Human development of the ability to learn from bad news

Christina Moutsiana[a,1], Neil Garrett[a], Richard C. Clarke[b], R. Beau Lotto[b], Sarah-Jayne Blakemore[c], and Tali Sharot[a]

[a]Department of Cognitive, Perceptual, and Brain Sciences, University College London, London WC1H 0AP, United Kingdom; [b]Institute of Ophthalmology, University College London, London EC1V 9EL, United Kingdom; and [c]Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, United Kingdom

Edited by Michael S. Gazzaniga, University of California, Santa Barbara, CA, and approved July 29, 2013 (received for review March 24, 2013)

Humans show a natural tendency to discount bad news while incorporating good news into beliefs (the "good news–bad news effect"), an effect that may help explain seemingly irrational risk taking. Understanding how this bias develops with age is important because adolescents are prone to engage in risky behavior; thus, educating them about danger is crucial. We reveal a striking valence-dependent asymmetry in how belief updating develops with age. In the ages tested (9–26 y), younger age was associated with inaccurate updating of beliefs in response to undesirable information regarding vulnerability. In contrast, the ability to update beliefs accurately in response to desirable information remained relatively stable with age. This asymmetry was mediated by adequate computational use of positive but not negative estimation errors to alter beliefs. The results are important for understanding how belief formation develops and might help explain why adolescents do not respond adequately to warnings.

decision making | learning | optimism

Human decision making is markedly influenced by beliefs of what might occur in the future. We form and update those beliefs based on information we receive from the world around us. However, even when we are presented with accurate information, cognitive biases and heuristics restrict our ability to make adequate adjustments to our prior beliefs (1–5).

One such bias, with important implications for well-being, is the human tendency to discount bad news. For instance, highlighting previously unknown risk factors for diseases is surprisingly ineffective at altering an individual's perception of their medical vulnerability (6). On the other hand, when people are informed that they are less at risk for encountering adverse events (e.g., car accidents or a sport injury) than they previously thought, they will alter their beliefs appropriately (7). One may view this process as a greater tendency to use priors as an anchor (1) when subsequent information is undesirable vs. desirable. It has also been suggested that this bias can exist because the computational and neural processes that govern learning from good and bad information are partially separable (7–10). When receiving desirable news (such as learning that they are more attractive than they imagined), people tend to integrate this information in a Bayesian manner (8). However, upon receiving undesirable news, people's posterior beliefs are noisy and deviate from Bayesian predictions, discounting the strength of the signal they receive (8). If altering beliefs in response to desirable and undesirable information is mediated by partially distinguishable mechanisms, they may also have different development trajectories.

Understanding how the ability to alter beliefs about vulnerability develops with age is of great importance because adolescents are especially prone to engage in risky and dangerous behavior (11, 12). This can lead to grave outcomes. Indeed, the primary cause of adolescent mortality in the Western world is from accidents related to risk taking (13, 16). It is therefore critical to educate adolescents about risk and inform them of the likelihoods of unwanted outcomes.

When communicating risk, however, we must take into account the many changes in information processing and memory that take place in childhood and adolescence (for review, see refs. 17 and 18). For example, both gist and verbatim memory improve during development, with the former improving more rapidly (ref. 19; see ref. 18 for implications for policy). However, we do not know whether children and teenagers incorporate facts into prior beliefs in a valence-dependent manner. If they do, then this bias may be one of the reasons why campaigns targeted at adolescents that highlight the dangers of careless driving (20), unprotected sex (21), and alcohol and drug abuse (22) have limited impact. A prominent explanation for this limitation is that the obtained knowledge fails to alter subsequent behavior due to competing emotional and social factors that influence action (23). In younger individuals, emotional factors are especially likely to influence action because brain systems involved in emotion regulation and cognitive control have yet to mature (24, 25). However, increased emotional influence may bias not only action selection but also the mechanism by which information is integrated into existing beliefs.

Here we test whether the ability to alter beliefs in response to good and bad news develops differently with age. To that end, 59 volunteers between the ages of 9 and 26 were tested on the belief update task (7, 26). They estimated their likelihood of experiencing 40 adverse life events (e.g., passenger in a car accident or home burglary). After each trial, participants were presented with the actual frequency of that event in their representative population (Fig. 1*A*). Participants were then asked in a second session to reestimate their likelihoods for all life events. This enabled us to quantify how participants adjust their beliefs in response to new information in two instances: (*i*) when they learn that the average likelihood of encountering a negative life event is lower than their own estimate (good news; Fig. 1*B*) and (*ii*) when it is greater (bad news; Fig. 1*C*).

Formal models suggest that learning from information that disconfirms one's expectations is mediated by a prediction error signal that quantifies a difference between expectation and outcome (27). We have previously shown that an analogous mechanism underpins belief updating in this task (7). Specifically, the difference between participants' initial estimations and the information provided (that is, estimation error = estimation − probability presented) predicts subsequent updates, as would be expected from learning models (27). The strength of this association is indicative of learning. In adults, such learning is valence-dependent, being greater for information that offers

---

Author contributions: T.S. designed research; C.M., N.G., R.C.C., and R.B.L. performed research; C.M., N.G., and T.S. contributed new reagents/analytic tools; C.M., N.G., S.-J.B., and T.S. analyzed data; and C.M., S.-J.B., and T.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence should be addressed. E-mail: christina.moutsiana@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1305631110/-/DCSupplemental.

www.pnas.org/cgi/doi/10.1073/pnas.1305631110

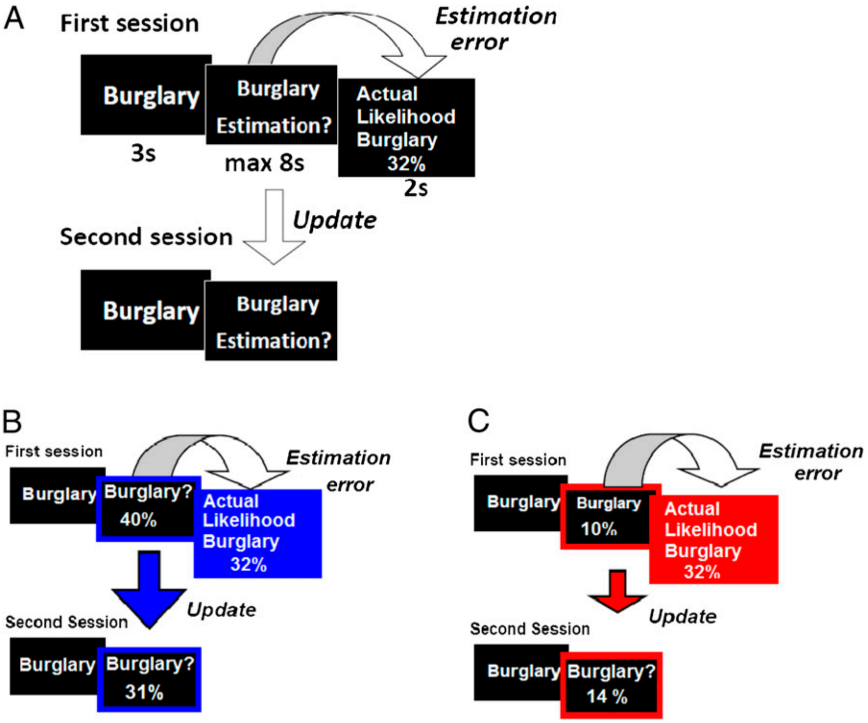

Fig. 1. Paradigm. (*A*) In each trial, participants were presented with a short description of 1 of 40 adverse events and asked to estimate how likely this event was to occur to them. They were then presented with the average probability of that event occurring. The second session was the same as the first except that the average probability of the event to occur was not presented. Examples of trials for which the participant's estimate was (*B*) higher or (*C*) lower than the average probability. Here, for illustration purposes only, the blue and red frames denote the participant's response (either an overestimation or underestimation, respectively), and the blue and red filled boxes denote information that calls for an adjustment in (*B*) a desirable or (*C*) an undesirable direction.

an opportunity to adopt a more optimistic outlook than for information that calls for a more pessimistic outlook (7). If children and adolescents exhibit an even stronger bias, they may be particularly compromised at integrating undesirable information, such as that provided by health and safety campaigns, into their beliefs.

Results

Learning from Good and Bad News. Our results revealed a marked asymmetry in how belief updating develops with age. Trials were divided into ones in which participants received desirable information (i.e., the probability presented was better than the estimate of their own probability; Fig. 1*B*) or undesirable information (i.e., the probability presented was worse; Fig. 1*C*). For each trial, an estimation error (estimation error = estimation − probability presented) and an update term were calculated (for desirable trials, update = first estimation − second estimation; for undesirable trials, update = second estimation − first estimation). For each participant we then calculated the correlation between estimation errors and subsequent updates across trials, which was the learning score.

Although participants learned from the information presented to them [mean Pearson correlation between an individual's estimation errors and update was significantly different from 0; $t(51) = 26.6$, $P < 0.001$], their ability to do so was differentially related to age as a function of valence. Specifically, the younger the participant, the worse they were at appropriately adjusting beliefs about the future in response to undesirable information, but the ability to adjust beliefs appropriately in response to desirable information did not alter significantly within the age range tested (Fig. 2*A*). This was evident from a positive correlation between age and learning scores in trials when participants received bad news ($r = 0.42$, $P < 0.005$) and no correlation between age and learning scores when receiving good news ($r = 0.21$, $P > 0.1$).

As demonstrated in Fig. 2 *B* and *C*, learning scores are calculated by quantifying the relationship, on a trial-by-trial basis for each participant, between an estimation error and subsequent update [as done previously (7)]. The resulting Pearson correlation coefficient makes up the learning score. This score, however, could be affected by a restricted range in the predictor and/or criterion (e.g., if the range of the predictor happens to be smaller in younger participants, this may artificially produce an association between age and learning scores). Thus, we adjusted the correlations reported above to control for the SD and mean of the two variables that make up the learning score: the initial estimation error (predictor) and subsequent update (criterion). After controlling for these factors, partial correlation coefficients confirmed a positive relationship between age and learning from bad news ($r = 0.45$, $P = 0.001$) and no correlation between age and learning from good news ($r = 0.11$, $P > 0.4$). The partial correlation between age and learning from bad news was significantly greater than the partial correlation between age and learning from good news (Steiger's $Z = 1.97$, $P < 0.05$).

Thus, reduced accurate learning from undesirable information about future likelihoods in late childhood/adolescence could not simply be explained by a general improvement of learning ability with age. Notably, the results cannot be explained by younger participants not understanding the task or having problems with calculating probabilities. This is because they performed exactly the same task in trials when they received good news and in trials when they received bad news. However, reduced accurate learning was more strongly associated with younger age for bad news trials than for good news trials.

For completeness, we also separately examined the correlation between age and the SD and mean of the criterion and predictor that make up the learning score (Fig. 3*A*). We found that the SD of the initial estimation error (i.e., the tendency to underestimate or overestimate one's likelihoods relative to the average population) did not correlate with age (for the undesirable trials, $r = 0.12$, $P > 0.3$; for the desirable trials, $r = -0.09$, $P > 0.5$). However, the SD of the update correlated negatively with age (for the undesirable trials, $r = -0.44$, $P = 0.001$; for the desirable trials, $r = -0.24$, $P = 0.09$). In other words, a positive correlation between age and learning from bad news was observed despite the range in update being smaller with increasing age. There was no significant relationship between age and the initial estimation

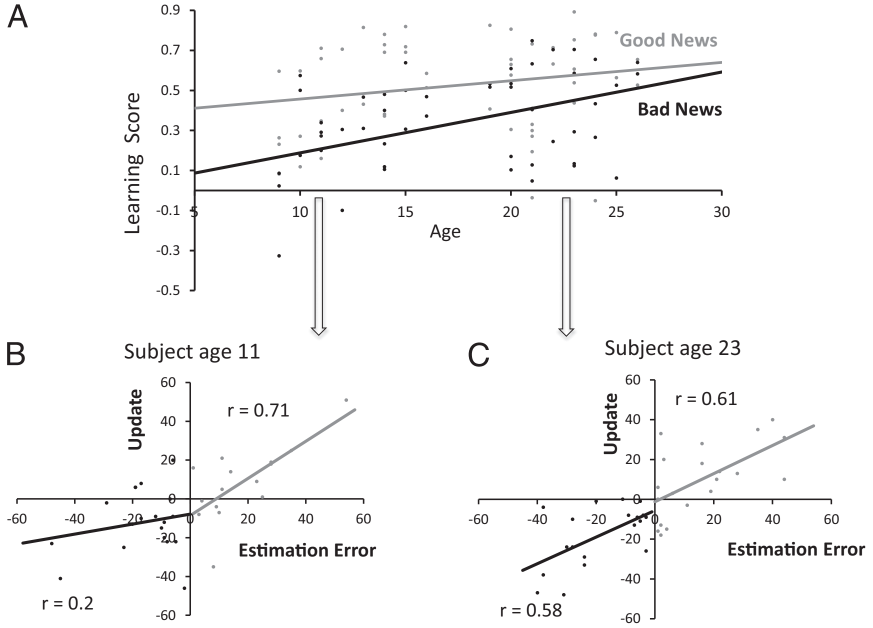

Fig. 2. Relationship between age and learning from good and bad news. (*A*) Correlation across subjects between age and learning from good and bad news (i.e., trials for which the information presented was better or worse than expected). (*B* and *C*) Learning is defined as the correlation between estimation errors and update across trials for each subject. Data from two subjects demonstrate this association for trials in which the subject received good news (thus, estimation errors are positive) and trials in which the subject received bad news (thus, estimation errors are negative). The slope of each line is the learning score of that subject. In this example, learning from bad news is worse than learning from good news in the younger participant but does not differ as much for the older participant.

error (for the undesirable trials, r = 0.048, P = 0.74; for the desirable trials, r = −0.25, P = 0.076), nor with the amount of accurate update (for the undesirable trials, r = −0.22, $P > 0.1$; for the desirable trials, r = 0.052, $P > 0.5$). These last findings suggest that although the magnitude of the estimates and the updating of those estimates are relatively stable in the age range tested, for younger subjects the amount of update was less precisely associated with the undesirable information received.

Valence-Dependent Effects on the Development of Learning Cannot Be Explained by Priors, Memory, Past Experience, Negativity, Familiarity, or Reaction Times. To examine whether the relationship between age and learning could be explained by any other factor, we tested for a correlation between age and all other variables recorded (Fig. 3*A*). We do not perform corrections for multiple comparisons because the aim of these analyses was to identify potential confounding factors; thus, by not using Bonferroni corrections, our analyses are more stringent.

Prior beliefs. An important question is whether subjects' prior beliefs of vulnerability differed across the ages tested. Consistent with previous findings (28), subjects perceived themselves as slightly less likely than their subpopulation to encounter harm. Specifically, their priors (i.e., pre-news estimation of likelihoods) tended to be slightly lower than the population statistics [$t(51)$ = 1.83, P = 0.07]. However, there was no significant relationship between age and priors [for the undesirable trials, r = −0.09, P = 0.5, and for the desirable trials, r = −0.17, P = 0.24; these correlations did not differ from each other (P = 0.7)]. In other words, younger individuals did not have more of an optimistic prior than older individuals. This is consistent with previous studies (29, 30) and runs contrary to the belief that children/adolescents perceive themselves as less vulnerable than adults do (31). Note that we only elicited one number for each stimulus that represents the subjects' prior. The shape of the distribution from which this number is taken is unknown.

Memory. Memory is a particularly important variable because memory for risk information is known to influence updating of beliefs regarding vulnerability (32), and inaccurate memory can often provide an explanation for judgment biases, as in the case of the hindsight bias (33). Furthermore, differences in subsequent memory can indicate attention differences during encoding.

Thus, we conducted multiple analyses to examine whether memory differed across age in a valence-dependent manner. Memory errors correlated with age for both desirable and undesirable trials (for the desirable trials, r = −0.40, P = 0.003; for the undesirable trials, r = −0.39, P = 0.004). However, these correlations did not differ from each other (Z = 0.08, $P > 0.5$).

Second, for each stimulus we calculated the correlation between memory and age and then excluded stimuli for which the two were significantly correlated (six stimuli were excluded). Removing these stimuli ensured that memory was no longer correlated with age (for the good news, r = 0.08, $P > 0.5$, and for the bad news, r = −0.05, $P > 0.5$; Z = 0.05, P > 0.5). We then repeated the main analysis without these stimuli and replicated our previous results: there was a positive correlation between age and learning from bad news (partial correlation controlling for magnitude and the variance of estimation error and update; r = 0.43, P = 0.003) but no correlation between age and learning from good news (partial correlation as above; r = 0.059, $P > 0.7$), and the two correlations were significantly different from each other (Z = 2.15, $P < 0.05$). This further suggests that the valence-dependent asymmetry in learning does not result from age differences in memory for the information presented.

Third, to directly examine the effect of memory for the presented information on updating of beliefs, we conducted an analysis relating the two. Specifically, for each subject we correlated across trials the accuracy of the subject's memory for the information given (i.e., memory errors; *Materials and Methods*) with their updating of risk estimates. As expected, the two were significantly correlated [$t(51)$ = 4.32, $P < 0.001$]. However, the relationship between them did not alter with age (for the good news trials, r = 0.004, P = 0.98; for the bad news trials, r = −0.059, P = 0.68), and the two correlations did not differ from each other (Z = 0.3, $P > 0.5$). Thus, although the relationship between memory and updating is significant, it does not alter with age in a valence-dependent manner.

Other variables. There was no significant correlation between age and past experience with the events (for the undesirable trials, r = 0.04, P = 0.77; for the desirable trials, r = 0.04, P = 0.78), nor with how negative the events were perceived to be [for the undesirable trials, r = 0.02, P = 0.89, and for the desirable trials, r = −0.14, P = 0.34; these last two correlations were, however, significantly different from each other ($P < 0.05$)], nor with

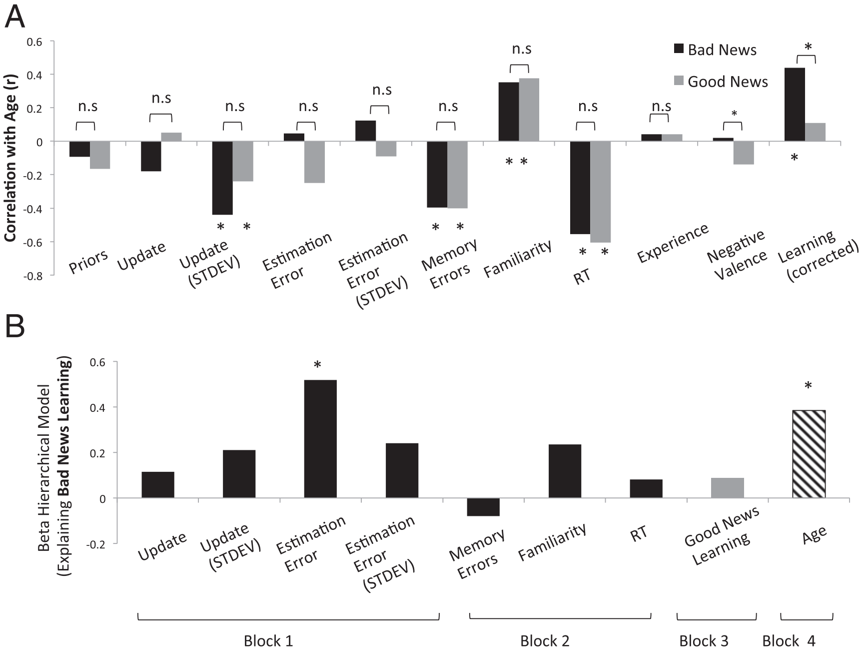

Fig. 3. Controlling for other variables. (*A*) Each bar represents the magnitude of the Pearson correlation coefficient (r) between that variable and age. Bars are plotted separately for good news trials and bad news trials. For learning, bars represent partial correlation coefficients that control for the magnitude and standard division across trials of the two factors that compose the learning score (i.e., corrected learning score); these differed significantly between good and bad news trials. Lower asterisks represent factors that showed a significant correlation with age for either good or bad news. (*B*) These factors were entered in a hierarchical regression model explaining learning from bad news. The four blocks correspond to the order of the variables that were entered. First, we controlled for the magnitude and SD of the two factors that compose the learning score (block 1); then, we controlled for any additional variable that correlated with age (block 2); next, we controlled for learning from good news (block 3); and finally, we introduced age (block 4). Betas are plotted after the final block is entered. Age significantly accounted for additional variance in learning from bad news over and above all other predictors. $^*P < 0.05$; n.s., no significant difference between bars.

depression scores ($r = -0.088$, $P = 0.56$). For both desirable and undesirable trials, age correlated significantly with reaction times (for the undesirable trials, $r = -0.56$, $P = 0.0001$; for the desirable trials, $r = -0.60$, $P = 0.0001$) and familiarity ratings (for the undesirable trials, $r = 0.35$, $P = 0.011$; for the desirable trials, $r = 0.38$, $P = 0.007$). The younger the participant, the slower the reaction times, and the less likely the participant was to be familiar with the stimuli presented.

Controlling for all relevant variables independently and together. To control for the variables that were found to be related to development, when examining the correlation between learning from bad news and age, we first controlled for each variable separately by conducting multiple partial correlation analyses. The results, displayed in Table S1, show that the relationship between learning from bad news and age remained significant in all these cases and also when controlling for gender.

Second, we conducted a hierarchical linear regression (Fig. 3*B*). Learning from undesirable information was entered as the dependent measure. In the first block of predictors (i.e., independent measures) we controlled for variables that were most likely to influence the learning score. The learning score was calculated as the relationship between update and estimation error, and thus, we first controlled for the magnitude and SD of these two variables (block 1). Next, we controlled for independent variables that were found to be related to age (memory, familiarity, and reaction time; block 2). In the third block we controlled for a variable that was not related to age but theoretically important for the investigation, learning from good news (this was entered to show again that the association between age and learning is specific to undesirable trials; block 3). Finally, after controlling for all these variables, we asked whether age can explain any additional variance in learning from bad news (block 4). We found that age accounted for additional variance in learning from undesirable information over and above all other predictors [$R^2_{\text{change}} = 0.064$, $F_{\text{change}\ (1,41)} = 4.63$, $P = 0.037$]. The results show that even after controlling for all these additional variables, the relationship between age and learning from undesirable information remains significant (Beta = 0.38, $P < 0.05$). Learning from desirable information did not provide additional explanatory variance in predicting learning from undesirable information [$R^2_{\text{change}} = 0.018$, $F_{\text{change}\ (1,41)} = 1.18$, $P = 0.28$]. This highlights the possibility that differential processes may be at play when adjusting beliefs in response to good and bad news.

Discussion

Our results reveal that the ability to alter beliefs of vulnerability appropriately in response to new information develops in a valence-dependent manner. Younger age was associated with inaccurate learning from information suggesting that the future is bleaker than expected, whereas the ability to learn from information suggesting that the future is brighter than expected was stable by comparison. Specifically, accurate updating in response to bad news improved linearly with development in the age range tested but stayed relatively constant in response to good news. This may be one factor that explains why adolescents can be especially resistant to warnings about danger: they have greater difficulty in learning from bad news than do adults.

The valence-dependent effects on learning were specific to integrating information regarding one's own risk, rather than learning about risk in the general population. In particular, memory for average risk in the population improved with age in a similar manner for desirable and undesirable information. Thus, although memory accuracy alters with age, which may also indicate changes in attention to the information provided, these changes were not valence-dependent. Furthermore, the tendency to overestimate or underestimate vulnerability did not alter with development, and the participants' prior beliefs were stable across age. The quantitative amount by which participants updated their beliefs in response to desirable or undesirable information did not alter with age, either. What did alter with age was the degree to which the amount of update was proportional to the error made, but only when the error called for an adjustment in a negative direction. Thus, when receiving negative information, younger individuals appeared to adjust their beliefs in a seemingly random manner, which was disproportionate to the data in front of them. This may result in inaccurate beliefs. For example, a young individual might assign enhanced weight for relatively low risk factors (i.e., noninvasive physical contact with an HIV-positive person) and reduced weight for high-risk factors (i.e., unprotected sex with a non-HIV carrier, which can result in unwanted pregnancy and sexually transmitted diseases), resulting in suboptimal decision making.

It is important to note that our results relate to the ability to adjust beliefs in response to information about possible future likelihoods. Such information is relatively unconstrained and open to interpretation. Slow development of learning from undesirable information may be restricted to such cases, and we do not know whether it will generalize to instances where outcomes are experienced (such as experienced punishments or losses). In fact, there is evidence that younger individuals exhibit better Pavlovian reversal learning from punishment than reward (34) but do better at adjusting behavior in response to positive feedback than negative feedback in an instrumental learning task (35–38). Furthermore, the likelihoods presented in the current study were related mostly to physical danger. It is possible that the ability to adjust beliefs regarding social outcomes, for example, has a different developmental trajectory, which might also differ between genders.

The relative tendency in late childhood/early adolescence not to adjust beliefs accurately in response to negative information may be related to the developmental of the frontal cortex. We have previously reported that people who show activity in the right inferior frontal gyrus, which closely tracks unexpected negative errors in this task, are more likely to update their beliefs accordingly (7). In contrast, adjustments in response to positive information in this task were supported not only by frontal regions but also by nonfrontal regions (7). In other words, a larger network of regions was sensitive to positive estimation errors, rendering learning from desirable information less susceptible to maturation of a particular region.

It has also been shown that the ability to adjust beliefs in this task is modulated by dopamine (39). The dopaminergic system undergoes substantial changes during the age range tested here (12, 40). These changes may have a significant effect on the process by which beliefs are adjusted in response to good and bad news, although the exact mechanism is unknown and likely to be complex. Cohen et al. (41), for example, have shown that neural prediction error signals to rewards in the striatum peak in adolescence. This may underlie the accurate computational use of positive estimation errors in adolescence observed here.

It is of interest to examine the developmental trajectory of learning from bad news beyond young adulthood into late adulthood. Both dopamine and frontal lobe function alter as we age (42), and valence-dependent effects on information processing and decision making have been demonstrated with aging. For example, it has been shown that younger adults have a negativity bias in attending to information relative to older adults (43) and are more inclined to pursue knowledge, whereas the elderly focus more on emotional satisfaction and meaning [known as the socioemotional selectivity theory (44)]. Such differences in goals and attention result in distinct effects of positive and negative mood on risk taking in younger and older adults (45). Thus, learning from bad news may follow an inverse U shape across the life span, peaking in young adulthood and declining with old age.

The asymmetry in the development of accurate learning from negative and positive information may be in accord with both the need for exploration in late childhood and adolescence and the heightened risk taking in that age group (11, 12). Indeed, exploratory behavior, which is important for acquiring new skills and independence, may involve taking a certain amount of risk (12). However, it can also result in aversive outcomes, which is why vast resources have been dedicated to educating adolescents about the consequences of their risky behavior (20–22). Our results show that this approach may be inherently limited because the ability to appropriately adjust beliefs about vulnerability in response to undesirable information develops disproportionally late between late childhood and adulthood. However, reframing the information to highlight positive outcomes of desired behaviors (e.g., the positive effect of reduced alcohol consumption on sports performance), rather than dangers of undesired ones, may have a larger impact (18, 46).

Materials and Methods

Participants. Fifty-nine volunteers (ages 9–26, 33 females) were recruited via a University College London Web site and the Science Museum in London. See Table S2 for age distribution. Participants gave informed consent and were compensated for their time. The study was approved by the University College London Research Ethics Committee.

Participants completed the Beck Depression Inventory (BDI) scale. It has been shown that depressed individuals tend to overestimate the likelihood of encountering negative events, such as those included in our study (47), in their lives. Thus, to control for this, we excluded subjects with BDI scores higher than 12. For participants younger than 18 y of age the BDI was revised such that one question regarding sexual behavior was excluded, and the score was adjusted accordingly. Four participants failed to complete the BDI. Seven participants with known health conditions and/or with BDI scores higher than 12 were excluded from the analysis (three adults and four adolescents), leaving 52 participants in the analysis.

Stimuli. Forty short descriptions of negative life events (e.g., passenger in a car accident or home burglary; Table S3) were presented in random order. All events were shown to all participants. Because of the developmental nature of the current study we used events that were associated with comparable likelihoods across age. This was verified to the best of our ability by online resources (as indicated below) and by a medical doctor [we note that medical doctors are themselves subject to biases and errors when reporting risk statistics (e.g., ref. 48)]. For each adverse event the average probability of that event occurring at least once was calculated from data compiled from online resources (including the Office for National Statistics and PubMed; see Table S4 for a list of online resources). Very rare or very common events were not included; all events probabilities lay between 10% and 70%. To ensure that the range of possible overestimation was equal to the range of possible underestimation, participants were told that the range of probabilities lay between 3% and 77%.

Procedure. The paradigm was adapted from our previous studies (7, 26, 39). Participants completed a practice session before beginning the main experiment.

In each trial, 1 of 40 adverse life events was presented for 3 s, and participants were asked to estimate how likely the event was to happen to them in the future. Participants had up to 8 s to respond. They were then presented with the actuarial frequency of the event in a demographically similar population for 2 s (Fig. 1).

In a second session, immediately after the first, participants were asked again to provide estimates of their likelihood of encountering the same events so that we could assess how they updated their estimate in response to the information presented. Participants then rated all stimuli on prior experience [for the question "Has this event happened to you before?" the responses ranged from 1 (never) to 6 (very often)], familiarity [for the question "Regardless if this event has happened to you before, how familiar do you feel it is to you from TV, friends, movies, and so on?" the responses ranged from 1 (not at all familiar) to 6 (very familiar)], and negativity [for the question "How negative would this event be for you?" the responses ranged from 1 (not negative at all) to 6 (very negative)].

To test memory for the information presented, subjects were asked at the end to provide the actual probability previously presented of each event. Memory errors were calculated as the absolute difference between the probability previously presented and the participants' recollection of that statistic:

$$\begin{aligned}\text{memory error} = |&\text{actual probability presented}\\ &- \text{recollection of probability presented}|.\end{aligned} \quad [1]$$

To ensure that the younger participants had a basic understanding of percentages, they were asked to complete two separate tasks before they began the experiment. First, they completed 10 trials in which they were presented with 100 cartoon faces, some of them yellow and some red. They were asked to identify the percentage of red faces out of three possible answers. Following this task the experimenter asked each participant to indicate which of two percentages was larger or smaller. Only one participant did not successfully complete all trials in both tasks and thus was not given the main task to complete.

Data Analysis. Trials were divided into ones in which participants received desirable information [i.e., the probability presented was lower than the estimate of their own probability (Fig. 1*B*)] or undesirable information [i.e., the probability presented was higher (Fig. 1*C*)]. For each trial an estimation error (estimation error = estimation − probability presented) and an update term were calculated (for desirable trials, update = first estimation − second estimation; for undesirable trials, update = second estimation − first estimation). Thus, positive updates indicate a change toward the actuarial frequency, and negative updates indicate a change away from the actuarial frequency.

For each participant we calculated the strength of the association between the estimation errors (predictor) and subsequent updates (criterion) across trials in which participants received desirable information and separately across trials for which they received undesirable information (see also ref. 7). This resulted in two Pearson coefficient scores for each participant. We then examined the relationship between these scores and age to determine whether this association is enhanced or reduced with development in the two conditions. Because small correlation coefficients can be a result of restriction in range in the predictor, the criterion, or both, we controlled for the SD and mean of the criterion and predictor by adding those as variables in a partial correlation analysis when examining the association between age and learning.

To determine whether the association between age and learning was different for undesirable and desirable trials, we then compared the two resulting partial correlations using Fisher's transformation, changing *r* to a *Z* score, and performed a Steiger's *Z* test, which examines if two correlated correlations within a single population are different from each other. The test compared the strength of the partial correlations between age and learning from desirable information with the partial correlation between age and learning from undesirable information, while taking into account the correlation between learning from desirable information and undesirable information.

Next, we controlled for any additional variables that were significantly correlated with age. To that end, we tested for a correlation between age and all stimuli ratings (ratings of familiarity, past experience, and negativity), memory, and reaction times. This was done separately for desirable and undesirable trials. We used a hierarchical linear regression model to test for the strength of the association between age and learning from undesirable information controlling for the other factors (this was not done for desirable trials because the correlation between age and learning from desirable trials was found not significant).

All *P* values are for two-tailed hypotheses unless otherwise stated.

ACKNOWLEDGMENTS. We thank Dr. Nicholas Wright for professional medical advice in developing stimuli. This work was supported by a Wellcome Trust Career Development Fellowship (to T.S.). S.-J.B. is funded by a Royal Society University Research Fellowship.

1. Tversky A, Kahneman D (1974) Judgement under uncertainty: Heuristics and biases. *Science* 185(4157):1124–1131.
2. Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *J Risk Uncertainty* 5:297–323.
3. Fischhoff B (2007) An early history of hindsight research. *Soc Cognit* 25:10–13.
4. Epley N, Gilovich T (2001) Putting adjustment back in the anchoring and adjustment heuristic: Differential processing of self-generated and experimenter-provided anchors. *Psychol Sci* 12(5):391–396.
5. Kahneman D, Tversky A (1973) On the psychology of prediction. *Psychol Rev* 80: 237–257.
6. Weinstein ND, Klein WM (1995) Resistance of personal risk perceptions to debiasing interventions. *Health Psychol* 14(2):132–140.
7. Sharot T, Korn CW, Dolan RJ (2011) How unrealistic optimism is maintained in the face of reality. *Nat Neurosci* 14(11):1475–1479.
8. Eil D, Rao JM (2011) The good news-bad news effect: Asymmetric processing of objective information about yourself. *Am Econ J Microecon* 3(2):114–138.
9. Kuhnen CM, Knutson B (2011) The influence of affect on beliefs, preferences, and financial decisions. *J Financ Quant Anal* 46(3):605–626.
10. Kuhnen CM (2012) Asymmetric learning from financial information. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2083173.
11. Steinberg L (2010) *Adolescence* (McGraw-Hill, New York), 9th Ed.
12. Blakemore S-J, Robbins TW (2012) Decision-making in the adolescent brain. *Nat Neurosci* 15(9):1184–1191.
13. Blum RW, Nelson-Mmari K (2004) The health of young people in a global context. *J Adolesc Health* 35(5):402–418.
14. Williams PG, Holmbeck GN, Greenley RN (2002) Adolescent health psychology. *J Consult Clin Psychol* 70(3):828–842.
15. Patton GC, et al. (2009) Global patterns of mortality in young people: A systematic analysis of population health data. *Lancet* 374(9693):881–892.
16. Viner RM, et al. (2012) Adolescence and the social determinants of health. *Lancet* 379(9826):1641–1652.
17. Reyna VF (2011) Across the life span. *Communicating Risks and Benefits: An Evidence-Based User's Guide*, eds Fischhoff B, Brewer NT, Downs JS (Food and Drug Admin, US Dep of Health and Hum Serv, Silver Spring, MD), pp 111–119.
18. Reyna VF, Farley F (2006) Risk and rationality in adolescent decision-making: Implications for theory, practice, and public policy. *Psychol Sci Public Interest* 7(1):1–44.
19. Reyna VF, Kiernan B (1994) Development of gist versus verbatim memory in sentence recognition: Effects of lexical familiarity, semantic content, encoding instructions, and retention interval. *Dev Psychol* 30:178–191.
20. Osius E, Rosenthal J (2009) *The National Research Council Institute of Medicine's Adolescent Health Services: Highlights and Considerations for State Health Policymakers* (Natl Acad for State Health Policy, Portland, ME).
21. Trenholm C, et al. (2007) *Impacts of Four Title V, Section 510 Abstinence Education Programs* (US Dep of Health and Hum Serv, Washington, DC).
22. Ennett ST, Tobler NS, Ringwalt CL, Flewelling RL (1994) How effective is drug abuse resistance education? A meta-analysis of Project DARE outcome evaluations. *Am J Public Health* 84(9):1394–1401.
23. Steinberg L (2008) A social neuroscience perspective on adolescent risk-taking. *Dev Rev* 28(1):78–106.
24. Luna B, et al. (2001) Maturation of widely distributed brain function subserves cognitive development. *Neuroimage* 13(5):786–793.
25. Giedd JN, et al. (1999) Brain development during childhood and adolescence: A longitudinal MRI study. *Nat Neurosci* 2(10):861–863.
26. Korn CW, Sharot T, Walter H, Heekeren HR, Dolan RJ (2013) Depression is related to an absence of optimistically biased belief updating about future life events. *Psychol Med*, 10.1017/S0033291713001074.
27. Sutton RS, Burto AG (1998) *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).
28. Weinstein ND (1980) Unrealistic optimism about future life events. *J Pers Soc Psychol* 39:806–820.
29. Millstein SG (1993) Perceptual, attributional, and affective process in perceptions of vulnerability through the life span. *Adolescent Risk Taking*, eds Bell NJ, Bell RW (Sage, Thousand Oaks, CA), pp 55–65.
30. Quadrel MJ, Fischhoff B, Davis W (1993) Adolescent (in)vulnerability. *Am Psychol* 48(2):102–116.
31. Elkind D (1967) Egocentrism in adolescence. *Child Dev* 38(4):1025–1034.
32. Fischhoff B, MacGregor D (1983) Judged lethality: How much people seem to know depends upon how they are asked. *Risk Anal* 3:229–236.
33. Reyna VF (2005) Fuzzy-trace theory, judgment, and decision-making: A dual-processes approach. *Human Learning and Memory: Advances in Theory and Application: The 4th Tsukuba International Conference on Memory*, eds Izawa C, Ohta N (Lawrence Erlbaum Assoc, Mahwah, NJ), pp 239–256.
34. van der Schaaf ME, Warmerdam E, Crone EA, Cools R (2011) Distinct linear and non-linear trajectories of reward and punishment reversal learning during development: Relevance for dopamine's role in adolescent decision making. *Dev Cogn Neurosci* 1(4): 578–590.
35. Crone EA, Ridderinkhof KR, Worm M, Somsen RJ, van der Molen MW (2004) Switching between spatial stimulus-response mappings: A developmental study of cognitive flexibility. *Dev Sci* 7(4):443–455.
36. Crone EA, Zanolie K, Van Leijenhorst L, Westenberg PM, Rombouts SA (2008) Neural mechanisms supporting flexible performance adjustment during development. *Cogn Affect Behav Neurosci* 8(2):165–177.
37. Huizinga M, Dolan CV, van der Molen MW (2006) Age-related change in executive function: Developmental trends and a latent variable analysis. *Neuropsychologia* 44(11):2017–2036.
38. Van Duijvenvoorde ACK, Zanolie K, Rombouts SA, Raijmakers MEJ, Crone EA (2008) Evaluating the negative or valuing the positive? Neural mechanisms supporting feedback-based learning across development. *J Neurosci* 28(38):9495–9503.
39. Sharot T, Guitart-Masip M, Korn CW, Chowdhury R, Dolan RJ (2012) How dopamine enhances an optimism bias in humans. *Curr Biol* 22(16):1477–1481.
40. Casey BJ, Jones RM, Hare TA (2008) The adolescent brain. *Ann N Y Acad Sci* 1124: 111–126.
41. Cohen JR, et al. (2010) A unique adolescent response to reward prediction errors. *Nat Neurosci* 13(6):669–671.
42. Chowdhury R, et al. (2013) Dopamine restores reward prediction errors in old age. *Nat Neurosci* 16(5):648–653.
43. Rozin P, Royzman E (2001) Negativity bias, negativity dominance, and contagion. *Pers Soc Psychol Rev* 5:296–320.
44. Carstensen LL, Isaacowitz DM, Charles ST (1999) Taking time seriously. A theory of socioemotional selectivity. *Am Psychol* 54(3):165–181.
45. Chou KL, Lee TMC, Ho AHY (2007) Does mood state change risk taking tendency in older adults? *Psychol Aging* 22(2):310–318.
46. Gibbons FX, Houlihan AE, Gerrard M (2009) Reason and reaction: The utility of a dual-focus, dual-processing perspective on promotion and prevention of adolescent health risk behaviour. *Br J Health Psychol* 14(Pt 2):231–248.
47. Strunk DR, Lopez HL, DeRubeis RJ (2006) Depressive symptoms are associated with unrealistic negative predictions of future life events. *Behav Res Ther* 44(6):861–882.
48. Reyna VF, Adam MB (2003) Fuzzy-trace theory, risk communication, and product labeling in sexually transmitted diseases. *Risk Anal* 23(2):325–342.