019). How adolescents and adults translate motivational value to action: Age-related shifts in strategic physical effort exertion for monetary rewards. *PsyArXiv*. doi:10.31234/osf.io/zcw4s

Rubia, K., Smith, A. B., Woolley, J., Nosarti, C., Heyman, I., Taylor, E., & Brammer, M. (2006). Progressive increase of frontostriatal brain activation from childhood to adulthood during event-related tasks of cognitive control. *Hum Brain Mapp*, 27(12), 973-993. doi:10.1002/hbm.20237

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217-240. doi:10.1016/j.neuron.2013.07.007

Smid, C. R., Kool, W., Hauser, T. U., & Steinbeis, N. (2020). Model-based decision-making and its metacontrol in childhood. *PsyArXiv*. doi:10.31234/osf.io/ervsb

Smittenaar, P., FitzGerald, T. H., Romei, V., Wright, N. D., & Dolan, R. J. (2013). Disruption of dorsolateral prefrontal cortex decreases model-based in favor of model-free control in humans. *Neuron*, 80(4), 914-919. doi:10.1016/j.neuron.2013.08.009

Stan Development Team. (2018). RStan: the R interface to Stan (Version 2.19.2). Retrieved from <http://mc-stan.org>

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453-458. doi:10.1126/science.7455683

Ullsperger, M., Danielmeier, C., & Jocham, G. (2014). Neurophysiology of performance monitoring and adaptive behavior. *Physiol Rev*, 94(1), 35-79. doi:10.1152/physrev.00041.2012

Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., . . . Daw, N. D. (2019). Hippocampal Contributions to Model-Based Planning and Spatial Memory. *Neuron*. doi:10.1016/j.neuron.2019.02.014

Voon, V., Baek, K., Enander, J., Worbe, Y., Morris, L. S., Harrison, N. A., . . . Daw, N. (2015). Motivation and value influences in the relative balance of goal-directed and habitual behaviours in obsessive-compulsive disorder. *Transl Psychiatry*, 5, e670. doi:10.1038/tp.2015.165

Worbe, Y., Palminteri, S., Savulich, G., Daw, N. D., Fernandez-Egea, E., Robbins, T. W., & Voon, V. (2016). Valence-dependent influence of serotonin depletion on model-based choice strategy. *Mol Psychiatry*, 21(5), 624-629. doi:10.1038/mp.2015.46

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: review and analysis of the unique effects of losses over gains. *Psychol Bull*, 139(2), 497-518. doi:10.1037/a0029383

Valence bias in metacontrol of decision making in adolescents and young adults

Supplementary information

Table S1

Effects on model-free weights in the non-developmental model

Effect	b	CI
stakes	0.03	[0.01, 0.04]
valence	0.04	[0.02, 0.05]
stakes X valence	0.01	[-0.02, 0.04]

Table S2

Effects on model-free weight in the developmental model

Effect	b	CI
age	-0.01	[-0.06, 0.04]
age X stakes	0.05	[0.002, 0.09]
age X valence	0.01	[-0.04, 0.05]
age X stakes X valence	0.08	[-0.02, 0.17]

Table S3

Effects on model-based weights in the non-developmental model separately for both age groups

Effect	Adolescents		Young adults	
	b	CI	b	CI
stakes	0.08	[0.05, 0.11]	0.11	[0.07, 0.14]
valence	-0.04	[-0.07, -0.02]	-0.01	[-0.04, 0.01]
stakes X valence	0.07	[0.03, 0.11]	0.05	[0.002, 0.09]

Table S4

Effects on model-based weights in the non-developmental model assuming different parameter values for eligibility trace decay (λ)

Effect	$\lambda = 0$		$\lambda = 0.5$	
	b	CI	b	CI
stakes	0.10	[0.08, 0.12]	0.09	[0.07, 0.11]
valence	-0.02	[-0.04, -0.005]	-0.03	[-0.05, -0.02]
stakes X valence	0.06	[0.04, 0.09]	0.06	[0.03, 0.09]

Table S5

Effects on model-based weights in the developmental model assuming different parameter values for eligibility trace decay (λ)

Effect	$\lambda = 0$		$\lambda = 0.5$	
	b	CI	b	CI
age	0.21	[0.13, 0.29]	0.21	[0.13, 0.30]
age X stakes	0.09	[0.02, 0.17]	0.08	[-0.00, 0.16]
age X valence	0.01	[-0.03, 0.06]	0.01	[-0.04, 0.06]
age X stakes X valence	0.02	[-0.05, 0.10]	0.005	[-0.08, 0.09]

Figure S1. Metacontrol of model-free decision making and its development across adolescence. (A) Mean model-free weights across all participants. Error bars represent Bayesian 95% credible intervals. (B) Model-free weights in all experimental conditions as a function of age. The shaded areas represent 95% credible intervals. Points represent means of the participant-level distributions (jittered along the x-axis).

Figure S2. Mean model-based weights separately for both age groups. Error bars represent 95% credible intervals.