UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Impaired learning to dissociate advantageous and disadvantageous risky choices in adolescents

Jepma, M.; Schaaf, J.V.; Visser, I.; Huizenga, H.M.

DOI
10.1038/s41598-022-10100-7

Publication date
2022

Document Version
Final published version

Published in
Scientific Reports

License
CC BY

Link to publication

Citation for published version (APA):
Jepma, M., Schaaf, J. V., Visser, I., & Huizenga, H. M. (2022). Impaired learning to dissociate advantageous and disadvantageous risky choices in adolescents. *Scientific Reports*, *12*, Article 6490. https://doi.org/10.1038/s41598-022-10100-7

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:15 Oct 2026

scientific reports

Check for updates

OPEN

Impaired learning to dissociate advantageous and disadvantageous risky choices in adolescents

Marieke Jepma[1,2]✉, Jessica V. Schaaf[1,2], Ingmar Visser[1] & Hilde M. Huizenga[1]

Adolescence is characterized by a surge in maladaptive risk-taking behaviors, but whether and how this relates to developmental changes in experience-based learning is largely unknown. In this preregistered study, we addressed this issue using a novel task that allowed us to separate the learning-driven optimization of risky choice behavior over time from overall risk-taking tendencies. Adolescents (12–17 years old) learned to dissociate advantageous from disadvantageous risky choices less well than adults (20–35 years old), and this impairment was stronger in early than mid-late adolescents. Computational modeling revealed that adolescents' suboptimal performance was largely due to an inefficiency in core learning and choice processes. Specifically, adolescents used a simpler, suboptimal, expectation-updating process and a more stochastic choice policy. In addition, the modeling results suggested that adolescents, but not adults, overvalued the highest rewards. Finally, an exploratory latent-mixture model analysis indicated that a substantial proportion of the participants in each age group did not engage in experience-based learning but used a gambler's fallacy strategy, stressing the importance of analyzing individual differences. Our results help understand why adolescents tend to make more, and more persistent, maladaptive risky decisions than adults when the values of these decisions have to be learned from experience.

Although adolescence is a developmental period of good physical health and strength, it is associated with relatively high morbidity and mortality rates[1]. This downside of adolescence is largely due to an increased engagement in hazardous activities—such as dangerous behavior in traffic, substance and alcohol abuse, and unprotected sex[2–6]—which is often referred to as risk taking[7]. Indeed, risk taking has been defined in the developmental literature as the engagement in behaviors with potentially undesirable outcomes[8]. Importantly, however, some risky actions are also associated with beneficial outcomes that outweigh the potential negative ones, in which case risk taking can be advantageous. Participating in an international exchange program, for example, may result in failure and disappointment, but is more likely to promote personal growth and happiness; hence is a form of risk taking that is generally considered positive. Thus, to optimize positive and minimize negative life outcomes, adolescents' ability to dissociate advantageous from disadvantageous risky behaviors may be more important than their inclination to take risk per se.

In many daily-life situations, one's a priori knowledge about the potential outcomes of risky actions, and their probabilities, may be vague or incomplete. For example, although an adolescent may know that excessive alcohol consumption can have both positive and negative outcomes, he or she may not have a clear idea of the likelihood and severity of those outcomes. In situations like this, people can rely on experience-based learning to obtain a better estimate of risky actions' expected value (i.e., the sum of all possible outcomes multiplied by their respective probabilities). This learning process allows people to optimize their risky choice behavior over time. For instance, by repeatedly experiencing that binge drinking results in bad hangovers, someone may learn that these negative outcomes outweigh the positive ones and therefore decide to stop binge drinking.

There is extensive evidence that experience-based learning processes differ between adolescents and adults[9–16], and most of these studies suggest that adolescents are suboptimal learners[17]. Relating these findings to the risk-taking domain, it is an interesting possibility that developmental changes in learning may contribute to the increased incidence of maladaptive risk-taking behaviors in adolescents as compared to adults. Broadly consistent with this idea, two task features that have been found to promote risk taking in adolescents are also

[1]Department of Psychology, University of Amsterdam, Amsterdam, the Netherlands. [2]These authors contributed equally: Marieke Jepma and Jessica V. Schaaf. ✉email: mariekejepma@gmail.com

key ingredients of experience-based learning: Uncertainty about outcome probabilities (i.e., ambiguity)[18–21] and the presence of outcome feedback following each choice[7,22,23].

Many developmental studies have examined risk taking independent from learning using 'decisions from description' tasks that involve choices between two options—e.g., gambles or lotteries—that differ in their degree of risk[24–30]. Experience-based learning is not required in these tasks because participants are provided with complete information about all possible outcomes and their corresponding probabilities, creating a situation of explicit or known risk. Recent studies have also used tasks with incomplete information about the risky option's outcome probabilities, creating a situation of unknown risk or ambiguity[18–21]. However, as these tasks did not include outcome feedback, ambiguity reduction through experience-based learning was not possible. Together, the results from these studies suggest that, in choice situations that do not involve learning, adolescents make similar choices as adults when choice options differ in known risk, but are more likely to choose ambiguous options.

In another class of risk-taking tasks—including the Balloon Analogue Risk Task (BART)[31]—participants decide when to stop a series of increasingly risky choices, in a stepwise manner, and receive positive or negative outcome feedback following each choice. Some studies have found that adolescents take more risk than adults on these tasks[22,23]. The optimal stopping point in the BART can in principle be learned from experience over the course of the task. However, with two recent exceptions[32,33], developmental studies using this task focused on average measures of risk taking across all trials, leaving open whether and how developmental changes in experience-based learning may contribute to age-related differences in risk taking.

This question has been addressed, however, by studies using the Iowa Gambling task (IGT)[34] or child-friendly versions of this task[35], which measure risk taking in an experience-based learning context. Participants in the IGT make repeated choices between four options, represented as decks of cards. Two options are more risky (higher outcome variability) and have a low expected value, whereas the other two options are less risky (lower outcome variability) and have a high expected value. Participants do not receive prior information about the options' outcome magnitudes or probabilities, but have to learn which options are most beneficial through trial-and-error. Healthy adults typically develop a preference for (one of) the advantageous, less risky, choice options over the course of the task, and developmental studies have found that adolescents' choice behavior improves less over time[36–38]. These findings have been taken to suggest that adolescents, as compared to adults, are impaired at learning the negative value of risky choices, resulting in more persistent maladaptive risky choice behavior.

However, the IGT is a rather complex task, and people's performance on this task reflects multiple cognitive processes that are difficult to disentangle[7,39]. First, as the risky options are also the disadvantageous ones (i.e., those with lower expected value), it is difficult to determine whether a more persistent preference for these options reflects a stronger tendency to make risky choices, a weaker sensitivity to expected value, and/or an impairment in experience-based learning. Second, the IGT requires participants to track and integrate gain and loss magnitudes and probabilities for four different choice options in parallel, which places a high demand on working memory. As working-memory function and its underlying neural circuitry continue to mature into late adolescence[40–44], age-related differences in IGT performance may also reflect these developmental changes in working memory.

Present study. In the present preregistered study, we aimed to examine potential differences between adolescents (12–17 years old) and young adults (20–35 years old; matched with regard to educational level) in the learning component of risk taking, while avoiding the interpretation difficulties associated with the IGT. To this end, we developed a simple task that requires people to learn the expected value of a risky choice option based on its previously experienced outcomes. Participants in this task make repeated choices between two options: a 'sure' (riskless) option and a risky option. While choices for the sure option always result in a fixed small reward, choices for the risky option can result in either a larger reward or no reward at all. Because participants are repeatedly presented with the same choice options, and directly observe the outcome of each choice, they can learn the expected value of the risky option over time and adjust their choice behavior accordingly. To minimize working-memory demands, participants are fully informed about the reward magnitudes for the sure and the risky option, and about the sure option's reward probability (100%); hence the only variable that has to be learned is the risky option's reward probability. Importantly, the risky option's reward probability varies across task blocks, such that the expected value of the risky option is lower than, higher than, or identical to that of the sure option, in different blocks (referred to as risk-disadvantageous, risk-advantageous, and risk-neutral blocks, respectively). This aspect of the task allows us to dissociate effects of expected value and risk, and makes it more representative of real life in which risk taking is sometimes adaptive and sometimes maladaptive.

We expected participants to choose the risky option more often than the sure option during the first trials of each block, as this would allow them to estimate the risky option's expected value. Successful learning should then cause a decrease in risky choice behavior over trials in risk-disadvantageous blocks, but not in risk-advantageous blocks (a trial x block type interaction). Regarding developmental differences, we expected that adolescents would be less efficient learners than adults[17], and would therefore show a smaller adjustment of their risky choice behavior over trials as a function of the risky option's expected value (a trial x block type x age group interaction). Given evidence that adolescents are more tolerant to ambiguity than adults[18,19,21], we also expected adolescents to choose the risky option more often overall. Although we did not have specific hypotheses about the difference between early and mid-late adolescents in our task, work using related tasks has suggested differences between early and late adolescents[7]. To explore this issue, we administered our task to groups of early and mid-late adolescents.

We tested the hypotheses using multilevel regression. In addition to the regression analysis, we applied a set of computational models to participants' choice data. Whereas the regression analysis tests for the presence of

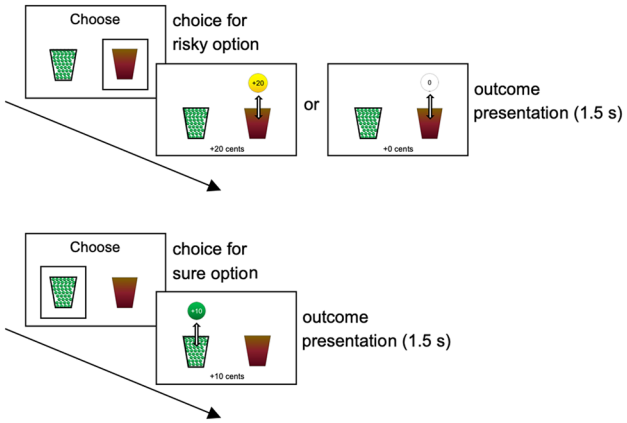

Figure 1. Trial outline in case of a risky (upper plot) and sure (lower plot) choice. This figure was created using Microsoft PowerPoint 16.16.27, https://www.microsoft.com/en-us/microsoft-365/powerpoint.

developmental differences in risk taking and learning efficiency, computational models of learning and decision-making can elucidate the nature of these differences. Specifically, we tested which out of several learning models—which differ in the complexity/optimality of the learning process and in how they can give rise to risk-sensitive behavior in our task—best explained each age group's choice data. When possible, we also examined whether the estimated model parameters differed between the age groups.

Results

Preregistration. We preregistered our main research questions, analyses, and exclusion criteria using AsPredicted (http://aspredicted.org/kw96t.pdf). Non-preregistered analyses are treated as exploratory, and indicated as such.

Experimental design. In each task block, participants made 20 choices between a sure and a risky option to earn monetary rewards. These options were depicted as vases filled with balls (Fig. 1; see Methods for task details). The 'sure' vase only contained balls with a '+10' label (worth €0.10). The 'risky' vase contained a mix of balls with a '0' label (worth €0) and balls with a '+20' label (worth €0.20), with unknown proportions. Participants completed ten blocks of this task. The proportion of risky choices that yielded a 20-cents outcome was either 0.3 (four blocks), 0.7 (four blocks), or 0.5 (two blocks), such that the expected value of the risky option was, respectively, lower than (6 cents), higher than (14 cents), or identical to (10 cents) that of the sure option. Participants were not informed about the risky option's expected value, but could acquire an estimate of this value over the course of each block based on the experienced outcomes.

To assess developmental trajectories of risk-taking behavior, we administered our task to groups of early adolescents (12–14-year-olds; N = 31), mid-late adolescents (15–17-year-olds; N = 39) and adults (20–35-year-olds; N = 35), and tested for linear and quadratic age effects across these three groups.

Many participants showed evidence of a gambler's fallacy-like strategy. We preregistered to exclude participants who switched to the sure option after experiencing a win outcome from the risky option on more than 30% of those trials, as this behavior is indicative of a gambler's fallacy instead of a learning strategy. This exclusion criterion resulted in the exclusion of 46% of the adults, 67% of the mid-late adolescents and 68% of the early adolescents. Although the proportion of participants showing gambler's fallacy-like behavior was numerically larger in the adolescent groups than in the adult group, a non-preregistered and thus exploratory analysis showed that this difference was not significant ($\chi^2(2, n = 105) = 4.5$, $p = 0.11$).

To verify that this behavior was maladaptive in the current task, we compared overall choice accuracy (proportion of choices for the option with the highest expected value, across all trials of the risk-advantageous and risk-disadvantageous blocks) for participants who showed vs. did not show this behavior. This non-preregistered and thus exploratory analysis showed that overall choice accuracy was lower in the participants who showed gambler's fallacy-like behavior (as defined by our exclusion criterion) than in the participants who did not (49% vs. 62%, $t(58.7) = 7.3$, $p < 0.001$), confirming the maladaptive nature of this behavior.

Next, we report three sets of analyses: (i) the preregistered regression analysis examining effects of trial, block type and age group on risky choices; (ii) the preregistered computational-modeling analyses (Methods; Supplemental Text 1); and (iii) a non-preregistered hence exploratory latent mixture-model analysis examining individual differences in strategy use (Methods; Supplemental Text 2). In the first two analyses we excluded all participants who met our preregistered exclusion criterion. In the third analysis, we included all participants. In addition, we repeated our preregistered regression analysis on a larger number of participants based on the mixture-model results (Supplemental Text 3).

The adaptive optimization of risky choice behavior improves from early adolescence to mid-late adolescence to adulthood. Participants made more risky choices when the risky option had

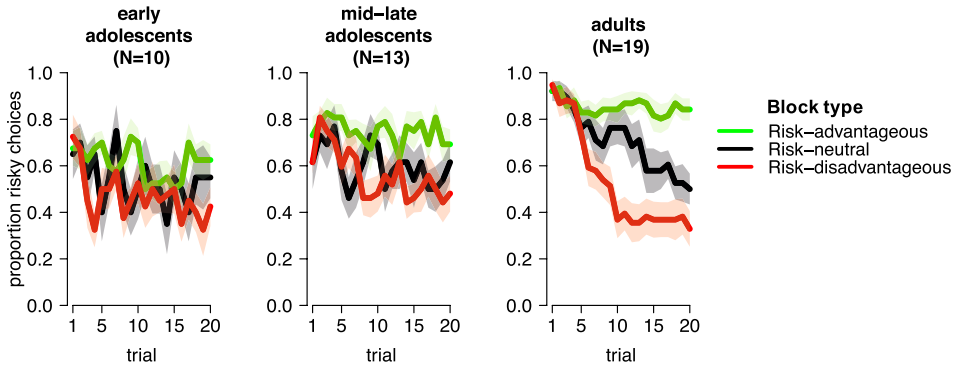

Figure 2. Mean proportion of choices for the risky option per trial, block type, and age group. Shaded areas represent 1 SEM. This figure was created using RStudio 1.1.463, https://www.rstudio.com.

a higher expected value (main effect of block type, $z = 19.2$, $p < 0.001$; Fig. 2). This effect emerged over trials in a nonlinear way (block type x trial-linear interaction, $z = 7.8$, $p < 0.001$; block type x trial-quadratic interaction, $z = 3.2$, $p = 0.002$), suggesting that participants learned to optimize their risk choice behavior over time. Regarding developmental differences, the speed at which participants optimized their choice behavior over trials increased across the three age groups in a linear way, as reflected in block type x trial-linear x age group-linear ($z = 6.1$, $p < 0.001$) and block type x trial-quadratic x age group-linear ($z = 3.6$, $p < 0.001$) interactions. Thus, adults optimized their risky choice behavior faster than mid-late adolescents, who in turn optimized faster than early adolescents.

The overall number of risky choices increased linearly across the three age groups as well (main effect of age group-linear, $z = 2.9$, $p = 0.004$). Thus, contrary to our hypothesis, the number of risky choices was lowest in the early adolescents and highest in the adults. Figure 2 shows that this was driven by more risky choices in older age groups when the risky option was advantageous and, in all conditions, during early trials.

As adolescents optimized their choice behavior rather poorly over trials, we performed additional, exploratory, analyses to test some alternative explanations for poor performance. Specifically, we ruled out that adolescents carried over the learned risk value from one block to the next (Supplemental Fig. 6), and that they stopped engaging as the task progressed (Supplemental Text 6).

Same analysis using a larger sample. To examine the robustness of the effects reported above, we repeated the analysis while including a larger number of participants—18 early adolescents, 25 mid-late adolescents and 28 adults—based on a less conservative exclusion criterion (see last Results section). Inclusion of this larger number of participants did not change the significance of the age-related effects reported above, with one exception: The overall number of risky choices no longer differed across the three age groups (main effect of age group-linear, $z = 1.7$, $p = 0.09$). However, the block type x trial x age group-linear interactions remained highly significant (p's < 0.001). Results from this non-preregistered, hence exploratory, analysis are fully reported in Supplemental Text 3 and Supplemental Fig. 1.

Computational modeling results. The regression results reported above suggest that the adaptation of risky choice behavior based on previously experienced outcomes improves between early and mid-late adolescence, and further improves between mid-late adolescence and adulthood. To examine whether these results reflect the use of different learning strategies, different parameter values of the same learning strategy, and/or differences in choice stochasticity across the different age groups, we applied a set of computational models to each age group's choice data, using a hierarchical Bayesian approach. Full descriptions of the models and their equations can be found in the Methods and Supplemental Text 1; here we provide a brief overview.

We used two broad families of learning models: Reinforcement-learning models (Models 1A-C) and Bayesian ideal-observer models (Models 2A-D). Bayesian ideal-observer models take into account the uncertainty of expectations, and are more sophisticated and flexible but also more computationally demanding than reinforcement-learning models. Based on evidence that adolescents use suboptimal learning strategies[14,17], we reasoned that adolescents' choices may be better described by reinforcement-learning models.

Within each model family, we compared different model versions (Table 1). The first version of each model family (Models 1A and 2A) does not contain an explicit mechanism to explain risk-sensitive behaviors. The second version of each model family (Models 1B and 2B) can account for risk-sensitive behavior via asymmetric learning from win and no-win outcomes[9,45]. The third version of each model family (Models 1C and 2C) can account for risk-sensitive behavior via nonlinear subjective utilities for different outcome magnitudes[46]. Finally, we included a fourth version of the Bayesian ideal-observer model (Model 2D) which can account for risk-sensitive behavior via positive or negative effects of uncertainty on expected value[47,48]. This last version cannot be implemented in reinforcement-learning models, as these do not represent uncertainty.

We combined all learning models with a softmax decision function, which translates expected values into choice probabilities. The 'inverse-temperature' parameter controls the sensitivity of choice probabilities to

	Parameters	
	Reinforcement learning (1)	**Bayesian ideal-observer (2)**
Basic models (A)	Learning rate α	Update rate π
Asymmetric learning (B): Stronger weighting of win outcomes promotes risk seeking	Learning rates for win and no-win outcomes: α_+ and α_-	Update rates for win and no-win outcomes: π_+ and π_-
Nonlinear utility function (C): Overvaluation of higher outcomes promotes risk seeking	Utility parameter κ $\kappa > 1$ and $\kappa < 1$ cause over- and undervaluation of higher outcomes, respectively	Utility parameter κ
Uncertainty affects value (D): Uncertainty bonus promotes risk seeking	Not applicable	Uncertainty parameter φ $\varphi > 0$ and $\varphi < 0$ cause uncertainty bonus and penalty, respectively

Table 1. Overview of the learning models. The initial value of the risky option (Q_1) was a free parameter in all models as well (not included in the table).

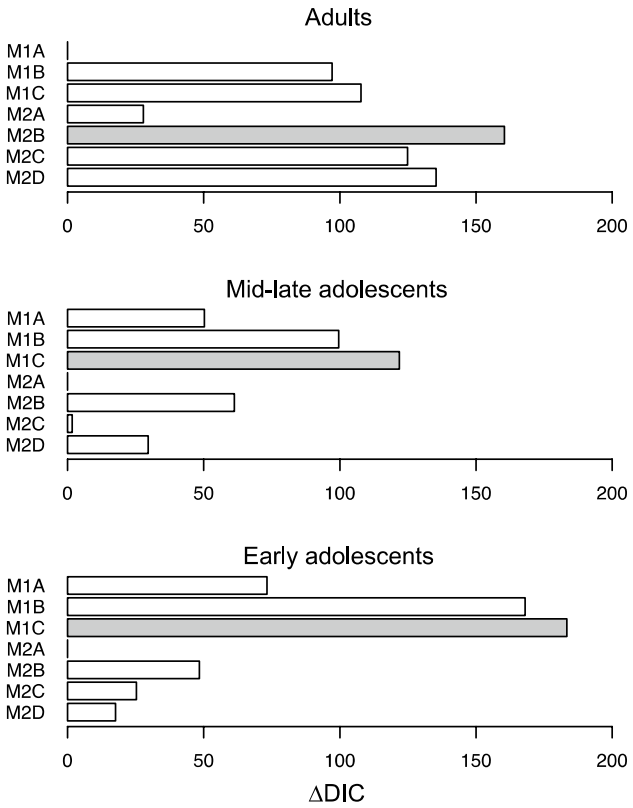

Figure 3. Model comparison. ΔDIC values reflect each model's goodness of fit relative to the worst-fitting model. The winning model in each age group is indicated in gray. M1A = basic reinforcement-learning, M1B = reinforcement-learning with 2 learning rates, M1C = reinforcement-learning with nonlinear utility function; M2A = basic Bayesian ideal-observer, M2B = Bayesian ideal-observer with 2 update rates, M2C = Bayesian ideal-observer with nonlinear utility function, M2D = Bayesian ideal-observer with uncertainty bonus. This figure was created using RStudio 1.1.463, https://www.rstudio.com.

differences in expected value. As the value of the inverse-temperature parameter increases, the probability that the option with the highest expected value is chosen also increases (i.e., choice stochasticity decreases).

To verify that our models were distinguishable and thus that potential developmental differences are meaningful[49], we performed a model-recovery analysis. This analysis indicated that our procedures were able to distinguish between reinforcement-learning and Bayesian ideal-observer models and between different risk-sensitive mechanisms with rather high accuracy (Supplemental Text 4; Supplemental Fig. 5).

Model comparison. We computed ΔDIC values by subtracting each model's DIC from the DIC of the worst-fitting model, such that higher ΔDIC values indicate a better fit (Fig. 3). For the adults, the Bayesian ideal-observer models outperformed their corresponding reinforcement-learning models. Conversely, for the early and mid-late adolescents' data, the reinforcement-learning models outperformed their corresponding Bayesian

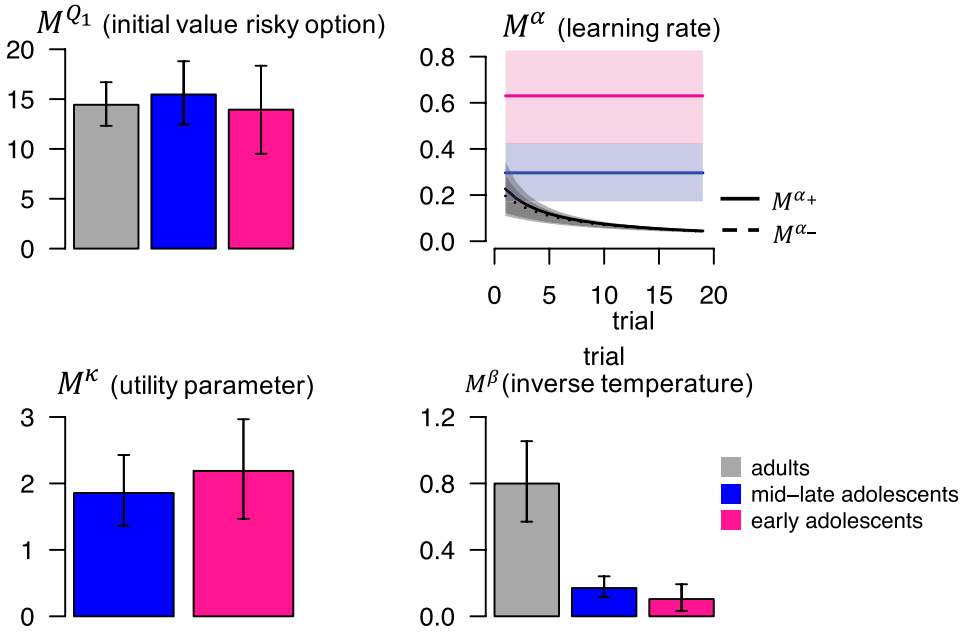

Figure 4. Posterior medians of the winning models' group-level mean parameters, per age group. Error bars and shaded areas indicate 95% HDIs of the posterior distributions. This figure was created using RStudio 1.1.463, https://www.rstudio.com. Note Q_1 is expressed in cents. For the adult group, Q_1 was computed using the posterior median of the initial a parameter (a_1) of the winning Bayesian ideal-observer model, and trial-specific learning rates α_+ and α_- were computed using the posterior medians of the update-rate parameters (π_+ and π_-) of the winning Bayesian ideal-observer model (Supplemental Text 1). The lines for α_+ and α_- almost overlap.

ideal-observer models. This suggests that the adults reduced their learning rate over trials (within each block) as their expected-value estimates became more certain, whereas the adolescents used a constant learning rate. Note that reducing one's learning rate over trials is optimal when stimulus-outcome contingencies are stable, as in the current task[14]. Thus, the adolescents used a simpler, suboptimal, learning process than the adults. Furthermore, models that included a mechanism to account for risk-sensitive behavior (models 1B-C and 2B-D) performed better than their corresponding basic models (models 1A and 2A), in all age groups. The winning model for the adults was the Bayesian ideal-observer model with two update rates (Model 2B), and the winning model for both adolescent groups was the reinforcement-learning model with nonlinear utility function (Model 1C).

Supplemental Fig. 2 illustrates the fit of the winning models for each age group. We inspected the differences in model fit between the two model families visually and also compared the absolute deviation between the simulated and observed proportion of risky choices (Supplemental Fig. 2). Both inspections did not provide clear clues as to why the reinforcement-learning model with nonlinear utility function fitted the data better in both adolescent groups and the Bayesian ideal-observer model with asymmetric updating fitted better in the adults. We speculate that the difference in DIC values is due to individual differences in strategy use, as supported by our latent-mixture model analysis (see below).

Parameter estimates. To shed more light on the learning and choice mechanisms implemented in the winning model for each age group, we next examined the group-level mean (M) parameters of these models. Figure 4 shows the medians and 95% highest density intervals (HDIs) of these parameters' posterior distributions. Supplemental Fig. 3 shows the full posterior distributions.

The posterior distribution of the risky option's initial expected value (M^{Q_1}) was higher than 10—the value of the sure option—in all age groups. Such 'optimistic initial values' for the risky option promote exploration of this option, in line with participants' initial preference for the risky option (Fig. 2).

The posterior distribution of the learning-rate parameter (M^{α}) was higher for the early adolescents than for the mid-late adolescents (99.6% of the difference distribution lay above 0), suggesting that the early adolescents' value estimate for the risky option was driven more by the most recent outcome. The adults' data was best explained by a dynamic learning rate, which started out at a similar value as the mid-late adolescents' learning rate and then decreased over trials. Although the best-fitting model for the adults contained separate update rates for positive and negative outcomes, the group-level means of these parameters were highly similar (Fig. 4, solid vs. dotted black line in upper right plot). This suggests that some adults learned more from positive than negative outcomes, while others showed the opposite bias, but that there was no systematic learning asymmetry (at the individual level, π_+ was higher than π_- for 9 adults, and π_+ was lower than π_- for 10 adults).

The adolescents' data was best explained by a model with nonlinear utility function, and the posterior distribution of the parameter governing the shape of the utility function (M^{κ}) lay above 1 for both adolescent groups. This suggests that adolescents' subjective value of a 20-cent outcome was more than twice the subjective value

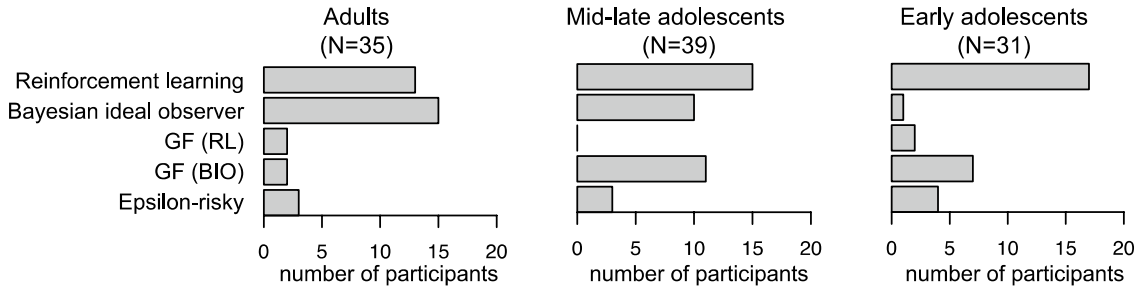

Figure 5. Inferred strategy use according to our mixture-model analysis. The plots show the number of participants per age group assigned to each of the 5 models included in the mixture model. GF = gambler's fallacy, RL = reinforcement learning, BIO = Bayesian ideal observer. This figure was created using RStudio 1.1.463, https://www.rstudio.com.

of a 10-cent outcome (convex utility function). This overvaluation of the higher outcomes promotes risk taking, which in our task should impair performance in the risk-disadvantageous blocks but improve performance in the risk-advantageous blocks. However, an advantage for the adolescents over the adults in the risk-advantageous blocks was not apparent in the choice data (Fig. 2), likely due to adolescents' higher choice stochasticity, as described below.

The posterior distribution of the inverse-temperature parameter (M^{β}) was higher in the adults than in both adolescent groups, but did not differ between the early and mid-late adolescents (9% of the difference distribution lay above 0). The lower inverse temperature in the adolescents than the adults may reflect that that adolescents' choices were driven less by the options' expected-value estimates (i.e., they were less prone to choose the option with the highest expected-value estimate), reflecting more stochastic choice behavior. Alternatively, the lower inverse temperature in the adolescents may reflect that the value-updating algorithms implemented in our models were less able to capture the adolescents' choices. Adolescents' lower inverse temperature can explain why they made fewer risky choices during the initial trials of a block: Even though the risky option's initial expected value was similarly optimistic in all age groups, adolescents' choices were less sensitive to expected values, which resulted in a weaker initial choice preference for the risky option. Adolescents' low inverse temperature may also explain why their overvaluation of the higher win outcomes (convex utility function)—which should increase the expected value of the risky option—was not translated in a substantial increase in risky choices.

In sum, our modeling results identified two mechanisms underlying the impaired adaptation of risky choice behavior over time in adolescents as compared to adults: (i) adolescents used a simpler, suboptimal, learning process involving a constant learning rate; and (ii) adolescents showed a higher degree of choice stochasticity (lower inverse temperature). In addition, the results suggest that the adolescents, but not the adults, overvalued the highest outcomes. The impaired performance of the early compared to the mid-late adolescents may be due to the elevated learning rate (causing excessive expectation updating) and numerically lower inverse temperature in the early adolescents.

Results of latent-mixture model analysis including all participants. We performed an additional (not preregistered hence exploratory) latent mixture-model analysis to infer the learning or non-learning strategy used by each individual participant, and to examine whether the prevalence of specific strategies differed between our three age groups (Supplemental Text 2). In this analysis, we included all participants.

Our mixture model included the two winning learning models from the previous analysis (reinforcement-learning model with nonlinear utility function and Bayesian ideal-observer model with two update rates). In addition, it included two gambler's fallacy models. These were identical to the two learning models except that they used negative learning/update rates. Thus, the gambler's fallacy models assume that the value of the risky option decreases following each win outcome and increases following each no-win outcome, reflecting the belief that outcomes that occurred more frequent in the past will be less likely in the future. Finally, we included a fifth model that does not consider the experienced outcomes. This model—the epsilon-risky model—has a fixed probability of choosing the risky option on each trial (determined by parameter ε), hence can account for guessing behavior and general tendencies to seek or avoid risk that are insensitive to outcome feedback.

Figure 5 shows the number of participants in each age group assigned to each model (Supplemental Fig. 4 shows the model assignment per participant). The proportion of participants assigned to a learning model (reinforcement learning or Bayesian ideal-observer) did not differ across the three age groups ($\chi^2(2, n = 105) = 4.0$, $p = 0.14$).

For the participants assigned to a learning model, the proportion of participants assigned to the reinforcement-learning vs. Bayesian ideal-observer model differed across the three age groups ($\chi^2(2, n = 71) = 11.1$, $p = 0.004$). While the proportion of participants assigned to the Bayesian ideal-observer model was highest in the adults and lowest in the early adolescents, the proportion of participants assigned to the reinforcement-learning model showed the opposite pattern, corroborating the model-comparison results from the previous analysis.

Finally, a few participants in each age group were assigned to the epsilon-risky model, suggesting that their tendency to choose the risky option was constant over time and independent of the experienced outcomes.

Comparison of mixture-model results and preregistered exclusion criterion. The number of participants assigned to a learning model by our mixture-model analysis (58% of the early adolescents, 64% of the mid-late adolescents, and 80% of the adults) was higher than the number of participants who used a learning strategy according to our preregistered criterion (32% of the early adolescents, 33% of the mid-late adolescents, and 54% of the adults). Our preregistered criterion assumed that participants who frequently switched to the sure option after a win outcome from the risky option did not use a learning strategy. However, frequent win-switch behavior does not necessarily reflect a lack of learning, but could also be due to a high degree of choice stochasticity, which may explain the discrepancy between our preregistered criterion and mixture-model results. As our computational models were designed to tease apart learning and choice processes, these are arguably better-suited to dissociate learners from non-learners than our rather crude preregistered criterion. Therefore, we repeated the regression analysis on risky choice behavior while including all participants who learned according to our mixture-model analysis (see 'Same analysis using a larger sample' section above, and Supplemental Text 3).

Discussion

Adolescence is characterized by an increase in maladaptive risk-taking behaviors, but whether and how this relates to developmental changes in experience-based learning is largely unknown. We addressed this question using a simple learning task involving repeated choices between a sure and a risky option. Importantly, the expected value of the risky option—which could be learned from experience—varied across task blocks, allowing us to dissociate the learning-driven optimization of risky choice behavior from general risk-taking tendencies.

Our exploratory mixture-model analysis suggested that approximately two-thirds of the participants used a learning strategy, while the remaining participants used a gambler's fallacy or other irrelevant choice strategy. Thus, a substantial proportion of the participants in each age group did not engage in the experience-based, incremental, type of learning that our task was designed to measure. This highlights the importance of considering individual differences in strategy use, for example using mixture models, before averaging data across participants. It also raises the questions whether and how the use of suboptimal (non-learning) strategies can be prevented, and to what degree such strategies contribute to maladaptive risk-taking in other learning tasks (such as the IGT) and in real life.

After excluding participants who used a gambler's fallacy-like strategy (see Limitations and future directions section below), we found that adolescents adapted their risky choice behavior less to the expected value of this behavior than adults. Computational modeling provided evidence that both adolescent groups, compared to the adult group, used (i) a simpler, suboptimal, learning process (involving a constant learning rate) and (ii) a more stochastic choice policy. The modeling results also suggested that adolescents, but not adults, overvalued the highest rewards. However, this last effect was overshadowed by adolescents' larger choice stochasticity such that a corresponding increase in risk taking was not observed. Together, these findings suggest that adolescents' suboptimal risky choice behavior was largely due to an inefficiency in core learning and choice processes, and did not reflect increased risk seeking per se.

We found that the optimization of risky choice behavior improves during adolescence as well. Early adolescents performed worse than mid-late adolescents, which could be explained by early adolescents' higher constant learning rate (causing excessive expectation updating) and more stochastic choice policy. Our findings thus suggest that the "settings" of learning and choice parameters improve between early and mid-late adolescence, and that further optimization—the transition to a Bayesian learning process (involving a dynamic learning rate) and a further decrease in choice stochasticity—occurs between mid-late adolescence and adulthood.

Our results are consistent with findings from the IGT that adolescents develop a weaker preference for the advantageous choice options than adults[36–38], and suggest that adolescents' suboptimal IGT performance may be related to (i) suboptimal learning strategies, (ii) enhanced choice stochasticity, and/or (iii) overvaluation of higher rewards. Our results are also consistent with two previous studies that examined the role of learning in adolescent risk taking on the BART[32,33]. Adolescents in these studies showed a weaker optimization of their risk-taking behavior (number of pumps) over the course of the task than adults, in line with our findings, suggesting that this developmental effect generalizes across different risk-taking paradigms.

There is an interesting discrepancy between our results and those from previous developmental studies examining 'decisions from description'[18,19,21]. These previous studies, which did not involve learning, found that adolescents were more willing to choose options with unknown outcome probabilities than adults, which was attributed to adolescents' larger tolerance for ambiguity. In contrast, during the first trial of each block in our study (i.e., before learning had occurred), adolescents chose the option with unknown outcome probabilities less often than adults. These opposite developmental effects suggest that age-related differences in ambiguity attitude are not fixed across situations, but depend on whether or not ambiguity can be reduced through learning (also see[50]). In our task, choosing the risky option during the first trials allowed participants to quickly estimate its expected value and thereby optimize future choices. That adolescents did this less often than adults is thus indicative of suboptimal, less goal-directed, (learning) behavior. Our modeling results suggested that this behavior was not due to a lower initial valuation of the risky option, but reflected adolescents' more stochastic choice policy which made them less prone to choose the option with the highest expected value.

A higher degree of choice stochasticity in younger age groups—as reflected in a lower inverse-temperature parameter—is a common finding in experience-based learning tasks[10,13,14,38,51,52]. This has been attributed to children's and adolescents' increased tendency to explore lower-valued but more uncertain options[15]. However, in our task, the adolescents chose the risky (and more uncertain) option less often than the adults, which is inconsistent with the notion of increased exploration. Instead, adolescents' more stochastic choice behavior in our task may simply reflect a higher degree of choice randomness, possibly related to attentional or motivational lapses. Alternatively, adolescents' more stochastic choice behavior may also reflect an increased tendency to

alternate between the choice options, regardless of how rewarding they are. Finally, it could be that adolescents more often held specific, idiosyncratic, beliefs about the nature of the task (e.g., the belief that a series of two good outcomes will be followed by a bad outcome), preventing them from learning the true values of the risky option and leading to an apparently high degree of choice stochasticity.

Limitations and future directions. An important open question concerns the degree to which the developmental effects found in our task generalize to real-life situations. Recent evidence suggests that developmental changes in learning are most pronounced in more demanding tasks, such as those with a higher cognitive load[53,54]. This suggests that the age-related differences in learning and choice mechanisms revealed in our relatively simple task may have even stronger effects on risk taking in more complex environments, such as those encountered in real life. However, most real-life risk-taking scenarios also involve social and emotional factors that were not captured by our task. For example, adolescents (but not adults) have been shown to take more risk when observed by peers, which has been attributed to a heightened sensitivity to the potential rewards of risky actions[55–57]. Thus, peer pressure likely increases risk seeking in adolescents via reward-valuation mechanisms not identified in our study. Whether peer pressure also affects adolescents' learning process, utility function, and/or choice stochasticity—and thereby the ability to optimize risky decisions over time—is an interesting question for future studies.

Another limitation of our task is that the worst possible outcome was gaining nothing (0 cents), but participants could never lose money. Like several previous studies[18,19,24,45], we thus focused on risk taking in the gain domain. As people are generally risk-averse in the gain domain but risk-seeking in the loss domain[58,59], we expect that participants in a loss version of our task would make more risky choices overall. Whether outcome valence also affects people's learning and choice processes, and whether outcome-valence effects change during development, is currently unknown. To address these questions, future developmental studies could compare gain and loss versions of our task.

Finally, due to the substantial proportion of participants engaging in irrelevant choice strategies, the sample size of our computational analyses was rather small, potentially affecting the reliability of our results[60]. However, one of the advantages of our hierarchical modeling approach, and especially our Bayesian framework[61,62], is that valid results can even be obtained with small sample sizes[63]. This theoretical account in combination with rather high model recoverability (Supplemental Text 4) strengthen the reliability of our results, regardless of the small sample size. Nevertheless, future studies are advised to replicate our results using larger sample sizes, especially when aimed at further investigation of individual differences (e.g.,[64]).

To conclude, by combining a novel experience-based risk-taking task with computational modeling, we demonstrated how adolescents' simpler learning process and more stochastic choice policy can give rise to an impaired optimization of risky choice behavior over time. Future work could examine the generalizability of our findings to risk-taking behaviors in social contexts and in the loss domain, which may eventually inform interventions aimed at reducing the detrimental outcomes associated with adolescents' maladaptive forms of real-life risk taking.

Methods

Participants. A total of 35 adults (mean age = 22.6; age range = 20–35; 63% female), 39 mid-late adolescents (mean age = 15.8; age range = 15–17; 74% female), and 31 early adolescents (mean age = 13.0; age range = 12–14; 48% female) completed the study. The adults were students, or former students, at universities or colleges of higher professional education. The mid-late adolescents were in the fourth or fifth year of high school, and the early adolescents were in the first or second year of high school (all pre-university or higher general secondary education).

Participants reported no history of psychiatric or neurological disorders. Adult participants received a fixed amount of €2.50 or course credits for their participation, plus a variable performance-dependent amount of maximally €2.50. Adolescent participants only received the variable performance-dependent amount of maximally €2.50 (the school we collaborated with did not allow larger rewards). All participants of age 16 and older provided written informed consent. Primary caretakers of participants younger than 16 were informed about the experiment and provided active written informed consent. All procedures were approved by the ethics committee of the Faculty of Social and Behavioural Sciences of the University of Amsterdam, and the study was performed in accordance with the relevant guidelines and regulations.

We preregistered to exclude participants who switched to the sure option after experiencing a win outcome from the risky option on more than 30% of those trials, as we reasoned that this behavior is indicative of a gambler's fallacy instead of a learning strategy, reflecting a misunderstanding of the task. Based on this criterion, we excluded 16 adults (46%), 26 mid-late adolescents (67%), and 21 early adolescents (68%), leaving a final sample of 19 adults, 13 mid-late adolescents, and 10 early adolescents. We performed our preregistered analyses on this sample. In addition, we performed a non-preregistered, hence exploratory, latent mixture-model analysis (described in Supplemental Text 2) including all participants. This analysis suggested that 28 adults, 25 mid-late adolescents, and 18 early adolescents used a learning, instead of a gamblers fallacy, strategy—a larger number than suggested by our preregistered criterion. Therefore, we also repeated our behavioral analysis on this larger sample.

General procedure. Adolescent participants were tested in a classroom or computer room at their high school, in groups of approximately 20 participants. Adult participants were tested in a room at the university, in groups of at least two participants. At least one experimenter was always present in the testing room as well. Participants performed the task individually on a laptop or PC. Before starting the task, participants received com-

puterized task instructions and performed a short practice block. The task lasted approximately 15 to 20 min, after which participants were reimbursed.

Experience-based risk-taking task. In each task block, participants made 20 choices between two options to earn monetary rewards. Choices for one option always paid off 10 cents; we refer to this option as the sure option. Choices for the other option paid off either 20 or 0 cents, we refer to this option as the risky option.

The two options were depicted as vases, displayed at the left and right side of the screen (Fig. 1). The 'sure' vase was filled with green balls with a '+10' label (worth €0.10). This vase was transparent such that participants could see that it contained green balls only. The 'risky' vase was opaque such that its content was invisible, and participants were instructed that it contained a mix of white balls with a '0' label (worth €0) and gold balls with a '+20' label (worth €0.20). However, the proportion of white and gold balls in this vase was unknown. We instructed participants that each time they chose a vase, one ball would be drawn from that vase and displayed on the screen, and participants would gain the money associated with that ball. Participants were also instructed that after a ball was drawn, it was returned into its vase and all balls were shuffled for the next trial (i.e., random draws with replacement).

On each trial, participants chose a vase by pressing a left ('z') or a right ('/') key. A black frame appeared around the chosen vase and 200 ms later one ball was displayed above that vase. The ball moved upwards and then back downwards (back inside the vase) as if it was drawn from the vase and then returned. The ball-drawing animation lasted 1.5 s; during this period the payoff amount was also displayed at the bottom of the screen (+10 cents, +20 cents, or +0 cents).

Unbeknownst to the participants, the proportion of risky choices that yielded 20 cents was either 0.3 (four blocks), 0.7 (four blocks), or 0.5 (two blocks). Specifically, out of every ten risky choices, either three, seven, or five choices paid off 20 cents and the remaining choices paid off 0 cents, in random order. Thus, the expected value of the risky option was lower than (6 cents), higher than (14 cents), or identical to (10 cents) the expected value of the sure option (10 cents), respectively. We refer to these three block types as risk disadvantageous, risk advantageous, and risk neutral. Participants performed the ten blocks in random order, with one constraint: Both the first five and the last five blocks had to contain one risk-neutral, two risk-disadvantageous, and two risk-advantageous blocks. A new risky vase, with a unique color, was introduced in each block. We instructed participants that the number of gold and white balls inside the opaque vase would vary across blocks, but would not change during a block.

Except for the outcome probabilities of the risky option, participants were fully informed about the task structure and procedure. After completing the task, participants received the average amount of money they had gained across the ten blocks (i.e., total amount divided by ten, rounded to the nearest 50 cents). In addition, the adults also received their fixed payment of €2.50 or course credits.

Behavioral analysis. We performed a multilevel logistic regression analysis on the single-trial choice data using the lme4 package[65] in R. The dependent variable was the binary choice variable (risky vs. sure choices, coded as 1 vs. 0, respectively). We tested for effects of block type (coded as -1, 0 and 1 for the risk-disadvantageous, risk-neutral and risk-advantageous blocks), block-specific trial (linear and quadratic effects), age group (linear and quadratic effects), and their interactions. In addition to these fixed effects, we modeled random intercepts and random slopes for the within-subject fixed effects (separately for each age group). This model failed to converge. Therefore, as preregistered, we removed the random slopes and only modeled random intercepts for the within-subject fixed effects.

Computational models. We examined the ability of a range of learning models to capture each age group's trial-to-trial choice data. Supplemental Text 1 contains a detailed description of all models, the model equations, and our hierarchical Bayesian parameter estimation methods; here we provide an intuitive description of the models (also see Table 1).

All models estimate the expected value of the risky option—and update this estimate following each new outcome. Note that the expected value of the sure option did not have to be estimated, as its payoff was fixed and known in advance (100% probability of €0.10; hence its expected value is always 10 cents).

We used two broad families of learning models: Reinforcement-learning models (Q-learning models; Models 1A-C) and Bayesian ideal-observer models (beta-binomial models; Models 2A-D). Reinforcement-learning models represent expected values as point values, which are updated in response to prediction errors using a constant learning rate[66]. In contrast, Bayesian ideal-observer models represent expected win probabilities as beta distributions, naturally track both the mean and uncertainty of expectations[67–71]. These latter models determine their effective learning rate on each trial as a function of the current uncertainty (higher uncertainty results in stronger updating). Because they take into account the uncertainty of expectations, Bayesian ideal-observer models are more sophisticated and flexible but also more computationally demanding than reinforcement-learning models. Bayesian ideal-observer models have been shown to describe adults' choices in a reward-learning task better than reinforcement-learning models[68], but these two model families have not been compared in adolescent samples. Based on evidence that adolescents, as compared to adults, use suboptimal learning strategies[14,17], we reasoned that adolescents' choice data may be better described by reinforcement-learning models than Bayesian ideal-observer models.

Within each model family, we compared different model versions that vary in whether and how they can account for risk-sensitive behavior (Table 1). The first, most basic, version of each model family (Models 1A and 2A) does not contain an explicit mechanism to explain risk-avoidance or risk-seeking behaviors. These models have two free parameters: The learning or update rate which controls the speed of learning, and the initial

expected value of the risky option which captures the initial preference for the risky option (before any learning has taken place). The other versions are extended versions of these basic models.

The second version of each model family (Models 1B and 2B) can account for risk-sensitive behavior via asymmetric learning. Specifically, it contains separate learning/update rates for win and no-win outcomes[9,45]. If win outcomes are weighted stronger than no-win outcomes, this leads to an overestimation of the risky option's expected value, promoting risk seeking. The opposite learning asymmetry has the opposite effect, promoting risk avoidance.

The third version of each model family (Models 1C and 2C) does not allow for asymmetric learning, but allows for nonlinear subjective utilities for different win magnitudes[46]. The subjective utility of the risky option's win outcome (20 cents) can be higher or lower than 20 cents (convex vs. concave subjective utility curve, respectively), depending on the setting of a utility parameter[45]. Values of this parameter smaller than 1 produce undervaluation of the risky option's win outcomes, promoting risk aversion, whereas values of this parameter larger than 1 have the opposite effect.

Finally, we included a fourth version of the Bayesian ideal-observer model (Model 2D) which allows the value of the risky option to increase or decrease as a function of its uncertainty[47,48]. The expected value of the risky option is uncertain, especially during the first trials of each block, while the expected value of the sure option is not. Therefore, a positive effect of uncertainty (uncertainty bonus) will increase the value of the risky option and hence promote risky choices, whereas a negative effect of uncertainty (uncertainty penalty) will have the opposite effect. The direction and magnitude of the uncertainty effect is controlled by a free parameter. Note that an uncertainty bonus/penalty cannot be implemented in reinforcement-learning models, as these do not represent uncertainty.

Decision function. We combined all learning models with a softmax decision function, which translates the expected-value difference of the risky and the sure option into a probability of choosing the risky option. The sensitivity of choice probabilities to differences in expected value is controlled by an 'inverse-temperature' parameter. If this parameter is 0, both options are equally likely to be chosen, irrespective of their expected values. As the value of the inverse-temperature parameter increases, the probability that the option with the highest expected value is chosen also increases.

Model comparison. We compared the performance of the different models using the deviance information criterion (DIC)[72], separately for each age group. The DIC is a hierarchical modeling generalization of the AIC which is easily computed in hierarchical Bayesian model-selection problems using Markov chain Monte Carlo (MCMC) sampling. It provides an index of the goodness of fit of a model, penalized by its effective number of parameters. Models with smaller DIC are better supported by the data.

Exploratory latent-mixture model analysis. We performed an additional model-based analysis on all participants—including those who were excluded from the previous analyses—to examine the choice strategies used by each individual participant: A latent mixture-model analysis[73–75]. This non-preregistered hence exploratory analysis is described in Supplemental Text 2.

Received: 20 August 2021; Accepted: 25 March 2022
Published online: 20 April 2022

References

1. Dahl, R. E. Adolescent brain development: A period of vulnerabilities and opportunities. Keynote address. *Ann. N. Y. Acad. Sci.* **1021**, 1–22. https://doi.org/10.1196/annals.1308.001 (2004).
2. Windle, M. *et al.* Transitions into underage and problem drinking: Developmental processes and mechanisms between 10 and 15 years of age. *Pediatrics* **121**(Suppl 4), S273-289. https://doi.org/10.1542/peds.2007-2243C (2008).
3. Wakefield, M. A. *et al.* Effect of restrictions on smoking at home, at school, and in public places on teenage smoking: Cross sectional study. *BMJ* **321**, 333–337. https://doi.org/10.1136/bmj.321.7257.333 (2000).
4. Gittes, E. B. & Irwin, C. E. Jr. Sexually transmitted diseases in adolescents. *Pediatr. Rev.* **14**, 180–189. https://doi.org/10.1542/pir.14-5-180 (1993).
5. Jonah, B. A. Accident risk and risk-taking behaviour among young drivers. *Accid. Anal. Prev.* **18**, 255–271. https://doi.org/10.1016/0001-4575(86)90041-2 (1986).
6. Rai, A. A. *et al.* Relative influences of perceived parental monitoring and perceived peer involvement on adolescent risk behaviors: An analysis of six cross-sectional data sets. *J. Adolesc. Health* **33**, 108–118. https://doi.org/10.1016/s1054-139x(03)00179-4 (2003).
7. Defoe, I. N., Dubas, J. S., Figner, B. & van Aken, M. A. A meta-analysis on age differences in risky decision making: Adolescents versus children and adults. *Psychol. Bull.* **141**, 48–84. https://doi.org/10.1037/a0038088 (2015).
8. Boyer, T. W. The development of risk-taking: A multi-perspective review. *Dev. Rev.* **26**, 291–345 (2006).
9. van den Bos, W., Cohen, M. X., Kahnt, T. & Crone, E. A. Striatum-medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cereb. Cortex* **22**, 1247–1255. https://doi.org/10.1093/cercor/bhr198 (2012).
10. Palminteri, S., Kilford, E. J., Coricelli, G. & Blakemore, S. J. The computational development of reinforcement learning during adolescence. *PLoS Comput. Biol.* **12**, e1004953. https://doi.org/10.1371/journal.pcbi.1004953 (2016).
11. Cohen, J. R. *et al.* A unique adolescent response to reward prediction errors. *Nat. Neurosci.* **13**, 669–671. https://doi.org/10.1038/nn.2558 (2010).
12. Hauser, T. U., Iannaccone, R., Walitza, S., Brandeis, D. & Brem, S. Cognitive flexibility in adolescence: Neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *Neuroimage* **104**, 347–354. https://doi.org/10.1016/j.neuroimage.2014.09.018 (2015).

13. Javadi, A. H., Schmidt, D. H. K. & Smolka, M. N. Adolescents adapt more slowly than adults to varying reward contingencies. *J. Cogn. Neurosci.* **26**, 2670–2681. https://doi.org/10.1162/jocn_a_00677 (2014).
14. Jepma, M., Schaaf, J. V., Visser, I. & Huizenga, H. M. Uncertainty-driven regulation of learning and exploration in adolescents: A computational account. *PLoS Comput. Biol.* **16**, e1008276. https://doi.org/10.1371/journal.pcbi.1008276 (2020).
15. Bolenz, F., Reiter, A. M. F. & Eppinger, B. Developmental changes in learning: Computational mechanisms and social influences. *Front. Psychol.* **8**, 2048. https://doi.org/10.3389/fpsyg.2017.02048 (2017).
16. Byrnes, J. P., Miller, D. C. & Reynolds, M. Learning to make good decisions: A self-regulation perspective. *Child Dev.* **70**, 1121–1140 (1999).
17. Nussenbaum, K. & Hartley, C. A. Reinforcement learning across development: What insights can we draw from a decade of research?. *Dev. Cogn. Neurosci.* **40**, 100733. https://doi.org/10.1016/j.dcn.2019.100733 (2019).
18. Blankenstein, N. E., Crone, E. A., van den Bos, W. & van Duijvenvoorde, A. C. Dealing with uncertainty: Testing risk- and ambiguity-attitude across adolescence. *Dev. Neuropsychol.* **41**, 77–92. https://doi.org/10.1080/87565641.2016.1158265 (2016).
19. Tymula, A. *et al.* Adolescents' risk-taking behavior is driven by tolerance to ambiguity. *Proc. Natl. Acad. Sci. U S A* **109**, 17135–17140. https://doi.org/10.1073/pnas.1207144109 (2012).
20. Tymula, A., Rosenberg, B. L. A., Ruderman, L., Glimcher, P. W. & Levy, I. Like cognitive function, decision making across the life span shows profound age-related changes. *Proc. Natl. Acad. Sci. U S A* **110**, 17143–17148. https://doi.org/10.1073/pnas.1309909110 (2013).
21. van den Bos, W. & Hertwig, R. Adolescents display distinctive tolerance to ambiguity and to uncertainty during risky decision making. *Sci. Rep.* **7**, 40962. https://doi.org/10.1038/srep40962 (2017).
22. Figner, B., Mackinlay, R. J., Wilkening, F. & Weber, E. U. Affective and deliberative processes in risky choice: Age differences in risk taking in the Columbia Card Task. *J. Exp. Psychol. Learn. Mem. Cogn.* **35**, 709–730. https://doi.org/10.1037/a0014983 (2009).
23. Mitchell, S. H., Schoel, C. & Stevens, A. A. Mechanisms underlying heightened risk taking in adolescents as compared with adults. *Psychon. Bull. Rev.* **15**, 272–277. https://doi.org/10.3758/pbr.15.2.272 (2008).
24. Van Leijenhorst, L., Westenberg, P. M. & Crone, E. A. A developmental study of risky decisions on the cake gambling task: Age and gender analyses of probability estimation and reward evaluation. *Dev. Neuropsychol.* **33**, 179–196. https://doi.org/10.1080/87565640701884287 (2008).
25. Burnett, S., Bault, N., Coricelli, G. & Blakemore, S. J. Adolescents' heightened risk-seeking in a probabilistic gambling task. *Cogn. Dev.* **25**, 183–196. https://doi.org/10.1016/j.cogdev.2009.11.003 (2010).
26. Eshel, N., Nelson, E. E., Blair, R. J., Pine, D. S. & Ernst, M. Neural substrates of choice selection in adults and adolescents: Development of the ventrolateral prefrontal and anterior cingulate cortices. *Neuropsychologia* **45**, 1270–1279. https://doi.org/10.1016/j.neuropsychologia.2006.10.004 (2007).
27. Jansen, B. R., van Duijvenvoorde, A. C. & Huizenga, H. M. Development of decision making: Sequential versus integrative rules. *J. Exp. Child Psychol.* **111**, 87–100. https://doi.org/10.1016/j.jecp.2011.07.006 (2012).
28. Dror, I. E., Katona, M. & Mungur, K. Age differences in decision making: To take a risk or not?. *Gerontology* **44**, 67–71. https://doi.org/10.1159/000021986 (1998).
29. Scheres, A. *et al.* Temporal and probabilistic discounting of rewards in children and adolescents: Effects of age and ADHD symptoms. *Neuropsychologia* **44**, 2092–2103. https://doi.org/10.1016/j.neuropsychologia.2005.10.012 (2006).
30. Levin, I. P., Weller, J. A., Pederson, A. A. & Harshman, L. A. Age-related differences in adaptive decision making: Sensitivity to expected value in risky choice. *Judgm Decis. Mak* **2**, 225–233 (2007).
31. Lejuez, C. W. *et al.* Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *J. Exp. Psychol.-Appl.* **8**, 75–84. https://doi.org/10.1037//1076-898x.8.2.75 (2002).
32. Osmont, A. *et al.* How does explicit versus implicit risk information influence adolescent risk-taking engagement?. *J. Behav. Decis. Mak.* **30**, 1093–1103 (2017).
33. Humphreys, K. L. *et al.* Risky decision making from childhood through adulthood: Contributions of learning and sensitivity to negative feedback. *Emotion* **16**, 101–109. https://doi.org/10.1037/emo0000116 (2016).
34. Bechara, A., Damasio, A. R., Damasio, H. & Anderson, S. W. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* **50**, 7–15. https://doi.org/10.1016/0010-0277(94)90018-3 (1994).
35. Crone, E. A. & van der Molen, M. W. Developmental changes in real life decision making: Performance on a gambling task previously shown to depend on the ventromedial prefrontal cortex. *Dev. Neuropsychol.* **25**, 251–279. https://doi.org/10.1207/s15326942dn2503_2 (2004).
36. Van Duijvenvoorde, A. C., Jansen, B. R., Bredman, J. C. & Huizenga, H. M. Age-related changes in decision making: Comparing informed and noninformed situations. *Dev. Psychol.* **48**, 192–203. https://doi.org/10.1037/a0025601 (2012).
37. Overman, W. H. *et al.* Performance on the IOWA card task by adolescents and adults. *Neuropsychologia* **42**, 1838–1851. https://doi.org/10.1016/j.neuropsychologia.2004.03.014 (2004).
38. Christakou, A. *et al.* Neural and psychological maturation of decision-making in adolescence and young adulthood. *J. Cogn. Neurosci.* **25**, 1807–1823. https://doi.org/10.1162/jocn_a_00447 (2013).
39. Schonberg, T., Fox, C. R. & Poldrack, R. A. Mind the gap: Bridging economic and naturalistic risk-taking with cognitive neuroscience. *Trends Cogn. Sci.* **15**, 11–19. https://doi.org/10.1016/j.tics.2010.10.002 (2011).
40. Geier, C. F., Garver, K., Terwilliger, R. & Luna, B. Development of working memory maintenance. *J. Neurophysiol.* **101**, 84–99 (2009).
41. Huizinga, M., Dolan, C. V. & van der Molen, M. W. Age-related change in executive function: Developmental trends and a latent variable analysis. *Neuropsychologia* **44**, 2017–2036 (2006).
42. Luna, B. Developmental changes in cognitive control through adolescence. *Adv. Child Dev. Behav.* **37**, 233–278. https://doi.org/10.1016/s0065-2407(09)03706-9 (2009).
43. McAuley, T. & White, D. A. A latent variables examination of processing speed, response inhibition, and working memory during typical development. *J. Exp. Child Psychol.* **108**, 453–468. https://doi.org/10.1016/j.jecp.2010.08.009 (2011).
44. Kwon, H., Reiss, A. L. & Menon, V. Neural basis of protracted developmental changes in visuo-spatial working memory. *Proc. Natl. Acad. Sci. U S A* **99**, 13336–13341. https://doi.org/10.1073/pnas.162486399 (2002).
45. Niv, Y., Edlund, J. A., Dayan, P. & O'Doherty, J. P. Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *J. Neurosci.* **32**, 551–562. https://doi.org/10.1523/JNEUROSCI.5498-10.2012 (2012).
46. Bernoulli, D. Exposition of a new theory on the measurement of risk. *Econometrica* **22**, 23–36 (1954).
47. Kakade, S. & Dayan, P. Dopamine: Generalization and bonuses. *Neural Netw.* **15**, 549–559. https://doi.org/10.1016/s0893-6080(02)00048-5 (2002).
48. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B. & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879. https://doi.org/10.1038/nature04766 (2006).
49. Wilson, R. C. & Ten Collins, A. G. simple rules for the computational modeling of behavioral data. *Elife* https://doi.org/10.7554/eLife.49547 (2019).
50. Rosenbaum, G. M., Grassie, H. L. & Hartley, C. A. Valence biases in reinforcement learning shift across adolescence and modulate subsequent memory. *Elife* https://doi.org/10.7554/eLife.64620 (2022).
51. Rodriguez, B. J. M., Heekeren, H. R. & van den Bos, W. The computational basis of following advice in adolescents. *J. Exp. Child Psychol.* **180**, 39–54. https://doi.org/10.1016/j.jecp.2018.11.019 (2019).

52. Decker, J. H., Lourenco, F. S., Doll, B. B. & Hartley, C. A. Experiential reward learning outweighs instruction prior to adulthood. *Cogn. Affect. Behav. Neurosci.* **15**, 310–320. https://doi.org/10.3758/s13415-014-0332-5 (2015).
53. Master, S. L. *et al.* Distentangling the systems contributing to changes in learning during adolescence. *Dev. Cogn. Neurosci.* **41**, 100732. https://doi.org/10.1016/j.dcn.2019.100732 (2020).
54. Jepma, M., Schaaf, J. V., Visser, I. & Huizenga, H. M. Effects of advice on experienced-based learning in adolescents and adults. *J. Exp. Child Psychol.* **211**, 105230. https://doi.org/10.1016/j.jecp.2021.105230 (2021).
55. Gardner, M. & Steinberg, L. Peer influence on risk taking, risk preference, and risky decision making in adolescence and adulthood: An experimental study. *Dev. Psychol.* **41**, 625–635. https://doi.org/10.1037/0012-1649.41.4.625 (2005).
56. Chein, J., Albert, D., O'Brien, L., Uckert, K. & Steinberg, L. Peers increase adolescent risk taking by enhancing activity in the brain's reward circuitry. *Dev. Sci.* **14**, F1-10. https://doi.org/10.1111/j.1467-7687.2010.01035.x (2011).
57. Rhodes, N., Pivik, K. & Sutton, M. Risky driving among young male drivers: The effects of mood and passengers. *Transp. Res. Part F* **28**, 65–76 (2015).
58. Tversky, A. & Kahneman, D. The framing of decisions and the psychology of choice. *Science* **211**, 453–458. https://doi.org/10.1126/science.7455683 (1981).
59. Kahneman, D. & Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica* **47**, 263–292 (1979).
60. Button, K. S. *et al.* Power failure: Why small sample size undermines the reliability of neuroscience. *Nat. Rev. Neurosci.* **14**, 365–376. https://doi.org/10.1038/nrn3475 (2013).
61. Katahira, K. How hierarchical models improve point estimates of model parameters at the individual level. *J. Math. Psychol.* **73**, 37–58 (2016).
62. Shiffrin, R. M., Lee, M. D., Kim, W. & Wagenmakers, E. J. A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cogn. Sci.* **32**, 1248–1284. https://doi.org/10.1080/03640210802414826 (2008).
63. Depaoli, S. & van de Schoot, R. Improving transparency and replication in Bayesian statistics: The WAMBS-Checklist. *Psychol. Methods* **22**, 240–261. https://doi.org/10.1037/met0000065 (2017).
64. Bartlema, A., Lee, M., Wetzels, R. & Vanpaemel, W. A Bayesian hierarchical mixture approach to individual differences: Case studies in selective attention and representation in category learning. *J. Math. Psychol.* **59**, 132–150 (2014).
65. Bates, D., Maechler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).
66. Rescorla, R. A. & Wagner, A. R. in *Classical conditioning II: Current research and theory* (eds Abraham H. Black & William F. Prokasy) 64–99 (Appleton-Century-Crofts, 1972).
67. Stankevicius, A., Huys, Q. J., Kalra, A. & Series, P. Optimism as a prior belief about the probability of future reward. *PLoS Comput Biol* **10**, e1003605. https://doi.org/10.1371/journal.pcbi.1003605 (2014).
68. de Boer, L. *et al.* Attenuation of dopamine-modulated prefrontal value signals underlies probabilistic reward learning deficits in old age. *Elife* https://doi.org/10.7554/eLife.26424 (2017).
69. Tzovara, A., Korn, C. W. & Bach, D. R. Human Pavlovian fear conditioning conforms to probabilistic learning. *PLoS Comput. Biol.* **14**, e1006243. https://doi.org/10.1371/journal.pcbi.1006243 (2018).
70. Wise, T., Michely, J., Dayan, P. & Dolan, R. J. A computational account of threat-related attentional bias. *PLoS Comput. Biol.* **15**, e1007341. https://doi.org/10.1371/journal.pcbi.1007341 (2019).
71. Lamba, A., Frank, M. J. & FeldmanHall, O. Anxiety impedes adaptive social learning under uncertainty. *Psychol. Sci.* **31**, 592–603. https://doi.org/10.1177/0956797620910993 (2020).
72. Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & Van Der Linde, A. Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **64**, 583–639 (2002).
73. Steingroever, H., Jepma, M., Lee, M. D., Jansen, B. R. J. & Huizenga, H. M. Detecting strategies in developmental psychology. *Comput. Brain Behav.* **2**, 128–140 (2019).
74. Lee, M. D. Bayesian methods for analyzing true-and-error models. *Judgm. Decis. Mak.* **13**, 622–635 (2018).
75. Lodewyckx, T. *et al.* A tutorial on Bayes factor estimation with the product space method. *J. Math. Psychol.* **55**, 10058 (2011).

Acknowledgements

We thank Lisa Beense, Annie Johansson, Julia Koudijs, Mark Melaard, Quintin Roman, Henrike van den Berg, and Tess van Dorp for help with data collection. This work was funded by the Dutch National Science Foundation, NWO (VICI 453-12-005).

Author contributions

All authors conceptualized the research; MJ collected the data, analyzed the data and wrote the manuscript; JVS revised the manuscript, supervised by HMH and IV; all authors reviewed the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-022-10100-7.

Correspondence and requests for materials should be addressed to M.J.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2022