



UvA-DARE (Digital Academic Repository)

The computational basis of following advice in adolescents

Rodriguez Buritica, J.M.; Heekeren, H.R.; van den Bos, W.

DOI

[10.1016/j.jecp.2018.11.019](https://doi.org/10.1016/j.jecp.2018.11.019)

Publication date

2019

Document Version

Final published version

Published in

Journal of Experimental Child Psychology

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Rodriguez Buritica, J. M., Heekeren, H. R., & van den Bos, W. (2019). The computational basis of following advice in adolescents. *Journal of Experimental Child Psychology*, 180, 39-54. <https://doi.org/10.1016/j.jecp.2018.11.019>

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.



ELSEVIER

Contents lists available at ScienceDirect

Journal of Experimental Child Psychology

journal homepage: www.elsevier.com/locate/jecp



The computational basis of following advice in adolescents



Julia M. Rodriguez Buritica^{a,b,*}, Hauke R. Heekeren^a, Wouter van den Bos^{c,d}

^a Department of Education and Psychology, Freie Universität Berlin, 14195 Berlin, Germany

^b Center of Lifespan Psychology, Max Planck Institute for Human Development, 14195 Berlin, Germany

^c Center for Adaptive Rationality, Max Planck Institute for Human Development, 14195 Berlin, Germany

^d Department of Developmental Psychology, University of Amsterdam, 1012 WX Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Received 3 April 2018

Revised 13 November 2018

Available online 2 January 2019

Keywords:

Development

Advice taking

Social reinforcement learning

Learning from experience

Exploration

Adolescence

ABSTRACT

Advice taking helps one to quickly acquire knowledge and make decisions. This age-comparative study (in children [8- to 10-year-olds], adolescents [13- to 15-year-olds], and adults [18- to 22-year-olds]) investigated developmental differences in how advice, experience, and exploration influence learning. The results showed that adolescents were initially easily swayed to follow peer advice but also switched more rapidly to exploring alternatives like children. Whereas adults stayed with the advice over the task, adolescents put more weight on their own experience compared with adults. A social learning model showed that although social influence most strongly affects adolescents' initial expectations (i.e., their priors), adolescents showed higher exploration and discovered the other good option in the current task. Thus, our model resolved the apparently conflicting findings of adolescents being more *and* less sensitive to peer influence and provides novel insights into the dynamic interaction between social and individual learning.

© 2018 Elsevier Inc. All rights reserved.

Introduction

Advice taking can foster learning and decision making, particularly if decisions are complex and the given information is incomplete (Biele, Rieskamp, Krugel, & Heekeren, 2011). For instance, when you

* Corresponding author at: Department of Education and Psychology, Freie Universität Berlin, 14195 Berlin, Germany.

E-mail address: buritica@zedat.fu-berlin.de (J.M. Rodriguez Buritica).

arrive in a new city and must pick a restaurant for dinner, to make a more informed decision, you ask others who have previously visited this city for advice. After following the initial advice, you might also start gathering firsthand experience. For instance, after visiting the recommended restaurant a few times, you might learn that you do not like it that much and start exploring other options.

Taking advice is a common form of social learning¹ and influences one's own subsequent learning experiences (Biele et al., 2011; Biele, Rieskamp, & Gonzalez, 2009; Goodyear et al., 2016). However, the degree to which social information and one's own experience are used in learning varies significantly across development (Rodriguez, Heekeren, Li, & Eppinger, 2018). Adolescents in particular show a specific susceptibility to social influence, particularly to that of their own peers (Albert, Chein, & Steinberg, 2013; Rodman, Powers, & Somerville, 2017; Sebastian, Viding, Williams, & Blakemore, 2010). In this study, we aimed to better understand how this adolescent-specific susceptibility to peer influence shapes adolescents' subsequent learning behavior and decision making.

Developmental differences in susceptibility to social influence

Recent experimental studies have started to outline developmental differences in how susceptibility to social influence shapes behavior (Lourenco et al., 2015; Monahan, Steinberg, & Cauffman, 2009; Steinberg & Monahan, 2007; Steinberg, 2008; Sumter, Bokhorst, Steinberg, & Westenberg, 2009). Consistent with conventional wisdom, these studies show that social information becomes increasingly influential on behavior during adolescence, particularly when it comes from peers (Albert et al., 2013; Blakemore & Mills, 2014; Jones et al., 2014). More specifically, the tendency to rely on social information is particularly prominent during early adolescence (lasting until approximately 15 years of age; see also Blakemore & Mills, 2014; Gardner & Steinberg, 2005; Gunther Moor, van Leijenhorst, Rombouts, Crone, & Van der Molen, 2010; Sebastian et al., 2010). Many of these experimental studies compared behavior in the presence or absence of a passively observing peer (Casco et al., 2015; Chein, Albert, O'Brien, Uckert, & Steinberg, 2011; Gardner & Steinberg, 2005; Powers et al., 2018; Somerville et al., 2018). However, adolescents often make important decisions in the absence of peers, which raises an important question—how, and for how long, does social information influence learning and decision making when peers are no longer present?

Developmental differences in instruction- and experience-based learning

To date, little is known about *how* or for *how long* social information from peers affects adolescents' learning and decision making. To our knowledge, only one study partly addressed this question. In this study, Decker et al. showed that in a learning task adolescents relied less on intentionally false instructions compared with adults (Decker, Lourenco, Doll, & Hartley, 2015). Instead, adolescents relied more on their own experience than on the feedback presented during the learning episode of the task. The authors suggested that instruction biases learning through the top-down influence of the prefrontal cortex and that, due to decreased striatal–prefrontal connectivity (Imperati et al., 2011; Liston et al., 2006), adolescents are less influenced by instructions.² However, developmental differences during instruction-based learning could also have led to behavioral differences in this study. For instance, developmental differences in explorative behavior during learning can contribute to a reduced instruction bias in children and adolescents. Previous research has shown that children and adolescents explore more options than adults (Christakou et al., 2013). Thus, adolescents might have learned earlier that the instructions were false and that better options were available, which raises another important issue. If

¹ Previous studies have shown that advice and social information via observation, written advice, or spoken advice can influence participants' choices and that their influence does not necessarily depend on direct personal interaction (Biele et al., 2009, 2011; McElreath et al., 2008). Moreover, use of good advice (i.e., if the advice is reliable) does not even depend on whether the advice is received from a machine or a human agent (see Goodyear et al., 2016).

² In contrast to task instructions that contain *necessary* information that participants *should* use for doing an experimental task, advice (from another human) contains social information that is more *optional* information that participants *could* use for doing a specific task. Bonaccio and Dalal (2006) defined *advice* as "recommendation, from the advisor, favoring a particular option" (p. 128). Furthermore, the authors noted that "the person who receives the advice ... must decide what to do with it" and is "responsible for making the final decision" (p. 128).

we want to understand how social information affects adolescents in the absence of peers, we must understand how it interacts with their own experience, particularly given the well-established findings on developmental differences in learning from experience (Crone, Richard, & Van der Molen, 2004; Crone, Somsen, Zanolie, & Van der Molen, 2006; Eppinger, Mock, & Kray, 2009; Ferdinand & Kray, 2014; Hämmerer, Li, Müller, & Lindenberger, 2010; van den Bos, Cohen, Kahnt, & Crone, 2012; van Duijvenvoorde, Zanolie, Rombouts, Raijmakers, & Crone, 2008). One consistent result from these studies is that performance increases across adolescence. In addition, these studies show that children have greater difficulty in using negative feedback for learning (Crone et al., 2004; Eppinger et al., 2009; Hämmerer et al., 2010; van den Bos et al., 2012; van Duijvenvoorde et al., 2008). Thus, to understand the unique and lasting effect of social information on adolescent behavior, we also must consider developmental differences in experience-based learning.

The current study

To resolve these outstanding issues, we used a specifically designed social reinforcement learning (RL) task in combination with computational modeling. The task involves both social information- and experience-based learning (modified Iowa Gambling Task after Biele et al., 2011). Computational models have the advantage that they can be used to test different theories about the latent processes that underlie social influence and learning (van den Bos, Bruckner, Nassar, Mata, & Eppinger, 2018).

In the current task, participants (children, adolescents and adults) needed to choose one of four card decks that were associated with gains and losses (see Fig. 1A). Unbeknownst to the participants, two of the four decks were associated with higher expected positive values (“good decks”) than the other two




Fig. 1. Experimental design. (A) Participants received advice before they were asked to play a four-armed bandit task. Each trial started with the presentation of four card decks, one of which needed to be selected within a maximum of 4 s. Thereafter, the associated feedback was presented. Before a new trial started, a fixation cross was displayed for 1 s. (B) Payoff distribution (i.e., different payoff values on the x axis displayed per frequency on the y axis) and average payoff for the bad decks (in red) and the good decks (in blue). Error bars represent standard errors of the mean. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

(“bad decks”) (see Fig. 1B). At the beginning of the experiment, participants received good advice (i.e., for one of the good decks) from a same-aged peer. After the advice was given, the participants were free to draw from each of the decks as often as they liked. Both good decks were equally good (equal payoff distributions); thus, if the participants explored the other decks, they could over time learn that there was another equally good deck. Furthermore, if the initial advice was completely ignored, participants were expected to draw equally often from each of the good decks. However, Biele et al. (2011) showed that, using the exact same task, adults keep preferring the advised good deck. Thus, the task was designed such that the extent to which the other good deck is selected captures the influence of the advice. In that sense, the other good (nonrecommended) deck served as a “nonsocial” condition. Finally, we used computational modeling to separate the effects of advice, experience (positive and negative), and exploration on learning behavior (for more details on models, see Method). As noted above, each of these aspects could contribute to developmental differences in following advice.

Based on previous studies, we expected the susceptibility to peer influence to peak in adolescents (Steinberg & Monahan, 2007) compared with children and adults. Furthermore, we expected that influence would be particularly strong at the beginning of the task before learning by experience took over (Decker et al., 2015). In addition, we expected that throughout learning, children and adolescents would show more exploration than adults (Christakou et al., 2013; Decker et al., 2015). Thus, the two younger groups should be faster in discovering that there is another (equally) good deck in the experiment. Finally, children were expected to show less optimal performance overall due to their difficulty in using negative feedback for learning (van Duijvenvoorde et al., 2008). We have developed several computational RL models to capture the interaction among these different elements (social influence, exploration, and experience; see Method for detailed model predictions).

Method

Participants

The effective sample of the study consisted of 25 adults aged 18–22 years (13 female; mean age = 20.32 years, $SD = 1.15$), 24 adolescents aged 13–15 years (12 female; mean age = 13.71 years, $SD = 0.75$), and 24 children aged 8–10 years (10 female; mean age = 9.08 years, $SD = 0.83$) (see online supplementary material for a justification of age group selection). Our sample size estimation was based on a study by van den Bos et al. (2012) that used a feedback-based learning task, similar RL models to analyze the data (except that there was no advice), and comparable age ranges (see supplementary material for further details). In our analyses, we treated the groups as homogeneous developmental stages and used the contrast or group values shown in the linear $(-1, 0, 1)$, quadratic $(-1, 2, -1)$, and emergent $(-2, 1, 1)$ regression models. (Here, we follow standard practice in developmental studies; see, e.g., van den Bos, van Dijk, Westenberg, Rombouts, & Crone, 2011.)

The data of 1 adult were excluded due to technical problems during data collection, and the data of 1 adolescent were excluded from further analyses due to many missing responses (2 standard deviations from the group mean). All participants had normal or corrected-to-normal vision, had no neurological or psychological disorders, and were native German speakers. Participants received compensation of 15 Euros. Prior to the experiment, we obtained informed consent from the participants and their parents (in the case of children and adolescents). The study was approved by the ethics committee of the Max Planck Institute for Human Development.

Participants took part in one behavioral session where we assessed psychometric covariate measures and performance in the advice task. Participants in each age group were tested with a small group of strangers in the same room to make the social aspect of the task more salient. To control for potential performance effects of viewing one another (Chein et al., 2011; Weigard, Chein, Albert, Smith, & Steinberg, 2014), however, each participant was tested in a separate booth; participants could neither see nor interact with the other participants in the study.

To ensure the age representativeness of our sample, general cognitive abilities of the sample were described using (a) the numbers task (Gold, Carpenter, Randolph, Goldberg, & Weinberger, 1997) as an indicator of working memory (WM) and (b) a short version of the CFT (Culture Fair Test; Weiß, 2006)

as an indicator of fluid intelligence. The numbers task involved an auditory presentation of a mixed series of alternating numbers that needed to be repeated in the correct or reversed order. The series starts with a sequence of two numbers and on each step increases with one additional number. The short version of the CFT involved completing a matrix of patterns in which one item is missing. As with the Raven's Progressive Matrices (Raven & Court, 1998), the difficulty of the patterns increases. Children had lower WM scores compared with adolescents and adults (children: $M = 7.52$, $SD = 1.71$; adolescents: $M = 10.23$, $SD = 1.71$; adults: $M = 10.66$, $SD = 2.18$), $F(2, 70) = 19.77$, $p < .001$, $\eta_p^2 = .36$, whereas adolescents and adults did not differ in their WM scores, $t(48) = 0.87$, $p = .39$. Fluid intelligence scores increased with age (children: $M = 9.56$, $SD = 2.68$; adolescents: $M = 12.29$, $SD = 2.25$; adults: $M = 13.48$, $SD = 1.48$), $F(2, 70) = 20.57$, $p < .001$, $\eta_p^2 = .37$. These age differences are consistent with findings from larger population-based lifespan samples (Li et al., 2004).

Experimental design

In the learning task (Biele et al., 2011; Peirce, 2007), participants were asked to choose one of four card decks (see Fig. 1A) (according to Biele et al., 2011). The goal of the task was to maximize cumulative rewards. Participants could select one card deck within a maximal response time window of 4 s and received feedback (displayed for 1 s) thereafter. After a short fixation cross was displayed for 1 s, a new trial of a total of 210 trials started.

Unbeknownst to the participants, the four decks consisted of two more beneficial "good" decks and two less beneficial "bad" decks (see Fig. 1B). Participants received one recommendation for one of the good decks (counterbalanced across participants) at the beginning of the task. Thus, a preference for the recommended deck over the other good deck would be a clear indicator of advice taking. Participants were told that another peer (who participated in a previous session) gave a recommendation after they played the task. Unbeknownst to the participants, the advice was controlled by the experimenter and was always a good recommendation. Participants were verbally (and in writing) debriefed about the fact that the social information was controlled by the experimenter. Although participants were not explicitly asked, no participants indicated that they did not believe the information was from a previous participant.

Furthermore, to investigate the effect of advice following positive and negative feedback, each of the four card decks (the two good ones and the two bad ones) was associated with 50% losses and 50% gains (see frequency distribution of payoffs in Fig. 1B). The bad decks were associated with higher losses on average, but also with slightly higher gains, than the good decks (see payoff distribution in Fig. 1B). The resulting average payoff for the good decks was 10 points, whereas the payoff for the bad decks averaged 2.5 points (see Fig. 1B). Thus, selecting either of the two good decks would result in equal payoffs.

Data analysis

For the statistical analyses, we excluded trials that exceeded the response deadline (4 s) from further analyses (children: $M = 7.42$ trials, $SD = 8.05$; adolescents: $M = 6.83$ trials, $SD = 7.56$; adults: $M = 4.24$ trials, $SD = 5.09$). Importantly, the three age groups did not differ from one another in terms of missing responses, $F(2, 70) = 1.439$, $p = .24$, $\eta_p^2 = .04$.

The preference for the recommended good deck (coded as 1) over the other good deck (coded as 0) was analyzed with a mixed-effects logistic regression using the lme4 package in R (Bates, Mächler, Bolker, & Walker, 2015) with independent predictors of age group (adults, adolescents, or children), trial (1–210, z-transformed), and their interactions. We included per-participant and trial adjustments to the intercept (random intercepts).³ To compare the outcomes of the different learning strategies, we calculated expected payoff by multiplying the percentage of choices per deck times the average payoff of that deck for 210 trials. To correct for multiple comparisons using multiple t tests, a Bonferroni correction was applied.

³ This logistic regression model fit the data best (BIC = 10,179) compared with (a) including only per-trial adjustments to the intercept (BIC = 10,283) or (b) not including random intercepts (BIC = 11,347). All models revealed the same fixed effects and interactions as reported in the article.

General cognitive abilities were analyzed using univariate analysis of variance (ANOVA) and independent-samples t tests. Effect sizes (η_p^2) are reported. Age differences in parameter estimates of the best fitting model (i.e., “prior + bonus dual RL” model) were compared using multiple regression analyses with the different age trends as a predictor (van den Bos, Rodriguez, Schweitzer, & McClure, 2015).

Modeling social influence

RL models have previously been used to capture the interaction among the different mechanisms (advice, experience, and exploration) during social learning in adults (Behrens, Hunt, & Rushworth, 2009; Biele et al., 2009). To further separate the effects of advice, experience (positive and negative), and exploration on behavior in the four-armed bandit task, we fit nine different RL models to the behavioral data.

RL model

The basic RL model uses the prediction error to update the beliefs associated with each choice option (e.g., Deck A, B, C, or D). The prediction error (δ_t) compares the current outcome (r_t) with the expected outcome (w_t):

$$\delta_t = r_t - w(i)_t$$

where $w(i)_t$ is the expectation associated with the chosen option i . Whenever feedback is better (worse) than expected, the model will generate a positive (negative) prediction error, which is used to increase (decrease) the predicted value for the chosen option $w(i)_{t+1}$. The effect of the prediction error on forming new beliefs is scaled by the learning rate ($0 < \alpha < 1$):

$$w(i)_{t+1} = w(i)_t + \alpha \cdot \delta_t$$

A high learning rate indicates that new information has a much stronger effect on future behavior than less recent information (i.e., effect on advice).

To model trial-by-trial choices, we used the soft-max choice rule to compute the probability (P) of choosing one of the decks (i.e., i) based on one's own predictions about the outcomes of all decks j on trial t (Montague, Hyman, & Cohen, 2004):

$$P(i)_t = \frac{e^{\theta \cdot w(i)_t}}{\sum_{j=1}^N e^{\theta \cdot w(j)_t}}$$

where θ is a free parameter that indicates the sensitivity of the participant to the differences in decision weights. The lower the θ parameter, the more exploratory the choices appear. Therefore, it was our expectation that children and adolescents would show lower estimated values of the θ parameter. To test various models of social influence on these basic learning processes, we tested and compared several learning models.

Bonus model

To investigate the influence of advice on learning, we compared models that implemented different mechanisms in which advice influences behavior. First, previous studies in adults have suggested that the recommended option is associated with a constant bonus (2011; Biele et al., 2009). Due to this bonus, the recommended option remains attractive throughout the experiment. Thus, this model differs from the basic RL model by assuming that there is a constant bonus associated with choosing the recommended option. Accordingly, the reward function was modified such that

$$r_t = r_t + (\lambda(i) \cdot \beta \cdot \mu)$$

where $\lambda(i)$ is an indicator function that takes the value 1 if option i is recommended and the value 0 if option i is not recommended. The β parameter captures the extent to which social influence (i.e., across all trials) leads to an outcome bonus ($0 < \beta < \infty$), and μ is the expected payoff from choosing randomly among all options (i.e., 6.25 on average).

Bonus + decay model

Next, we extended the simple bonus models with the assumption that the outcome bonus associated with the recommended option would decline over time, for instance, due to forgetting (bonus + decay RL). To model this waning influence of advice over time, we modulated the bonus parameter as follows:

$$\beta_t = (\beta \cdot \mu) \cdot \left(\frac{1}{t}\right)^\pi$$

where π is the free parameter that captures how quickly the effect of influence decays, with smaller values indicating less decline ($0 < \pi < \infty$). If forgetting played a role in social influence, we would expect this effect to be the strongest for the youngest age groups.

Prior model

The prior RL assumes that the recommendation sets an initial strong positive expectation (prior) for the recommended deck but has no further influence on choices. The initial reward expectation in the prior model is defined as

$$w(\text{recommended})_0 = \rho \cdot \mu$$

where ρ captures the social influence on the prior expectations before sampling (on $t = 0$) and μ is the expected payoff from choosing randomly among all options (i.e., 6.25).

Given that we expected that adolescents were specifically most sensitive to the peer advice at the beginning of the task, we predicted that they would have the highest estimated prior (ρ) but that all age groups would show a similar effect of the advice bonus (β).

Dual learning rate models

Based on previous developmental studies, we expected that gains and losses would be asymmetrically updated (Kahnt et al., 2009; van den Bos et al., 2012); thus, we extended all of the social influence models with two independent learning rates for positive feedback (α_{pos}) and negative feedback (α_{neg}). Here, we particularly expected to see higher negative learning rates for children compared with the older age groups (van den Bos et al., 2012).

Mixed models

Finally, given that the bonus and prior mechanisms are not mutually exclusive, we tested mixed prior + bonus models that assume people have an increased prior ρ for the recommended deck and the recommendation is also associated with a constant bonus β . In total, this process resulted in a model space of nine models that can capture different interactions among exploration (θ), experience (αs), and social influence (β and ρ).

Note that for all models except the models with a prior, the decision weight (w) for each option was always set to 0 at the beginning of the experiment. This approach reflects the assumption that the participants had the same expectation of rewards for each option at the beginning. In addition, following Biele et al. (2009, 2011), all models assumed that participants who received advice would always choose the recommended option in their first trial. This assumption was implemented by setting the probability of choosing the recommended deck at 100% for the first trial.

Model fitting procedure

The free parameters of the models were individually estimated by fitting the model predictions to the participants' decisions (see [supplementary material](#) for further details on the model fitting procedure). For model selection purposes, we computed the Bayesian information criterion (BIC) across all participants for the different models, where lower BIC values indicate better fit.⁴

⁴ We followed the standard practice of computing the total amount of evidence for each model and comparing it (Farrell & Lewandowsky, 2018). In other words, all of the models fit on the individual level, but we transformed the summed log likelihood, or model fit, of all participants together and then computed the model BIC.

Results

Choice behavior

In a first step, we checked to what extent each deck was selected above chance level (i.e., 25%) averaged across all trials (i.e., 210) using one-tailed t tests. As seen in Fig. 2A, the adults and adolescents selected the recommended deck above chance ($t_s > 2.60$, $p_s \leq .01$), but the children did not, $t(23) = 1.53$, $p = .14$. The adolescents chose bad decks below chance ($t_s > -2.73$, $p_s \leq .01$), whereas the adults and children did not ($t_s > -2.00$, $p_s \geq .006$) (see Fig. 2A). As expected, most individuals selected the advised deck on the first trial (children: 85%; adolescents: 100%; adults: 86%).

In a second step, we compared how likely the recommended deck was to be selected over the other good deck per trial using mixed-effects logistic regression. As seen in Fig. 2B, the recommended deck was less likely to be selected over the other good deck by the children and adolescents compared with the adults ($b = -0.55$) (see Table 1). With more sampling time, the recommended deck was less likely to be chosen over the other good deck by the adolescents; that is, with each trial, the recommended deck, rather than the other good deck, was selected less often by $b = -0.16$ (see Table 1 and Fig. 2B). Finally, a post hoc test revealed that both adults and adolescents profited similarly equally by selecting the good decks more frequently, showing no significant difference in expected payoff ($p = .70$) (see Fig. 2C). Thus, only adolescents, but not adults, had a higher expected payoff compared with children ($p = .06$ and $p = .20$, Bonferroni corrected, respectively).

Modeling social influence

The model comparisons yielded several important findings. First, as expected, the model comparisons showed that all social models were better than the baseline RL model (see Fig. 3A). Second,




Fig. 2. (A) Percentage chosen per deck and age group averaged over all 210 trials. Error bars represent standard errors of the mean. (B) Percentage chosen of the recommended deck (in red) and the other good deck (in blue), presented by bin (of 5 trials each) separately for each age group. Shaded areas represent standard errors of the mean. (C) Total payoff per age group. (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

Table 1
Summary of mixed-effects logistic regression.

	B (SE)	95% CI for odds ratio		
		Lower	Odds ratio	Upper
Intercept	0.80 ^{***} (0.19)	1.54	2.23	3.25
Age.f2	-0.45 (0.27)	0.38	0.64	1.07
Age.f3	-0.55 ⁺ (0.27)	0.34	0.58	0.98
Trial	0.06 (0.05)	0.86	0.94	1.03
Age.f2 * Trial	-0.16 ^{**} (0.06)	0.76	0.86	0.96
Age.f3 * Trial	0.05 (0.06)	0.93	1.05	1.19

Note. Here, the reference group is adults. Age.f2 compares adolescents with adults, and Age.f3 compares children with adults. $R^2 = .143$ (Tjur's D). CI, confidence interval.

⁺ $p < .05$.
^{**} $p < .01$.
^{***} $p < .001$.




Fig. 3. Modeling social influence. (A) BICs for model comparison. The relative difference in BIC values for each model compared with the model with the lowest overall BIC value is shown. The Bayes factor for comparing the best model (lowest BIC) and second-best model is 6494, which indicates that the best fitting model is very strongly favored over all other models tested. Note that this model is also the winning model if we performed these comparisons separately at the level of age groups. (B) Parameter estimates for the prior + bonus dual RL model. Median estimates presented separately for the three age groups and the two learning rates (α_{gain} and α_{loss}), exploration (θ), prior (ρ), and bonus (β) are shown. Error bars represent standard errors of the mean. Recovery analyses further confirmed these parameter estimates (see [online supplementary material and Fig. S2](#)). (C) Age differences in the parameter estimates of the prior + bonus dual RL model. Age represents the linear trend, Age2 the adolescent peak, and Age3 the adolescent emerging trend. * $p < .05$; ** $p < .01$.

consistent with previous findings (van den Bos et al., 2012), the dual learning rate models were generally better at capturing behavior than the single learning rate models (see Fig. 3A). Finally, social influence was described by two separate processes: (a) an increased prior expectation (ρ) and (b) a constant bonus (β) for the recommended option. Note that the parameter values were uncorrelated ($r_{\text{Pearson}} = .01, p = .99$), suggesting that these processes can indeed be dissociated.

For further quality control, we simulated the behavior of the “prior + bonus dual RL” model (see supplementary material and Fig. S1). For each age group, we simulated 100 agents using the medians of the parameter fits for that group. The results of these simulations indicate that the model is adequately able to capture the qualitative difference in learning behavior observed in each age group (e.g., the initial peak in adolescence; see supplementary material and Fig. S1).

To further explore the validity of the RL models and the model selection procedure, we performed model recovery analyses (see supplementary material and Table S1). In these analyses, we check whether data generated by one specific model will also result in the best model fit for that same model. If that possibility is true, the model is “recoverable.” These analyses indicated that the different computational models we compared indeed made distinguishable predictions about behavior and, thus, that the model comparison is feasible.

Parameter estimates of the winning model

To better understand the processes that underlie age differences in behavior, we further investigated age differences in the parameter estimates (see Fig. 3C and supplementary material for further details) reflecting the effect of advice, experience (positive and negative), and exploration. Kolmogorov–Smirnov tests indicated that the distributions of model parameters deviated significantly from a normal distribution (all $ps < .001$). For subsequent age analyses, all parameters were initially transformed to approximate normality via the Box–Cox transformation.

Social influence

For all age groups, the bonus parameter (β) was significantly greater than zero (all $ps < .001$), indicating a consistent preference for the advised deck over the other good deck. Similarly, all age groups showed a significantly large effect of the advice on setting higher prior expectations ($\rho > 0$) (all $ps < .001$). However, only the prior, not the bonus, showed age-related differences. More specifically, the adolescents’ increase in prior expectations based on advice was significantly greater than that of both the children and adults ($\beta_{\text{quadratic}} = .14, t = 2.30, p = .025$). This finding suggests that social influence most strongly affects adolescents’ initial expectations compared with children and adults (i.e., quadratic age pattern in priors) (see Fig. 3B). Thus, our findings support the view that adolescence can be a developmental period with a particularly high susceptibility to social influence (Blakemore & Mills, 2014; Jones et al., 2014; van Hoorn, van Dijk, Meuwese, Rieffe, & Crone, 2016).

Exploration

As expected, there was a significant linear increase in the choice sensitivity parameter (θ) ($\beta_{\text{linear}} = .05, t = 3.03, p = .003$), indicating that the older participants showed less exploratory behavior than the younger participants.

Experience

As seen in Fig. 3B, we found no significant age trend for the learning rate associated with gains (α_{win}); however, consistent with previous studies (van den Bos et al., 2012), we found that the children had a higher learning rate for negative outcomes compared with the adolescents and adults ($\alpha_{\text{loss}} : \beta_{\text{emerging}} = .28, t = 2.08, p = .041$), indicating that they were more sensitive to negative feedback than the other age groups. A high learning rate indicated that new experiences have a much stronger effect on future behavior than less recent experiences (i.e., effect of advice on experiences).

Model simulations

Next, we explored how these group-level strategies would fare in different environments, extrapolating based on our modeling results. More specifically, we simulated participant behavior, based on our parameter estimates, in an environment where there was another deck that was better than the recommended deck and in an environment where the other deck was worse. Given that the adolescents appeared to be the only participants who learned that there was another good deck in the original experiment, we expected their strategy to be particularly advantageous when simulating their behavior in an environment that included a better alternative. Indeed, as expected, the simulations indicated that the adolescents would be quicker at switching from the recommended deck to the optimal deck (see Fig. 4A). However, this behavior would not occur when the alternative was worse (see Fig. 4B). In that case, the exploratory behavior of the children and adolescents would be costly (see Fig. 4B).

Discussion

In this developmental study, we used a four-armed bandit task that included peer advice in children (8–10 years old), adolescents (13–15 years old), and young adults (18–22 years old). As expected, all age groups followed the advice at the beginning of the task; however, after a few trials, the behavior of the different age groups started to diverge. The most salient developmental differences suggest




Fig. 4. Model simulations. Simulated data using the best-fitting model and the median parameter values for each group are shown. The simulations are the result of 210 trials and 100 iterations. Shaded areas represent standard errors. (A) Simulation: Better alternative. Simulations in an environment where the other good deck was better than the recommended deck are shown. That is the expected value (EV) was slightly higher for the other good deck ($EV_{\text{advised}} = 10$ and $EV_{\text{alternative}} = 12.5$). The blue line represents choices from the recommended deck, and the red line represents the better alternative. Total payoff per age group for the better alternative is shown. Payoff is calculated by multiplying the percentage of choices per deck times the expected value (times number of trials [$n = 210$]). When there was a better alternative, the adolescents had a slight advantage over the adults. (B) Simulation: Worse alternative. Simulations in which the recommended deck was better than all other decks are shown. Again, the blue line is the recommended deck ($EV_{\text{advised}} = 10$) and the red line is the worse alternative ($EV_{\text{alternative}} = 7.5$). Total payoff per age group for the worse alternative is shown. Payoff is calculated by multiplying the percentage of choices per deck times the expected value (times number of trials [$n = 210$]). (For interpretation of the references to color in this figure legend, the reader is referred to the Web version of this article.)

that (a) adolescents are initially the most sensitive to advice, (b) adults most consistently follow advice, (c) younger participants' behavior is more exploratory, and (d) children have more difficulty in handling negative feedback during learning. Finally, these developmental differences were best described by a specific computational social influence model. This model provides more detailed insight into developmental differences underlying advice taking, learning from experience (positive and negative), and exploration. In addition, the model made the interesting prediction that adolescents' behavior can be optimal in certain environments. These novel findings are discussed in more detail below.

Developmental differences in susceptibility to social influence

Participants of all ages benefited from receiving good advice; however, consistent with previous developmental findings, the adolescents showed the highest susceptibility to peer influence (Chein et al., 2011; van Hoorn et al., 2016). More specifically, our analyses suggest that social influence most strongly affects adolescents' initial expectations compared with children and adults. That is, we found a quadratic age pattern in the prior parameter ρ (see Fig. 3B). Thus, our findings support the view that adolescence can be a developmental period with a particularly high susceptibility to social influence (Blakemore & Mills, 2014; Jones et al., 2014; van Hoorn et al., 2016). However, a recent study by Decker et al. (2015) also showed that for adolescents social influences might not have a long-lasting effect. Indeed, when we examine the behavior at the end of the task, we no longer observe adolescents choosing the advised option more often than the other good deck. Conversely, the adults showed a more consistent influence of advice over time. However, the modeling results show that for all age groups the advice had a similar long-lasting effect on choices (bonus parameter β did not show any age trends). Our modeling results suggest that further developmental differences in the effects of advice are due to differences in exploration- and experience-based learning.

Developmental differences in exploration

Consistent with previous developmental studies (Christakou et al., 2013; Decker et al., 2015), the children and adolescents showed more exploratory behavior compared with the adults (i.e., linear age pattern in sensitivity parameter θ) (see Fig. 3B). Higher exploration rates are assumed to be linked to the protracted maturation of prefrontal cognitive control functions (Decker et al., 2015; Thompson-Schill, Ramscar, & Chrysiou, 2009). Exploratory behavior in itself has many positive aspects, particularly in dynamic and unknown environments, and has been suggested to be an important adaptation in human development (Spear, 2000; Thompson-Schill et al., 2009) because this type of exploration aids in the learning of behavior relevant for adult functioning (Spear, 2000).

Indeed, in the current task, the exploratory behavior of the adolescents might have resulted in a benefit. That is, the adolescents increasingly selected the other good deck with learning and chose the bad decks below chance level (which was not true for the other two age groups) (see Fig. 2A). This observation suggests that their exploration might have led them to learn more about the expected value of each of the decks. Thus, adolescents ultimately used a strategy that led to an increased likelihood that both good decks, not only the recommended deck, were chosen more often than the bad decks (see Fig. 2B). Our simulations further supported this hypothesis (see Fig. 4A). That is, if the other good deck was associated with even higher expected values than the recommended deck, the adolescents' higher exploratory behavior would lead to higher learning rates compared with the adults' behavior and to higher total payoffs within the task (see Fig. 4A).

These findings highlight the importance of considering (a) the structure of the environment and (b) the age of the learner when making normative statements about certain types of behavior. This idea is consistent with findings showing that more goal-directed model-based learning strategies, as opposed to more reward-directed model-free learning strategies, apply differently in certain types of environments (Kool, Cushman, & Gershman, 2016; Kool, Gershman, & Cushman, 2017) and develop across adolescence (Decker, Otto, Daw, & Hartley, 2016; Potter, Bryce, & Hartley, 2017). Future studies that use different learning environments are needed to further explore the possible harms and benefits of adolescent versus adult learning strategies.

Developmental differences in experience-based learning

Finally, although the children also showed increased exploratory behavior, this change did not result in their choosing the other good deck more often (see Fig. 2). This finding might be the result of children's difficulties with using negative feedback for learning compared with adolescents and adults (Crone et al., 2004; van den Bos et al., 2012; van Duijvenvoorde et al., 2008). Our modeling results support this view by showing higher learning rates for losses in the children compared with the other two age groups (i.e., emerging age pattern in negative learning rate α_{loss}) (see Fig. 3B). Children's stronger reactivity to negative feedback has been linked to increased shifting behavior (Christakou et al., 2013), to a reduced capability to differentiate between response-dependent and uninformative negative feedback (Crone et al., 2004), to larger electrophysiological responses with respect to negative feedback (i.e., reflected in an enhanced feedback-related negativity compared with older age groups; Eppinger et al., 2009; Hämmerer et al., 2010), and higher negative learning rates in older age groups (Kahnt et al., 2009; van den Bos et al., 2009), indicating their difficulty in efficiently adapting their behavior to negative feedback signals. Thus, children might be less able to assess the information value of negative feedback and, therefore, less able to use these experiences to form accurate beliefs about future outcomes (van Duijvenvoorde et al., 2008). Our modeling results support this view by showing higher learning rates for losses (α_{loss}) in the children compared with the other two age groups (see Fig. 3B). Previous studies suggested that children's performance decreases as the probability of negative feedback increases (Eppinger et al., 2009).

In summary, although the children weighted (negative) experience more heavily than weight advice, they were not able to benefit from their experience to the same degree as the adolescents in this learning environment. In contrast, a combination of higher exploration rates and negative learning rates might have contributed to their overall lower performance. Future studies are needed to tease apart how much of children's learning deficit is driven by exploration and how much is driven by developmental differences in using negative feedback.

Limitations and future directions

Our findings provide unique insights into developmental differences underlying learning from advice and experience. However, in our experiments, social information was given once from an anonymous peer within each age group, and the recommendation was always good and controlled by the experimenter. This reliability might have affected the perception of the quality of advice. In addition, we did not include a measure about the belief of our participants in others' advice. Future studies should follow up on this potential variation by assessing more details about participants' beliefs in advice. Such an exploration might, for instance, further explain differences between how much peer manipulations affect behavior across development.

Furthermore, it would be interesting to examine whether adolescents' early sensitivity to good advice is specific to the identity of the source and to examine developmental differences in how advice taking depends on the expertise and quality of the advice. For instance, would the effect of advice be different or similar when it comes from peers versus adults or from an anonymous other (or even from a computer)? This investigation could be particularly important for school settings in which children and adolescents must learn from recommendations of teachers and peers. In school, peers also differ in popularity status, something that becomes extremely salient for the adolescents (Cillessen & Rose, 2005). Future studies would benefit from exploring how effects of advice on adolescent learning depend on the status of peers. One could speculate that the prior might be higher for high-status peers that could also influence the exploration of alternatives differently, whereas the reverse pattern might be observable for the advice of low-status peers.

Finally, we compare adolescents with prepubertal children and adults. However, adolescence is a very dynamic period that is marked by great changes in pubertal hormone levels, including an increase in the secretion of adrenal androgens and gonadal steroids. Recently, it has been suggested that these hormones might play an important role in how adolescents process social information (Blakemore, Burnett, & Dahl, 2010; Forbes & Dahl, 2010). Therefore, it might be of specific interest to study the effect of these hormones on social influence parameters.

When transferring the current findings to an applied context, one tempting interpretation of our results is that educational intervention programs should be aware of developmental differences in using instruction and experienced feedback for learning. Thus, different educational interventions might be more or less appropriate depending on the age of the learner. For instance, children show difficulty in using negative feedback for learning, and adolescents can rely too much on personal experiences versus teacher instruction. This point is specifically of interest for (guided) discovery-based learning, in which the success of the method depends on the right balance between exploratory learning and guided instruction (Mayer, 2004). It will be of great interest to pursue future studies to investigate the interaction between advice/instruction- and experience-based learning in classroom settings.

Conclusion

Taken together, our findings show that peer advice guides learning from one's own experience and that adolescents show the highest initial susceptibility to peer advice. Crucially, higher exploration rates enable adolescents to discover other opportunities. Thus, our results extend previous findings by showing that adolescents' more exploratory behavior could be—depending on the environmental structure—even more beneficial than less exploratory learning strategies. Taken together, these findings suggest a more nuanced view of developmental differences in advice taking because adolescence might be a unique period associated not only with higher peer influence on behavior but also with a healthy reliance on personal experience. These findings raise interesting questions concerning the features of the everyday environment of adolescents that afford such exploratory behavior and highlight the need to understand the structure of the environment in which development occurs.

Acknowledgments

We would like to thank Valerie Keller, Raphael Schultz and Jann Wäscher for help during data acquisition, Lukas Nagel for programming the task and our participants for their contribution to the study. This work was supported by the Freie Universität Berlin, the Jacobs Foundation, and an Open Research Area grant (ID 176), European Research Council grant (ERC-2018-StG-803338) and the Netherlands Organization for Scientific Research grant (NWO-VIDI 016.Vidi.185.068) awarded to W. VD.B. Julia Rodriguez Buritica's doctoral fellowship was supported by the Berlin School of Mind and Brain and Freie Universität Berlin. We are grateful to the Center of Lifespan Psychology at the Max Planck Institute for Human Development for making data collection possible.

Appendix A. Supplementary material

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jecp.2018.11.019>.

References

- Albert, D., Chain, J., & Steinberg, L. (2013). The teenage brain: Peer influences on adolescent decision making. *Current Directions in Psychological Science*, 22, 114–120.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1). <https://doi.org/10.18637/jss.v067.i01>.
- Behrens, T. E. J., Hunt, L. T., & Rushworth, M. F. S. (2009). The computation of social behavior. *Science*, 324, 1160–1164.
- Biele, G., Rieskamp, J., & Gonzalez, R. (2009). Computational models for the combination of advice and individual learning. *Cognitive Science*, 33, 206–242.
- Biele, G., Rieskamp, J., Krugel, L. K., & Heekeren, H. R. (2011). The neural basis of following advice. *PLoS Biology*, 9(6) e1001089.
- Blakemore, S.-J., Burnett, S., & Dahl, R. E. (2010). The role of puberty in the developing adolescent brain. *Human Brain Mapping*, 31, 926–933.
- Blakemore, S.-J., & Mills, K. L. (2014). Is adolescence a sensitive period for sociocultural processing? *Annual Review of Psychology*, 65, 187–207.
- Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, 101, 127–151.

- Cascio, C. N., Carp, J., O'Donnell, M. B., Tinney, F. J., Bingham, C. R., Shope, J. T., ... Falk, E. B. (2015). Buffering social influence: Neural correlates of response inhibition predict driving safety in the presence of a peer. *Journal of Cognitive Neuroscience*, *27*, 83–95.
- Chein, J., Albert, D., O'Brien, L., Uckert, K., & Steinberg, L. (2011). Peers increase adolescent risk taking by enhancing activity in the brain's reward circuitry. *Developmental Science*, *14*, F1–F10.
- Christakou, A., Gershman, S. J., Niv, Y., Simmons, A., Brammer, M., & Rubia, K. (2013). Neural and psychological maturation of decision-making in adolescence and young adulthood. *Journal of Cognitive Neuroscience*, *25*, 1807–1823.
- Cillessen, A. H. N., & Rose, A. J. (2005). Understanding popularity in the peer system. *Current Directions in Psychological Science*, *14*, 102–105.
- Crone, E. A., Richard, J., & Van der Molen, M. W. (2004). Developmental change in feedback processing as reflected by phasic heart rate changes. *Developmental Psychology*, *40*, 1228–1238.
- Crone, E. A., Somsen, R. J. M., Zanolie, K., & Van der Molen, M. W. (2006). A heart rate analysis of developmental change in feedback processing and rule shifting from childhood to early adulthood. *Journal of Experimental Child Psychology*, *95*, 99–116.
- Decker, J. H., Lourenco, F. S., Doll, B. B., & Hartley, C. A. (2015). Experiential reward learning outweighs instruction prior to adulthood. *Cognitive, Affective, & Behavioral Neuroscience*, *15*, 310–320.
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychological Science*, *27*, 848–858.
- Eppinger, B., Mock, B., & Kray, J. (2009). Developmental differences in learning and error processing: Evidence from ERPs. *Psychophysiology*, *46*, 1043–1053.
- Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. New York: Cambridge University Press.
- Ferdinand, N. K., & Kray, J. (2014). Developmental changes in performance monitoring: How electrophysiological data can enhance our understanding of error and feedback processing in childhood and adolescence. *Behavioural Brain Research*, *263*, 122–132.
- Forbes, E. E., & Dahl, R. E. (2010). Pubertal development and behavior: Hormonal activation of social and motivational tendencies. *Brain and Cognition*, *72*, 66–72.
- Gardner, M., & Steinberg, L. (2005). Peer influence on risk taking, risk preference, and risky decision making in adolescence and adulthood: An experimental study. *Developmental Psychology*, *41*, 625–635.
- Gold, J. M., Carpenter, C., Randolph, C., Goldberg, T. E., & Weinberger, D. R. (1997). Auditory working memory and Wisconsin Card Sorting Test performance in schizophrenia. *Archives of General Psychiatry*, *54*, 159–165.
- Goodyear, K., Parasuraman, R., Chernyak, S., Madhavan, P., Deshpande, G., & Krueger, F. (2016). Advice taking from humans and machines: An fMRI and effective connectivity study. *Frontiers in Human Neuroscience*, *10*. <https://doi.org/10.3389/fnhum.2016.00542>.
- Gunther Moor, B., van Leijenhorst, L., Rombouts, S. A. R. B., Crone, E. A., & Van der Molen, M. W. (2010). Do you like me? Neural correlates of social evaluation and developmental trajectories. *Social Neuroscience*, *5*, 461–482.
- Hämmerer, D., Li, S.-C., Müller, V., & Lindenberger, U. (2010). Life span differences in electrophysiological correlates of monitoring gains and losses during probabilistic reinforcement learning. *Journal of Cognitive Neuroscience*, *23*, 579–592.
- Imperati, D., Colcombe, S., Kelly, C., Di Martino, A., Zhou, J., Castellanos, F. X., & Milham, M. P. (2011). Differential development of human brain white matter tracts. *PLoS One*, *6*(8) e23437.
- Jones, R. M., Somerville, L. H., Li, J., Ruberry, E. J., Powers, A., Mehta, N., ... Casey, B. J. (2014). Adolescent-specific patterns of behavior and neural activity during social reinforcement learning. *Cognitive, Affective, & Behavioral Neuroscience*, *14*, 683–697.
- Kahnt, T., Park, S. Q., Cohen, M. X., Beck, A., Heinz, A., & Wrase, J. (2009). Dorsal striatal–midbrain connectivity in humans predicts how reinforcements are used to guide decisions. *Journal of Cognitive Neuroscience*, *21*, 1332–1345.
- Kool, W., Cushman, F. A., & Gershman, S. J. (2016). When does model-based control pay off? *PLoS Computational Biology*, *12* e1005090.
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost–benefit arbitration between multiple reinforcement-learning systems. *Psychological Science*, *28*, 1321–1333.
- Li, S.-C., Lindenberger, U., Hommel, B., Aschersleben, G., Prinz, W., & Baltes, P. B. (2004). Transformations in the couplings among intellectual abilities and constituent cognitive processes across the life span. *Psychological Science*, *15*, 155–163.
- Liston, C., Watts, R., Tottenham, N., Davidson, M. C., Niogi, S., Ulug, A. M., & Casey, B. J. (2006). Frontostriatal microstructure modulates efficient recruitment of cognitive control. *Cerebral Cortex*, *16*, 553–560.
- Lourenco, F. S., Decker, J. H., Pedersen, G. A., Dellarco, D. V., Casey, B. J., & Hartley, C. A. (2015). Consider the source: Adolescents and adults similarly follow older adult advice more than peer advice. *PLoS One*, *10*(6) e128047.
- Mayer, R. E. (2004). Should there be a three-strikes rule against pure discovery learning? *American Psychologist*, *59*, 14–19.
- McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*, 3515–3528.
- Monahan, K. C., Steinberg, L., & Cauffman, E. (2009). Affiliation with antisocial peers, susceptibility to peer influence, and antisocial behavior during the transition to adulthood. *Developmental Psychology*, *45*, 1520–1530.
- Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004). Computational roles for dopamine in behavioural control. *Nature*, *431*, 760–767.
- Pearce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*, 8–13.
- Potter, T. C. S., Bryce, N. V., & Hartley, C. A. (2017). Cognitive components underpinning the development of model-based learning. *Developmental Cognitive Neuroscience*, *25*, 272–280.
- Powers, K. E., Yaffe, G., Hartley, C. A., Davidow, J. Y., Kober, H., & Somerville, L. H. (2018). Consequences for peers differentially bias computations about risk across development. *Journal of Experimental Psychology: General*, *147*, 671–682.
- Raven, J. C., & Court, J. H. (1998). *Raven's progressive matrices and vocabulary scales*. Oxford, UK: Oxford Psychologists Press.
- Rodman, A. M., Powers, K. E., & Somerville, L. H. (2017). Development of self-protective biases in response to social evaluative feedback. *Proceedings of the National Academy of Sciences of the United States of America*, *114*, 13158–13163.

- Rodriguez, J. B., Heekeren, H. R., Li, S. C., & Eppinger, B. (2018). Developmental differences in the neural dynamics of observational learning. *Neuropsychologia*, *119*, 12–23.
- Sebastian, C., Viding, E., Williams, K. D., & Blakemore, S.-J. (2010). Social brain development and the affective consequences of ostracism in adolescence. *Brain and Cognition*, *72*, 134–145.
- Somerville, L. H., Haddara, N., Sasse, S. F., Skwara, A. C., Moran, J. M., & Figner, B. (2018). Dissecting "peer presence" and "decisions" to deepen understanding of peer influence on adolescent risky choice. *Child Development*. <https://doi.org/10.1111/cdev.13081>.
- Spear, L. P. (2000). The adolescent brain and age-related behavioral manifestations. *Neuroscience & Biobehavioral Reviews*, *24*, 417–463.
- Steinberg, L. (2008). A social neuroscience perspective on adolescent risk-taking. *Developmental Review*, *28*, 78–106.
- Steinberg, L., & Monahan, K. C. (2007). Age differences in resistance to peer influence. *Developmental Psychology*, *43*, 1531–1543.
- Sumter, S. R., Bokhorst, C. L., Steinberg, L., & Westenberg, P. M. (2009). The developmental pattern of resistance to peer influence in adolescence: Will the teenager ever be able to resist? *Journal of Adolescence*, *32*, 1009–1021.
- Thompson-Schill, S. L., Ramscar, M., & Chrysiou, E. G. (2009). Cognition without control: When a little frontal lobe goes a long way. *Current Directions in Psychological Science*, *18*, 259–263.
- Van Den Bos, W., Güroğlu, B., Van Den Bulk, B. G., Rombouts, S. A., & Crone, E. A. (2009). Better than expected or as bad as you thought? The neurocognitive development of probabilistic feedback processing. *Frontiers in human neuroscience*, *3*, 52.
- van den Bos, W., Bruckner, R., Nassar, M. R., Mata, R., & Eppinger, B. (2018). Computational neuroscience across the lifespan: Promises and pitfalls. *Developmental Cognitive Neuroscience*, *33*, 42–53.
- van den Bos, W., Cohen, M. X., Kahnt, T., & Crone, E. A. (2012). Striatum–medial prefrontal cortex connectivity predicts developmental changes in reinforcement learning. *Cerebral Cortex*, *22*, 1247–1255.
- van den Bos, W., Rodriguez, C. A., Schweitzer, J. B., & McClure, S. M. (2015). Adolescent impatience decreases with increased frontostriatal connectivity. *Proceedings of the National Academy of Sciences of the United States of America*, *112*, E3765–E3774.
- van den Bos, W., van Dijk, E., Westenberg, M., Rombouts, S. A. R. B., & Crone, E. A. (2011). Changing brains, changing perspectives: The neurocognitive development of reciprocity. *Psychological Science*, *22*, 60–70.
- van Duijvenvoorde, A. C. K., Zanolie, K., Rombouts, S. A. R. B., Raijmakers, M. E. J., & Crone, E. A. (2008). Evaluating the negative or valuing the positive? Neural mechanisms supporting feedback-based learning across development. *Journal of Neuroscience*, *28*, 9495–9503.
- van Hoorn, J., van Dijk, E., Meuwese, R., Rieffe, C., & Crone, E. A. (2016). Peer influence on prosocial behavior in adolescence. *Journal of Research on Adolescence*, *26*, 90–100.
- Weigard, A., Chein, J., Albert, D., Smith, A., & Steinberg, L. (2014). Effects of anonymous peer observation on adolescents' preference for immediate rewards. *Developmental Science*, *17*, 71–78.
- Weiß, R. (2006). *CFT 20-R: Grundintelligenztest Skala 2: Manual*. Göttingen, Germany: Hogrefe Verlag.