

Learning together: behavioral, computational, and neural mechanisms underlying social learning in adolescence Westhoff, B.

Citation

Westhoff, B. (2022, April 5). *Learning together: behavioral, computational, and neural mechanisms underlying social learning in adolescence*. Retrieved from https://hdl.handle.net/1887/3281632

Version:	Publisher's Version
License:	<u>Licence agreement concerning inclusion of doctoral</u> <u>thesis in the Institutional Repository of the University</u> <u>of Leiden</u>
Downloaded from:	https://hdl.handle.net/1887/3281632

Note: To cite this publication please use the final published version (if applicable).



Chapter 4

Uncertainty about others' trustworthiness increases during adolescence and guides social information sampling

This chapter is based on (manuscript under review): Ma, I., Westhoff, B., & van Duijvenvoorde, A. C. K. Uncertainty about others' trustworthiness increases during adolescence and guides social information sampling.

Abstract

Adolescence is a key life phase for developing well-adjusted social behavior. An essential component of well-adjusted social behavior is the ability to update our beliefs about the trustworthiness of others based on gathered information. Here, we examined how adolescents (N = 157, 10-24 years) sequentially sampled information about the trustworthiness of peers and how they used this information to update their beliefs about others' trustworthiness. Our Bayesian computational modeling approach revealed an adolescence-emergent increase in uncertainty of prior beliefs about others' trustworthiness. As a consequence, early to mid-adolescents (ages 10-16) gradually relied less on their prior beliefs and more on the gathered evidence when deciding to sample more information, and when deciding to trust. We propose that these age-related differences could be adaptive to the rapidly changing social environment of early and mid-adolescents. Together, these findings contribute to the understanding of adolescent social development by revealing adolescent-emergent flexibility in prior beliefs about others that drives adolescents' information sampling and trust decisions.

Introduction

Adolescence is a life-phase accompanied by a strong social reorientation (Nelson et al., 2016). Adolescents spend more time with peers (De Goede et al., 2009; Lam et al., 2014; Larson et al., 1996), are susceptible to peer influence (Albert et al., 2013; Gardner & Steinberg, 2005), and aim to achieve and maintain a positive peer status (Crone & Dahl, 2012; Gavin & Furman, 1989; LaFontana & Cillessen, 2010; Nelson et al., 2016; Steinberg & Silverberg, 1986). Violations of trust, such as social rejection, gossiping, and other negative peer interactions are exceptionally detrimental to adolescents' mental health and social development (Blakemore, 2008; Crone & Dahl, 2012). It is therefore imperative for adolescents to sample information about the trustworthiness of their peers to update their beliefs and adapt their behavior accordingly. For example, information can be sampled by asking close friends for their opinion about a specific peer or by observing how they treat others. When the sampled information indicates that the peer violates the trust of others, the adolescent should update their belief about that peer's trustworthiness (from likely trustworthy to likely untrustworthy) and adapt their behavior towards that peer accordingly.

However, sparse samples might not be representative of the peer's true trustworthiness. An untrustworthy peer might sometimes act in a trustworthy manner. Insufficient information can therefore result in erroneously trusting an untrustworthy peer or not trusting a trustworthy peer. Despite the relevance of trustworthiness information sampling to the social development during adolescence, not much is known about the age differences that take place in this process. Understanding how adolescents determine the quantity of their trustworthiness information samples, sheds light on the adaptive changes that underlie social development and may expose potential improvements.

From a social reorientation perspective, it might be intuitive to expect that adolescents focus more on peers and therefore excessively sample information about peers compared to children or adults. However, a recent study using a novel task and computational model identified three distinct factors that underlie the process of information sampling about others' trustworthiness in adults (Ma, Sanfey, et al., 2020) giving rise to more nuanced hypotheses about information sampling in adolescents. The main factors that were identified in the study were: 1. prior beliefs about trustworthiness, 2. uncertainty about the prior belief, and 3. uncertainty tolerance. The first factor, the *prior beliefs about trustworthiness*, is an individual's initial expectations about others' trustworthiness before any information is sampled. Past studies in adults show that biased prior beliefs subsequently biases how information is sampled and how the beliefs are updated in the light of new information (Chang et al., 2010; Fareri et al., 2012). For example, adults were more likely to update their beliefs about another person if the novel information was consistent with their prior beliefs (Fareri et al., 2012), and actively sample information to support their prior beliefs (Kaanders et al., 2021). Prior beliefs

about trustworthiness might show age differences across adolescence. Empirical studies have shown that initially placed trust increases from childhood to adulthood (Fett, Shergill, et al., 2014; Sutter & Kocher, 2007; van den Bos et al., 2010, 2011), suggesting a potential shift in prior beliefs about trustworthiness during adolescence (but see Flanagan & Stout, 2010). One of the aims of the current study is therefore to assess if prior beliefs about trustworthiness indeed shift during adolescence and affects information sampling or bias decisions to trust or not trust a peer.

The second important factor when sampling information is the uncertainty about prior beliefs (Ma, Sanfey, et al., 2020). This reflects the variation an individual expects in the trustworthiness of others. We first explain this concept in more detail before discussing potential changes during adolescence. For example, an individual with high uncertainty about their prior belief expects more variation in trustworthiness between different trustees. In contrast, an individual with low uncertainty about their prior belief expects that there will be little variation in trustworthiness between different trustees. There can be uncertainty about any prior belief; one individual might expect that all trustees are untrustworthy (low uncertainty), another could expect that everyone is trustworthy (low uncertainty), and yet another might expect that some are trustworthy while others are not (high uncertainty). Each new sample updates both the belief and the uncertainty. The updated result is a posterior belief and uncertainty about the posterior belief, respectively. Adults were shown to sample information until their posterior uncertainty dropped below a level to which they were tolerant to uncertainty (Ma, Sanfey, et al., 2020). Uncertainty about prior (and posterior) beliefs thereby influence the quantity of information samples, such that higher uncertainty likely result in more sampling (Ma, Sanfey, et al., 2020).



Figure 1. Illustration of how prior belief distributions update to posterior beliefs. The grid represents sampled information about trustworthiness. A green tile indicates that the sample resulted in an observation of trustworthiness, red tiles represent observations of untrustworthiness, and grey tiles are not sampled. At the start all tiles are grey as no samples have been drawn yet and therefore the current

belief distribution about trustworthiness is the prior belief distribution. The belief updates with each sample. The updated posteriors in the middle reflect an intermediate belief stage when there are 4 red and 1 green tile. The orange, green and blue lines in the plots represent three hypothetical subjects' prior distributions and their corresponding posterior distribution. These three hypothetical subjects were selected to show that the posterior beliefs can be quite different depending on the prior expectation (mean) and the prior uncertainty (variance) of the prior belief distribution. The orange prior distribution reflects the expectation that lower reciprocation probabilities are more likely. The observed outcomes exactly match that prior expectation. The posterior uncertainty therefore decreases but the posterior mean does not update. The blue prior distribution reflects a belief that higher reciprocation probabilities are more likely. The sample outcomes disconfirm this belief and the posterior therefore shows a large update and becomes more uncertain. The green prior distribution has a maximal uncertainty, i.e., a belief that all reciprocation probabilities are equally likely. The posterior then shows a both large update in the mean and a reduced uncertainty.

Little is known about the development of uncertainty about prior beliefs during adolescence, possibly because beliefs and especially uncertainty are difficult to observe directly in choice behavior and often require assessment through Bayesian computational models. Uncertainty about prior beliefs of trustworthiness likely change during adolescence, as changes take place in the set and frequency of social behaviors displayed by peers during early adolescence (e.g., courtship or competitive behavior such as gossiping). Transitioning from primary to high school exposes the adolescent to new peer groups with new group dynamics. Normatively, this novelty in the social environment should increase uncertainty about the generalizability of previously learned social behaviors (Piray & Daw, 2020; Stamps & Frankenhuis, 2016) (e.g., "childish" games such as playing tag may not be socially accepted anymore in high school). Specifically, uncertainty about prior beliefs should *increase* when the environment becomes more volatile, which leads to heightened sensitivity to new information, thereby allowing the individual to update their beliefs more with each new information sample (Behrens et al., 2007). Given the numerous changes in adolescents' social lives, we therefore expect an age-related *increase* in their uncertainty about their prior beliefs of peers' trustworthiness.

The third and final factor is *uncertainty tolerance*, which reflects the level of posterior uncertainty that an individual finds tolerable (Ma, Sanfey, et al., 2020). As mentioned earlier, this affects the sample quantity together with the uncertainty about prior beliefs, as adults sample until their posterior uncertainty drops below their uncertainty tolerance level (Ma, Sanfey, et al., 2020). Previous developmental studies showed individual and age-related differences in how tolerant adolescents are to uncertainty by using questionnaires (Boelen et al., 2010; Dekkers et al., 2017) and experimental risky choice tasks that vary the level of outcome uncertainty (Dekkers et al., 2017). In general, these studies with experimental tasks suggest that adolescents are more tolerant to uncertainty, which led them to explore risky gambles more often, compared to adults (Blankenstein et al., 2016; Tymula et al., 2012) (but see Blankenstein et al., 2018; Braams et al., 2019). One previous study explored how uncer-

tainty tolerance related to sampling information for monetary rewards. Findings showed that adolescents sample less information about lottery outcomes than children and adults, also suggesting that adolescents are more uncertainty tolerant than children and adults (Van Den Bos & Hertwig, 2017). Whereas these previous studies suggest that uncertainty tolerance is a trait that underlies risky choice, little is known about how an individual's level of uncertainty tolerance drives behavior in the social domain in adolescence. Given that information about peers is highly important for adolescents to successfully navigate their changing social environment, adolescents' uncertainty tolerance in non-social lottery tasks might not generalize to sampling information about peers and instead adolescents might become more uncertainty intolerant with age, especially from early to mid-adolescence.

In summary, here we examined how age differences in prior beliefs about trustworthiness, uncertainty about prior beliefs about trustworthiness, and uncertainty tolerance as factors that potentially may affect age-related differences in information sampling about others' trustworthiness. Participants (10-24 years, N = 157, 75 of which were boys) completed the Information Sampling Trust Game (ISTG, see Figure 2A). The Trust Game mimics characteristic consequences of trust, such that trusting is beneficial to all involved partners if reciprocated, but trust can also be betrayed. The ISTG extends this paradigm by allowing the participant to sample information about the trustee's history of trustworthiness prior to making a decision to trust or not trust.



Figure 2. Information Sampling Trust Game and data. **(A)** Task trial sequence example and payoff matrix. On each trial there are 2 players: the investor and a trustee. The participants played in the investor role and could sample a trustee's reciprocation history with other investors up to 25 times by turning tiles in a 5 by 5 grid. Green = reciprocated trust, red = betrayed trust, grey = not sampled. Investment

outcomes were not shown during the task. Six reciprocation probability conditions (r = 0.0, 0.2, 0.4, 0.6, 0.8, 1.0) generated the outcomes in the grid. It was clarified that they were playing with someone their own age, that the location of the tile was not informative, that each trial would be played with a new unknown trustee, and that the ratio green to red tiles would thus vary between trials. **(B)** Payoff matrix. Participants were told that if they invested, the trustee received the 6 tokens, which would be multiplied by 4 (24 tokens) and subsequently the trustee decided to either reciprocate by splitting the 24 tokens 50-50, or defect and keep all 24 tokens to themselves. Participants also had the option of not investing by keeping the initial endowment. **(C)** The number of samples (mean and standard error of the mean (s.e.m.)) as a function of reciprocation probability per age group (years). Age groups were created for visualization purposes only and analyses were conducted with age as continuous measure. **(D)** Proportion of investments as function of the generative reciprocation probability for each age group.

At the beginning of each trial, participants were endowed with 6 tokens which they could invest (entrust) in the trustee in a single-shot Trust Game (see Figure 2B for payoff matrix). Participants were told that these trustees previously played this game with 25 different investors in a different experiment and that their decisions to reciprocate or defect were stored in a covered 5 x 5 grid. Participants were given the opportunity to first sample information about the trustee's reciprocation history before deciding to either invest or not invest. Unbeknownst to participants, the grid outcomes were computer-generated and drawn from the following probabilities: 0.0, 0.2, 0.4, 0.6, 0.8, and 1.0, where 0.0 is completely untrustworthy (all red) and 1.0 is fully trustworthy (all green). Each subject sampled information about 60 different trustees (trials). There were no explicit sampling costs other than the time and effort involved in turning tiles. The outcomes of participants' trust decisions (invest or not invest) after sampling were not shown during the task to avoid changing meta-beliefs about the reliability of the acquired information. Instead, participants were told that 3 trials would be randomly selected at the end of the task and their average amount of tokens would be converted to money and paid to the participant (see Supplementary materials).

Results

Descriptive statistics

On average participants sampled 16.229 (SD = 7.532) of 25 times per trial. We expected based on non-social sampling studies (Fiedler & Juslin, 2006) that participants would sample more when the sample outcomes were less consistently green or red (i.e., when the outcome uncertainty was highest) and we examined the interaction with age. To this end we used a linear mixed effects model (LME4 in R (Bates et al., 2014; R Core Team, 2020) assessing the effects of outcome uncertainty (i.e., the variance in the Bernoulli distribution r (1-r) where r is the probability of reciprocation) and the linear and nonlinear age effects (polynomial models of linear and quadratic age effects) on the number of samples. In all our analyses, we report Chapter 4

the best fitting mixed-effects model results (see Supplementary materials for the full mixed effects model specification).

We found that participants sampled significantly more when the outcome uncertainty was higher, i.e., when the probability of reciprocation was closer to 0.5 (B = 2.103, p < .001, see Figure 2C). This effect interacted with the linear effect of age (B = 0.416, p < .001). There was no significant main effect of age (B = -0.467, p = .278). We conducted post-hoc Bonferroni-holm corrected analyses to further examine the interaction effect with age. This showed that the effect of outcome uncertainty was significantly less strong in early-adolescents while older age groups did not significantly differ from each other (see Supplementary materials for R code and all pairwise comparisons). Moreover, the effect of age on the number of samples was significantly stronger when the trustee was highly trustworthy (i.e., when the probability of reciprocation was closer 1.0) and when the trustee was highly untrustworthy (i.e., when the probability of reciprocation was closer to 0.0). This suggests that participants sampled less with age when trustees were very trustworthy or untrustworthy, (see Figure 2C).

Next, we used a mixed-effects model to examine whether the invest decisions (i.e., decisions to trust or not) were predicted by trustworthiness (i.e., the reciprocation probability) and whether this differed with age. As expected, we found a significant main effect of trust-worthiness (β = 3.414, *p* <0.001), showing that the likelihood of investing increased when the trustee was more trustworthy. There was a significant interaction between trustworthiness and the linear effect of age (β = 0.860, *p* < .001). Post-hoc linear mixed models per reciprocation probability were done to further examine the interaction with age. This showed that adolescents were more likely to invest in highly trustworthy trustees with age (reciprocation probability 0.8, *B* = 0.044, *p* < .001; reciprocation probability 1.0, *B* = 0.040, *p* < .001) and more likely to not invest in highly untrustworthy trustees (when the probability of reciprocation was 0.0, *B* = -0.020, *p* = .001, also see Supplementary materials for all results). Taken together, this suggests that younger participants were less likely to trust peers who were trustworthy and to some extent more likely to trust untrustworthy trustees compared to older adolescents, even though younger adolescents sampled more information about trustworthy peers than older participants (see Figure 2C and 2D).

Computational processes underlying trustworthiness information sampling

Information sampling and age differences therein were well captured by a Bayesian model of information sampling, called the Uncertainty model (Ma, Sanfey, et al., 2020). At its core, this model has a Bayesian belief distribution over the trustworthiness of the trustee. This belief distribution encompasses the prior belief and the uncertainty about the prior belief. With each sample this belief distribution is updated, resulting in the posterior belief distribution. The probability of stopping the information sampling increases as the updated uncertainty drops below the uncertainty tolerance level (see Methods for formal description).

We compared the Uncertainty model against three alternative computational models to test if trustworthiness information sampling strategies differed with age (see Methods for formal descriptions). The Sample Cost model and the Threshold model are alternative models developed in a previous study on the ISTG (Ma, Sanfey, et al., 2020) and the Count model was added to test if a simpler heuristic strategy is more prevalent in late childhood than at older ages. The *Sample Cost model* uses the Bayesian belief distribution to compute the normative solution for every state. The *Threshold model* is similar to the Uncertainty model without using a Bayesian beliefs distribution but instead it is based on the concept of sampling until the ratio between red and green tiles meets a subjective threshold. Finally, the *Count model* is the simplest model and tests if participants are insensitive to the gathered evidence and instead sample a fixed number of tiles with some variation.

The Uncertainty model fitted better than these alternative models for all ages (see Figure 3A), replicating previous findings in adults (Ma, Sanfey, et al., 2020). We assessed significance of the model fit difference by using bootstrapping to compute the 95% Cl's of the BIC differences (see Figure 3B). This showed that the Uncertainty model fitted significantly better than all other models (95%Cl of the summed BIC difference between the Uncertainty model and the Sample cost model was 95%CI [-3086, -527], the difference from the Threshold model was 95%CI [-3072, -636], and from the Count model 95%CI [-55095, -48774], see Figure 3B). We also performed Variational Bayesian Analyses (Rigoux et al., 2014) which returned a high probability that the Uncertainty model was most frequent in the comparison set (protected exceedance probability: Uncertainty model = 0.994, Threshold model = 0.005, Sample Cost model = 0.051, Count model = 0.000). Importantly, the Variational Bayesian Analyses for group comparisons returned a high probability that the five age groups have the same winning model (p(same winning model) = .907, see Figure 3C for model fit per age bin). We verified through model recovery that the models were distinguishable and through parameter recovery that the number of trials was sufficient to accurately estimate parameters (see Supplementary materials).

Chapter 4



Figure 3. Model fit results. **(A)** The Bayesian Information Criterion (BIC) per model summed over participants. Lower BIC values indicate a better fit, thus showing that the Uncertainty model fits best. BIC scores were computed for each participant and each model. Left image is zoomed in as the count model fitted much less well than the other models. Right image is not zoomed in. The grey area shows the zoom range. **(B)** 95% confidence intervals (CI) of the summed BIC difference between models. Zero indicates no difference between models. Negative values are in favor of the model before the subtraction sign, as lower BIC indicates a better fit. The Uncertainty model fits significantly better than the Sample Cost and Threshold models (95% CI does not contain zero). The BIC scores of a model pair were subtracted from each other for each participant, thereby obtaining one difference score per participant for each model pair. To assess significance, the 95% confidence interval of the BIC difference was computed using bootstrapping with 10⁵ iterations. **(C)** Uncertainty model fit across age. Age is grouped for visualization purposes. The shaded area is the s.e.m. of the data that was simulated with the participants' estimated parameters in the Uncertainty model. The line graph represents the mean and s.e.m. of the participants' actual data. The overlap between the shaded area and line graphs show that the Uncertainty model fitted well for each age group.

Prior beliefs, prior belief uncertainty, and uncertainty tolerance underlie agerelated differences in trustworthiness information sampling

Given our expected age-differences in prior beliefs, prior belief uncertainty, and uncertainty tolerance, we examined linear and non-linear effects of age on these model derived metrics using linear regression models. For each metric, we compared linear, polynomial (linear and quadratic), and logarithmic age effects and report the best fitting regression model results. We found that the *prior beliefs* increased monotonically with age (age linear, $\beta = 0.212$, *p* = .009). Though this effect was subtle, it shows that with age, participants expected peers to be more trustworthy (see Figure 4). Interestingly, the *prior belief uncertainty* strongly increased from early to mid-adolescence (ages 10-17 years) and then stabilized (Figure 4). The best fitting

age-model was polynomial, showing a significant linear ($\beta = 0.49$, p < .001) and quadratic effect of age ($\beta = -0.43$, p = .010). This suggest an adolescent-emergent increase in uncertainty about their prior beliefs of trustworthiness. Moreover, *Uncertainty tolerance* increased linearly with age ($\beta = 1.80$, p < .001, Figure 4), suggesting that adolescents gradually became more uncertainty tolerant with age. Finally, *decision noise* did not change with age (linear, $\beta = -0.12$, p = .587, quadratic, $\beta = 0.38$, p = .119). This shows that the degree to which participants' sampling choices followed the fitted Uncertainty model predictions did not change with age and therefore gives more confidence in the interpretability of the model derived metrics across adolescence.



Figure 4. Prior beliefs, prior uncertainty, and uncertainty tolerance as function of age. Note that a uniform prior corresponds to a prior uncertainty of 0.289, which is the width of the beta distribution for uniform priors. This thereby creates an upper bound for the prior uncertainty. Uncertainty tolerance has the same upper bound as prior uncertainty, as they both reflect the width of the belief distribution. Lower prior uncertainty values reflect a more certain prior. The Loess method was used to generate these plots.

No difference in prior beliefs in first and second task-half

To test if the prior beliefs changed during the task, we refitted the Uncertainty model to the first half and second half of the task separately. The prior belief is a model derived metric based on two free parameters in the Uncertainty model, called a_0 and b_0 (see Methods). Since different combinations of a_0 and b_0 can result in the same prior belief (but with different uncertainty about those prior beliefs), we conducted these analyses on the a_0 and b_0 parameter estimates rather than on the prior beliefs. Wilcoxon sign-rank tests showed no difference between the first and second task halves in either of the these two parameter estimates (a_0 estimate, z = 1.406, p = .160, median difference < 0.001, 95% CI [-0.122, 0.315]; b_0 estimate, z = 1.133, p = .257, median difference < 0.001, 95% CI [-0.478, 0.465]). Moreover, Spearman-rank correlations showed no significant relation between age and the difference between the first and second task halves in either of the two parameter (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter estimates (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter estimates (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter estimates (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter estimates (a_0 estimate) showed no significant relation between age and the difference between the first and second task halves in either of the two parameter estimates (a_0 estimate) showed no significant relation between age and the

 $r_s = -0.046$, p = .565; b_0 estimate $r_s = -0.058$, p = .573). Since uncertainty about the prior belief was also a metric derived from the a_0 and b_0 parameters (see Methods), this suggests that our findings of the prior beliefs and uncertainty about those beliefs were not likely confounded by changes over trials or age-related changes therein.

Discussion

Gathering information about outcomes of social interactions to adjust our beliefs about others is critical for successful and adaptive social behavior. Sampling and using information about others is particularly important during adolescence, as this is a developmental phase in which social cognition and peer relations rapidly develop. We found that adolescents adjusted their sampling quantity to the consistency of the sampled information by sampling more when the information was inconsistent. This effect of information inconsistency was less strong for early-adolescents compared to older adolescents. Moreover, adolescents trusted (i.e., invested) more often when peers were more trustworthy and this adaptive response to high trustworthiness became stronger with age. These behavioral findings show age differences in social information sampling and its use in trust decisions. The age differences in information over trustworthiness is updated with each new sample. The computational model showed that the age differences in sampling could be accounted for by age differences in prior beliefs about trustworthiness, uncertainty about those prior beliefs, and uncertainty tolerance.

We found a relation between age and prior beliefs, which indicated that with age, adolescents believed others to be more trustworthy before having sampled any information. This finding is consistent with previous developmental studies that have used (repeated) trust games and found that the first invested amount in the trust game tends to increase with age (Fett, Shergill, et al., 2014; Sutter & Kocher, 2007; van den Bos et al., 2010, 2011), suggesting more initially placed trust in others. Our computational model also showed that the relation between age and prior beliefs accounted for the, albeit subtle, finding that younger adolescents sampled more relative to older adolescents when the underlying reciprocation probabilities were either high or low (Figure 3C). Intuitively, when the sampled information indicated that the trustee was very trustworthy or very untrustworthy, it contradicted the prior beliefs of early-adolescents more than the older adolescents. This resulted in a sampling bias where early-adolescents sampled more relative to older adolescents when the trustee was highly trustworthy or highly untrustworthy, as they needed more information to be convinced of the information. The age-related increase in prior beliefs was subtle, yet possibly adaptive to age-related changes in the trustworthiness of peers, as earlier studies found that adolescents indeed become more trustworthy from early to mid-adolescence (Harbaugh et al., 2003; Sutter & Kocher, 2007).

The strongest and most striking age difference was found in the uncertainty about prior beliefs. Confirming our hypothesis, prior beliefs about others' trustworthiness rapidly became more uncertain with age from early to mid-adolescence. This adolescence-emergent increase in prior uncertainty led adolescents to rely less on their prior beliefs and more on their sampled information. Behaviorally, this resulted in adolescents rapidly adapting their sampling behavior to the information inconsistency from early to mid-adolescence. In other words, early to mid- adolescents became more open-minded about possible individual differences in trustworthiness between peers. The increase in uncertainty about prior beliefs is not likely explained by age-related differences in sampling strategies or task comprehension for at least four reasons. Firstly, we ruled-out alternative computational models that represented alternative sampling strategies, showing that participants of all ages used the same information sampling strategy in their decisions to continue sampling. Second, the instructions were read aloud and in-person by the experimenters who made sure all participants understood the task by asking comprehension guestions. Third, age differences in task comprehension cannot account for all age-related effects, such as the relation between age and prior beliefs. Finally, the prior belief distribution parameter estimates did not differ between the first and second task half, which indicates that the age-related increase in prior uncertainty was not due to age differences in sampling strategies or task comprehension as the task progressed.

The increase in uncertainty about prior beliefs may be adaptive given the numerous changes that take place in early to mid-adolescents' social environment, including changes in social behavior induced by the peers' pubertal stage and a transition of schools. Interest-ingly, given that the changes in the adolescent's environment occur mostly in social contexts but less so in non-social contexts (e.g., physics such as gravity mostly remain constant), this further suggests that learning flexibility in adolescents might be especially strong in social contexts. This pertains to learning about others or about the self in relation to a peer group. Future within-subjects studies are required to examine if the age-related increase in uncertainty about prior beliefs is indeed specific to social contexts, and in which social contexts this may be most pronounced.

We furthermore found that uncertainty tolerance increased linearly with age. This finding at first seems to contradict previous studies on uncertainty tolerance in non-social contexts, which suggest that adolescents are more uncertainty tolerant than adults (Blankenstein et al., 2018; Tymula et al., 2012; Van Den Bos & Hertwig, 2017). However, in the Uncertainty model, uncertainty updates with each sample and the probability to stop sampling increases once uncertainty drops below the individual's uncertainty tolerance level. Therefore, uncertainty tolerance should be interpreted in combination with uncertainty about prior beliefs. In our study, mid-adolescents relative to early-adolescents showed increased uncertainty about prior beliefs

Chapter 4

combined with a smaller increase in uncertainty tolerance. In the model this combination results in an increase in sampling (see model simulations in Figure 3C). After mid-adolescence, uncertainty about prior beliefs stabilized while uncertainty tolerance continued to increase. In the model, this combination resulted in a decrease in sampling from mid to late-adolescence (Figure 3C). Although we did not test non-social information sampling in this study, the idea of at least some degree of social specificity is corroborated by previous studies on non-social information sampling. Those studies found that adolescents gathered less information than children or early-adolescents prior to a risky financial decision (Bowler et al., 2020; Van Den Bos & Hertwig, 2017), which diverges from our findings in a trust context. Taken together, our findings of age-related differences in the factors that underlie information sampling fit with our notion that mid-adolescents may attempt to learn more about their social environment.

While reinforcement learning studies on adolescence are scarce, one advantage of this computational modeling approach is using formal model comparisons to assess age-related differences in decision-making strategies (Cohen et al., 2020; Kabotyanski et al., 2019; Palminteri et al., 2016). For example, a previous study showed that adolescents used different reinforcement learning strategies than adults by not benefitting from counterfactual feedback. which would have been challenging to conclude without the specific behavioral predictions that result from fitting computational models (Palminteri et al., 2016). In the current study, we ruled-out a family of normative models (i.e., Sample Cost model variants), heuristic models that were not based on Bayesian belief distributions (Threshold model variants), and the Count model for insensitivity to gathered evidence. We found that the Uncertainty model showed a good fit and fitted best for all ages, showing that across age, adolescents use their uncertainty in Bayesian belief distributions to update their beliefs about trustworthiness. This is consistent with a previous study on social information sampling costs in adults, where the Uncertainty model also fitted the data best (Ma, Sanfey, et al., 2020). We show that within this winning model, age-related differences in the parameter estimates accounted well for age differences in sampling behavior. Our computational modeling and model comparison approach therefore contributes to the field's understanding of age differences in cognitive strategies of social decision-making.

The potential applications of our approach extend to understanding how peer-status could influence prior beliefs about others' behavior. For instance, children who are frequently socially rejected by their peers may develop different prior beliefs over the trustworthiness of others compared to stably accepted children (e.g., those with experience of frequent rejection may have a highly certain prior belief that others are untrustworthy). Moreover, previous studies used behavioral economic games such as trust games to reveal aberrant social decision-making in psychiatric disorders (Hinterbuchinger et al., 2018; King-Casas & Chiu, 2012; Robson et al., 2019), including anxiety disorders, autism spectrum disorder (Izuma et al., 2011), borderline

personality disorder (King-Casas & Chiu, 2012; Seres et al., 2009; Unoka et al., 2009), and ADHD (Ma et al., 2016, 2017). In future studies, our task and models might further shed light on how individuals actively sample and use information to initiate or avoid social interactions and how this may depend on aberrant prior beliefs.

Methods

Participants and experiment procedure

A total of 157 adolescents (of which 75 boys) completed the experiment (range = 10-24 years, M = 17.50, SD = 4.34). The sample size was based on prior studies examining age-related differences in uncertainty tolerance within a comparable age-range (Blankenstein et al., 2016; Van Den Bos & Hertwig, 2017). Participants were screened for color blindness, psychiatric and neurological disorders, IQ was estimated by using subtests of the WISC and WAIS. IQ scores fell in the normal range (M = 107.5, SD = 10.9, range = 80-135), and did not correlate with age ($r_c = 0.119$, p = .138), parental social economic status (SES) was estimated by highest educational attainment of the caregiver (s). This sample generally showed a medium to high SES level and SES did not show a relationship with age (Kruskal-Wallis rank sum returned χ^{2} (4, n = 157) = 6.342, p = .175; low SES n = 8, medium SES n = 59, high SES n = 90). All procedures were approved by the institutional review board of the Leiden University Medical Center, and performed in accordance with the relevant guidelines and regulations. Written informed consent was given by adult participants, and by their legal guardians in the case of minors (minors provided written assent). This behavioral study was part of a larger imaging study. All participants performed the task in a quiet room near the neuroimaging labs of Leiden University. The task took approximately 30 minutes to complete (see Supplementary materials for payoff procedure).

Computational model

The Uncertainty model is based on the concept of sampling to reduce uncertainty until the subjective uncertainty tolerance is met. The model consists of four components: a prior belief distribution over the reciprocation probability (*r*), an evolving posterior distribution over *r*, the uncertainty tolerance, and decision noise. As explained above, individuals start with a prior belief distribution when nothing has been sampled yet. This prior belief distribution encompasses both the prior belief and the uncertainty about the prior belief. When information is sampled, the belief distribution is updated and called a posterior distribution. The posterior distribution is therefore a combination of the prior belief distribution and the sampled information. Information is sampled sequentially and each new sample results in a new update. The degree to which the prior beliefs and the uncertainty about prior beliefs update with each

new sample depends on the values of the prior belief and uncertainty about prior beliefs and whether the prior belief is confirmed or disconfirmed by the samples. Examples of how belief updates depend on prior belief distributions and samples are depicted in Figure 1.

Formally, in our model the *state* is defined by the number of turned green tiles (n_{\downarrow}) and the number of turned red tiles (n_{\cdot}) . The *actions* are to either sample or stop sampling until all 25 tiles are sampled. Specifically, the model assumes that people do not know the trustee's exact trustworthiness by using a Bayesian belief distribution over the possible range of *r*. The conjugate prior belief distribution over *r* is a beta distribution with parameters a_0 and β_0 . As information is sampled, the evolving posterior distribution is:

$$p(r|n_+, n_-) = \text{Beta}(r; \alpha, \beta) \tag{1}$$

The Beta distribution consists of parameters a and β . Here, $a = a_0 + n_1$ and $\beta = \beta_0 + n_2$.

As shown in Figure 1, the more uncertain a prior is, the more a new sample will reduce that uncertainty. In addition, if a sample is highly consistent with the prior expectation, the posterior mean will shift less than when a sample disagrees with the belief (Stamps & Frankenhuis, 2016). Uncertainty of the belief is operationalized using the standard deviation of the posterior distribution:

Uncertainty
$$(\alpha, \beta) = \sqrt{\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}}$$
 (2)

As sampling decreases uncertainty and at some point, the uncertainty will reach the subject's uncertainty tolerance k, i.e., how much uncertainty is tolerated by the subject. As the uncertainty reduces and approaches the tolerance, the probability that the subject takes another sample becomes smaller. These probabilities are given through the softmax function, allowing for decision noise τ :

$$p(\text{sample}|\alpha,\beta) = \frac{1}{1 + e^{\frac{Uncertainty(\alpha,\beta)-k}{\tau}}}$$
(3)

Where a larger k reflects more uncertainty tolerance and a larger τ reflects more decision noise.

We fitted the model for each participant individually using a log likelihood optimization algorithm as implemented in the fmincon routine in MATLAB (Mathworks) using 100 combinations of starting points to avoid local minima. Four free parameters are fitted for each subject: α_0 , β_0 , k, and τ . In addition, we used between-subjects Bayesian Model Selection

(Rigoux et al., 2014; Stephan et al., 2009) to assess variation in the best fitting model between different ages using five age groups.

Alternative models

We also considered three families of alternative models and found that these did not fit as well as the Uncertainty model (formal descriptions in Supplementary materials): The *Sample Cost model*, which uses the Bayesian belief distribution to compute the normative solution for every state. The *Threshold model* is a heuristic model that does not use Bayesian beliefs distributions. For model comparisons we calculated the difference between model evidence in terms of BIC for each model pair for each subject. The Count model is the simplest heuristic model. In this model, a fixed number of samples are drawn with some variation. Unlike the other models, the decision or stop sampling is therefore not dependent on the outcomes of previous samples. To assess the significance of the model fit differences, we used bootstrapping to compute the 95% confidence intervals of the summed difference in BIC using 10⁵ iterations (for within-model results see Table S1).

Data and code availability

The data of this study are openly available in Open Science Framework osf.io/upmwv and the code is available on Github on request <u>https://github.com/ili-ma/Social_Belief_Updates_Adolescence</u>.

Supplementary materials

Participant recruitment

Participants were recruited through flyers and presentations at local high schools in the Netherlands, as well as a recruitment website. Based on previous studies (see methods section), we aimed to recruit 160 participants equally divided across 5 age bins (10-12; 13-15; 16-18; 19-21; 22-24 years) with a balanced gender distribution for each age bin. Eventually, 159 people came to the laboratory to participate in our study protocol. One person did not start the study due to anxiety. One person could not finish the sampling experimental task due to time constraints. This led to the reported *N* of 157 adolescents.

Participant bonus fee

Computational models

Note that for all models, we tested different variants within each model, and selected the best fitting version for between model comparisons. Here, we denote the best fitting version of each model, for all versions see Ma et al., (2018).

The Sample Cost model

The Sample Cost model uses the Bayesian belief distribution over trustworthiness to compute the expected utility of sampling and stopping for every possible state in the task through forward reasoning. It consists of four components: prior beliefs over the trustworthiness (*r*), an evolving posterior distribution over *r* iterative maximization of future expected utility under this posterior distribution, and decision noise. The conjugate prior over *r* is a beta distribution with priors a_0 and β_0 , and parameters $a = n_+ + a_0$ and $\beta = n_- + \beta_0$, where n_+ is the number of green samples and n_- is the number of red samples. The posterior over *r* is:

$$p(r|n_+, n_-) = Beta(r; \alpha, \beta)$$

Investing results in either reciprocation (outcome = 1 with probability *r*, then the investment amount is multiplied by m = 2) or betrayal (outcome = 0 with probability 1 - *r*, then the investment amount is multiplied by m = 0). The agent does not know *r* and therefore has to marginalize over *r*, using the current posterior. This gives the conditional distribution of outcome given α and β :

$$p(\text{outcome} = 1|\alpha, \beta) = \int p(\text{outcome} = 1|r)p(r|\alpha, \beta)dr = \int rp(r|\alpha, \beta)dr = \frac{\alpha}{\alpha + \beta} \quad (1)$$

The expected utility of *not* investing is $U_0 = 1$. The expected utility of investing, U_1 , with a free parameter λ for risk attitude becomes:

$$U_{1}(\alpha,\beta) = \mathbb{E}[\text{outcome}|\alpha,\beta] - \lambda \text{Var}[\text{outcome}|\alpha,\beta]$$
$$= \frac{m\alpha}{\alpha+\beta} - \frac{\lambda m^{2}\alpha\beta}{(\alpha+\beta)^{2}}$$
(2)

The agent can decide between sampling (a = 1) or stopping (a = 0) at any time except when t = T+1, then all boxes are opened and only stopping (a = 0) is possible. The value of a state-action pair is given by the Bellman equations (Bellman, 1952). Specifically, the expected value of the state (α , β ,) is the higher of the expected utilities of not investing and investing:

$$Q_t(\alpha,\beta;\alpha=0) = max\{U_0, U_1(\alpha,\beta)\}$$
(3)

When t = T+1, the value of the state $(\alpha + \beta)$ is $V_t(\alpha, \beta) = Q_{T+1}(\alpha, \beta; \alpha = 0)$. At earlier times, V_t is the larger of the expected utilities:

$$V_t(\alpha,\beta) = \max\{Q_t(\alpha,\beta;\alpha=0), Q_t(\alpha,\beta;\alpha=1)\}$$
(4)

The expected value of a sampling action (a = 1) at time *t* in the state α , β subtracting the subjective cost of a sample *c* is:

$$Q_t(\alpha,\beta; a=1) = \frac{\alpha V_{t+1}(\alpha+1,\beta) + \beta V_{t+1}(\alpha,\beta+1)}{\alpha+\beta} - c$$
(5)

By starting with the final state (when *n*, we can obtain the Sample Cost solution for every possible state (dynamic programming). In the Sample Cost model, the decision variable (DV), is the difference between the utilities of sampling and stopping:

$$DV(\alpha,\beta) = Q_t(\alpha,\beta;\alpha=1) - Q_t(\alpha,\beta;\alpha=0)$$
(6)

The Sample Cost policy would be to sample when the DV is positive; however, we introduce decision noise through a softmax function:

$$p(\text{sample}|\alpha,\beta) = \frac{1}{1 + e^{-\frac{DV(\alpha,\beta)-k}{\tau}}}$$

This model version has six free parameters: a_0 , β_0 , c, λ , τ and k which were fitted for each subject.

The Threshold model

In the *Threshold Model*, the agent keeps track of the absolute difference between positive and negative information and stops sampling when this difference reaches a bound. However, to be consistent with our other models and with most of the value-based decision literature, we use soft rather than hard bounds. The bound *b* takes three possible values, depending on the sign of $n_{+} - n_{-}$:

$$b = \begin{cases} b_{+} & \text{if } n_{+} > n_{-} \\ \frac{b_{+} + b_{-}}{2} & \text{if } n_{+} = n_{-} \\ b_{-} & \text{if } n_{+} < n_{-} \end{cases}$$
(7)

The probability that the agent stops sampling is a logistic function of the difference between this decision variable and a bound *b*:

$$p(\text{sample}|n_+, n_-) = \frac{1}{1 + e^{-\frac{DV(n_+, n_-) - b}{\tau}}}$$
(8)

This model version has three free parameters: b_1 , b_2 , and τ which were fitted for each subject.

The Count model

In the count model, the agent takes a fixed number of samples k with some variation. This model represents a simple heuristic where the decision to sample another tile is not determined by the outcomes of the previous samples. We added this model to test if some individuals and especially younger participants might have a strategy of for example always sampling five tiles before making their decision to trust or not trust. The softmax function allows for decision noise:

$$p(\text{sample}|n_+, n_-) = \frac{1}{1 + e^{\frac{(n_+ + n_-) - k}{\tau}}}$$
(9)

This model has two free parameters: k and τ which were fitted for every subject.

Within-model comparisons results

For each model, we first fitted the basic version - the version with the fewest free parameters - and then tested whether adding a free parameter significantly improved the model fit (also see (Ma et al., 2018) for more details). The assessment of these additional free parameters was based on the trust game literature. Here, we describe the additionally tested free parameters for each model across all participants. The *Count model* is not included in the within-model comparisons, as it only had one version of the model, which had two free parameters; one for the fixed number of samples and one for variability.

Uncertainty model: prior beliefs improve the model fit. For the same reasons as described earlier, we tested the improvement in model fit when adding a prior belief, also estimated for each subject. In the basic version of our model, we used an uninformative, uniform prior (i.e., $a_0 = 1$ and $b_0 = 1$). In the model version with prior beliefs, we fitted a_0 and b_0 as free parameters for every subject. We found that allowing for a subjective prior improved the model fit (Table S1). The median of individual prior means was r = 0.47 (bootstrapped 95% CI [0.45, 0. 50]).

Sample Cost model: prior beliefs and risk attitude both improve the model fit. The repeated trust game literature suggests that subjective prior beliefs (Chang et al., 2010) and betrayal aversion (Aimone & Houser, 2012) both play important roles in determining individual differences in trust. In our paradigm, an individual's prior belief about trustworthiness is a beta distribution with two free parameters (a_0 and b_0). Betrayal attitude is operationalized as the variance of the outcome, multiplied by a free parameter *l*, which is subtracted from the expected utility of trusting. The median betrayal attitude parameter was estimated at 0.16 (bootstrapped 95% CI [0.10, 0.23]), which shows that participants were overall betrayal-averse.

Threshold model: asymmetric bounds but not collapsing bounds improved the model fit. The literature on Drift Diffusion Models in perception literature suggests that the model fits better with collapsing bounds that reflect an urgency signal (Tajima et al., 2016), or asymmetric bounds (Mulder et al., 2012) as positive sample outcomes might be weighted differently than negative outcomes. Using separate bounds for positive than for negative samples did indeed improve the model fit. Second, we tested the model when the bounds "collapse" to zero over time. This did not improve the model fit from the model with asymmetric bounds.

Table S1.1. Within-model comparison results.

		95% CI	
	Summed ΔBIC	Lower bound	Upper bound
Uncertainty basic vs. priors	2013	925	3316
Sample Cost basic vs. risk attitude	1597	1006	2293
Sample Cost risk attitude vs. risk attitude + priors	791	398	1205
Threshold basic vs. two bounds	44987	38479	51847
Threshold collapsing bound vs two bounds	847	501	1198

95% CI = Bootstrapped 95% confidence interval of the summed difference between model fits. Smaller BIC values indicate better fit. Thus, positive values indicate a better fit for the second model. The models were fitted to the data at the individual level using a log likelihood optimization algorithm as implemented in the fmincon routine in MATLAB (ÓMathworks). The optimization was iterated 100 times with varying initiations to avoid local minima. Because of the summation of the difference, large positive or negative numbers therefore reflect that one model wins consistently, i.e. for most subjects.

Between- model comparison results

We then compared the best fitting version of each model with the best fitting model of the other models. Table S2 shows the results which indicate that the Uncertainty model fitted best.

Pairwise comparison	Summed ABIC	Winning	Correlation between
between models	and 95% Cl	model	age and ∆BIC
Uncertainty - Sample Cost	-1685 [-3086, -527]	Uncertainty	r _s = -0.059, p = .464
Uncertainty - Threshold	-1741 [-3072, -636]	Uncertainty	$r_{\rm s} = 0.101, p = .221$
Uncertainty - Count	-51940 [-55095, -48774]	Uncertainty	r _s = -0.349, p < .001
Sample Cost - Threshold	-56 [-627, 504]	None	r _s = 0.031, p = .711
Sample Cost - Count	-50243[-52923, -47489]	Sample Cost	r _s = -0.385, p < .001
Threshold - Count	-50151[-52836, -47349]	Threshold	r _s = -0.364, p < .001

Table S1.2. Between model comparisons.

Lower BIC values indicate a better fit, thus showing that the Uncertainty model fits significantly better than all other models. BIC scores were computed for each participant and each model. The BIC scores of a model pair (left column) was then subtracted from each other, thereby obtaining one difference score per participant for each model pair. The middle column shows the sum of the difference across participants and the 95% confidence interval of the BIC difference, computed using bootstrapping with 10⁵ iterations. The last column shows the Spearman-rank correlation between age and each model pair's BIC difference.

Descriptive statistics number of samples

The descriptive statistics were obtained through generalized linear mixed models in R (package LME4). We assessed the fit of polynomial models (linear, quadratic) using the anova() function in R. For the number of samples, the linear age model fitted better than the polynomial (linear and quadratic) models.

For the analyses of the number of samples, we first calculated the outcome uncertainty, which is the variance in the Bernoulli distribution:

Outcome uncertainty = r(1-r)

Where *r* is the probability of reciprocation.

For transparency we report the full mixed effects model in R code here: SamplingModel <- Imer(number of samples ~ scale(outcome uncertainty) * scale(age) + (1| subject), data = df)

Dradictors	Sample decisions			
Fieuciois	β	95% CI	р	
Intercept	16.21	15.37 – 17.05	<.001	
Outcome uncertainty	2.10	2.00 - 2.20	<.001	
Age	-0.47	-1.31 – 0.37	.276	
Outcome uncertainty x Age	0.42	0.32 - 0.52	<.001	
Random Effects				
σ^2	24.46			
τ _{00 subject}	28.44			
ICC	0.54			
N _{subject}	157			
Observations	9420			
Marginal R ² / Conditional R ²	0.083 / 0.579			

Table S2.1. Results number of samples mixed model

Post hoc test for the number of samples

To further examine the interaction effect, we created age bins consistent with Figure 2C and compared the slopes for outcome uncertainty between the age groups using the emtrends() function. This showed that only the early adolescents differed significantly from all other age groups (Table S2.2). We used the following code:

emtrends(SamplingModel, pairwise ~ ageGroup, var="outcome uncertainty")

119

We also created bins for the probability of reciprocation consistent with Figure 2C and used the emtrends() function to compare the slopes of the age effect between the different reciprocation probabilities. The effect of age on the number of samples was strongest in the highest and lowest reciprocation probabilities (Table S2.2). We used the following code: emtrends(ReciprocationModel, pairwise ~ reciprocation probability bins, var="age")

Contrast age groups	estimate	t ratio	p-value
10-12 vs 13-15	-9.217	-5.722	<.0001*
10-12 vs 16-18	-12.081	-7.441	<.0001*
10-12 vs 19-20	-8.619	-5.351	<.0001*
10-12 vs 21-24	-14.565	-8.971	<.0001*
13-15 vs 16-18	-2.863	-1.778	.387
13-15 vs 19-20	0.598	0.375	.996
13-15 vs 21-24	-5.348	-3.320	.008
16-18 vs 19-20	3.462	2.149	.120
16-18 vs 21-24	-2.484	-1.530	.543
19-20 vs 21-24	-5.946	-3.691	.002*
Contrast reciprocation probability	estimate	t ratio	p-value
0-0.2	-0.087	-2.180	.247
0-0.4	-0.194	-4.833	<.0001*
0-0.6	-0.189	-4.712	<.0001*
0-0.8	-0.021	-0.521	.995
0-1.0	0.104	2.595	.099
0.2-0.4	-0.106	-2.653	.085
0.2-0.6	-0.102	-2.532	.115
0.2-0.8	0.067	1.659	.559
0.2-1.0	0.192	4.775	<.0001*
0.4-0.6	0.005	0.121	1.000
0.4-0.8	0.173	4.312	<.001*
0.4-1.0	0.298	7.428	<.0001*
0.6-0.8	0.168	4.191	<.001*
0.6-1.0	0.293	7.307	<.0001*
0.8-1.0	0.125	3.116	.023

Table S2.2. Results post-hoc pairwise comparison per age group

Note: *p-values are significant at the Bonferroni-holm corrected level.

Descriptive statistics invest decisions

For the invest decisions, the model with the linear effect of age also fitted better than the model with the polynomial effects of age (linear and quadratic). The mixed effects model was specified as follows:

InvestModel <- glmer(Invest decision ~ scale(reciprocation probability) * scale(age) + (1| subject), data = investdata, family = binomial, control = glmerControl(optCtrl = list(maxfun = 1e+9), optimizer = c("bobyqa")))

Predictors	Odds Ratios	95% CI	β	р
Intercept	0.62	0.53 - 0.72	-0.478	<.001
Reciprocation probability	30.38	26.15 - 35.28	3.414	<.001
Age	0.99	0.85 – 1.15	-0.011	.888.
Reciprocation probability x age	2.36	2.07 - 2.70	0.860	<.001
Random Effects				
σ^2	3.29			
τ _{00 subject}	0.73			
ICC subject	0.18			
N _{subject}	157			
Observations	9420			
Marginal R ² / Conditional R ²	0.755 / 0.800			

Table S2.3. Results invest decisions mixed model

Post hoc test for invest decisions

We were interested in whether the effect of age was present in all reciprocation probabilities. We therefore conducted post hoc tests per probability of reciprocation using the lmer() function in R for each probability of reciprocation separately:

posthoc_model <- Imer(Invest decision ~ scale(age) + (1| subject), data = df_reciprocationBin)

This showed that the effect of age was only significant when the probability of reciprocation was high (0.8 and 1.0) or lowest (0.0, see Table S2.3).

Probability of reciprocation	estimate	t value	n-value	
	cotimate	t value	p value	
0.0	-0.020	-3.333	.001*	
0.2	-0.018	-2.312	.022	
0.4	-0.010	-0.691	.491	
0.6	0.020	0.999	.319	
0.8	0.044	4.733	<.0001*	
1.0	0.040	5.108	<.0001*	

Table S2.4 post hoc test for the effect of age on invest decisions per reciprocation probability

Note: *p-values are significant at the Bonferroni-holm corrected level.

Expected reward increased with age

We examined if the expected values based on the trust decisions varied with age. We first normalized the outcomes on the endowment of 6 tokens. Thus, the expected value for not investing is 1. The expected value for investing is computed as the multiplier of 2 times the reciprocation probability. Since the reciprocation probabilities were 0.0, 0.2, 0.4, 0.6, 0.8, and 1.0, the possible average expected values ranged from $\frac{0+0.4+0.8+1+1+1}{6} = 0.67$ to $\frac{-1+1+1+1.2+1.6+2}{6} = 1.3$. All subjects (with exception of one), had an average expected reward higher than 1 (see Figure S1). This confirms that all subjects would - on average - have gained money on top of their original endowment by trusting when trusting was beneficial. However, younger adolescents earned less on this task than older adolescents, consistent with their deviations in low (0.0) and high (0.8, 1.0) investment probabilities. This was shown by a significant correlation between age and the average expected reward in the task ($r_s = 0.387$, p < .001).



Figure S1. Expected reward per subject as function of age. The expected reward increases with age. Error bars indicate the s.e.m. The line and shaded region indicates the mean and s.e.m across subjects. Color indicates age in months.

Age-related changes in Uncertainty model parameter estimates

Linear regressions were used to examine the relationship between age and the Uncertainty model parameter estimates using the lm() function in R. We first tested whether adding age as a quadratic term with the poly() function improved the model fits when compared to a linear term only using the anova() function in R (Phillips, 2017). For the prior uncertainty the quadratic term did improve the model fit (F(1,154) = 6.835, p = .010). For all other parameter estimates the quadratic effect of age did not improve the model fit (prior mean F(1,154) = 3.861, p = .051; uncertainty tolerance F(1,154) = 1.940, p = .166; decision noise F(1,154) = 2.255, p = .119). We applied a Bonferroni-holm correction for multiple testing, which includes the model improvement test for each of the four parameters, a linear age effect test for three of these parameters, and two test (linear and quadratic age effects) for one parameter, corresponding to a total of 9 tests.

Model recovery

We performed model recovery and verified that the models were distinguishable. To this end, we simulated data with each model and subsequently fitted the simulated data to each model. If the model is recoverable, then the best fitting model should be the model that generated the data. We show that this was indeed the case as the model fit was always best for the model that generated the data (Table S3.1).

Parameter recovery

To check if the parameters were recoverable, we simulated data with each model using varying parameter values as inputs. We randomly selected the estimates of 20 subjects to simulate data with an equal number of trials as in the actual task. We then fitted that simulated data to the model to check if the parameters values that generated the data were correctly estimated. We subtracted the estimated parameters of the simulation from the actual input parameters and calculated the bootstrapped 95% confidence interval and found no significant difference between the input parameters and the estimated parameter values. This returned: Uncertainty tolerance median difference = -0.0001, 95% CI [-0.000, 0.000]; a_0 median difference = -0.386, 95% CI [-0.590, 0.000]; b_0 median difference = -0.250, 95% CI [-0.580, 0.000]). This suggests that the parameters were indeed recoverable.

Table S3.1. Model recovery results

	95% CI		
	Summed ∆BIC	Lower bound	Upper bound
Data generated by the Uncertainty model			
Uncertainty vs. Sample Cost	-849	-1057	-654
Uncertainty vs. Threshold	-568	-861	-293
Uncertainty vs Count	-7014	-8129	-5859
Data generated by the Sample Cost model			
Sample Cost vs. Uncertainty	-795	-313	-541
Sample Cost vs. Threshold	-537	-908	-249
Sample Cost vs Count	-7350	-8675	-5978
Data generated by the Threshold model			
Threshold vs. Sample Cost	-827	-964	-689
Threshold vs. Uncertainty	-654	-894	-448
Threshold vs Count	-6102	-6923	-5179
Data generated by the Count model			
Count vs. Sample Cost	-2193600	-2195896	-2191370
Count vs. Uncertainty	-2563	-2917	-2224
Count vs. Threshold	-2190400	-2192665	-2188118

95% CI = Bootstrapped 95% confidence interval of the summed difference between model fits. Negative values indicate a better fit of the model that generated the data. The data were generated using the participants' parameter estimates and shows that all models were recoverable.