

Striatal structure and function predict individual biases in learning to avoid pain

Eran Eldar^{a,b,1}, Tobias U. Hauser^{a,b}, Peter Dayan^{c,2}, and Raymond J. Dolan^{a,b,2}

^aWellcome Trust Centre for Neuroimaging, University College London, London WC1N 3BG, United Kingdom; ^bMax Planck University College London Centre for Computational Psychiatry and Ageing Research, London WC1B 5EH, United Kingdom; and ^cGatsby Computational Neuroscience Unit, University College London, London W1T 4JG, United Kingdom

Edited by Thomas D. Albright, The Salk Institute for Biological Studies, La Jolla, CA, and approved March 15, 2016 (received for review October 6, 2015)

Pain is an elemental inducer of avoidance. Here, we demonstrate that people differ in how they learn to avoid pain, with some individuals refraining from actions that resulted in painful outcomes, whereas others favor actions that helped prevent pain. These individual biases were best explained by differences in learning from outcome prediction errors and were associated with distinct forms of striatal responses to painful outcomes. Specifically, striatal responses to pain were modulated in a manner consistent with an aversive prediction error in individuals who learned predominantly from pain, whereas in individuals who learned predominantly from success in preventing pain, modulation was consistent with an appetitive prediction error. In contrast, striatal responses to success in preventing pain were consistent with an appetitive prediction error in both groups. Furthermore, variation in striatal structure, encompassing the region where pain prediction errors were expressed, predicted participants' predominant mode of learning, suggesting the observed learning biases may reflect stable individual traits. These results reveal functional and structural neural components underlying individual differences in avoidance learning, which may be important contributors to psychiatric disorders involving pathological harm avoidance behavior.

avoidance learning | pain | individual differences | striatum | prediction errors

Pain conveys vital feedback on our actions, informing us whether an action compromises our safety and should be avoided. However, learning what to avoid doing, rather than what to do, could lead to maladaptive passive risk-avoidant behavior. For instance, when learning to ski, an overreaction to a painful fall could render a person overly cautious and hinder progress in skill acquisition. Likewise, failed investments might lead to an overconservative passive financial strategy, whereas social rejection might engender reclusive behavior. In some individuals, such as those with avoidant personality disorders, refraining too much from potentially harmful actions can manifest as a stable personality trait (1).

However, an opposite tendency, to learn predominantly from successful actions that helped avoid harm, might lead to maladaptive active behavior. Thus, in soccer, sporadic success in preventing goals by diving to the left or right before seeing where a penalty kick is heading is sufficient for goalkeepers to overwhelmingly prefer this suboptimal active strategy, when in fact the optimal strategy is to passively stay put (2). In the extreme, excessive repetitive activity so as to avoid harm may constitute compulsivity, a debilitating feature of obsessive-compulsive disorder (3).

Complementing previous studies of learning about abstract outcomes (4, 5), we investigated individual biases in learning to avoid pain. To this end, we used a novel gambling task that probes how participants adjust their choices in response to painful electrical shocks and, additionally, how they adjust their choices in response to success in preventing shocks. In line with studies of reward learning (6–9), learning in our task was best explained as driven by an outcome prediction error that reflects

the difference between expected and actual outcomes. Consistent with the expression of such a teaching signal, blood-oxygen level-dependent (BOLD) responses to outcomes in the striatum were modulated by expectation. However, striatal response to shocks were qualitatively different in negative learners (i.e., those who predominantly learned from shocks) compared with positive learners (i.e., those who predominantly learned from success in avoiding shocks). Specifically, striatal activity was consistent with an aversive prediction error signal in negative learners and with an appetitive prediction error signal in positive learners. The degree to which a participant tended to learn from success in avoiding than experiencing shocks was predicted by the structure of a participants' striatum, specifically by higher gray matter density where the response to shocks was consistent with a prediction error signal.

Results

Individual Biases in Learning from Pain and Its Prevention. To test for individual differences in learning to avoid pain, we tasked 41 participants to play a card game in which their goal was to minimize the number of painful electrical shocks they might receive. Participants could avoid shocks by gambling that the number they were about to draw will be higher than the number the computer had drawn. An unsuccessful gamble resulted in shock and a successful gamble led to its avoidance. Alternatively, participants could always decline the gamble and opt for a fixed 50% known probability of receiving a shock (Fig. 1A). Importantly, participants played with three different decks of cards and

Significance

Our ability to learn how to avoid harm is critical for maintaining physical and mental health. However, excessive harm avoidance can be maladaptive, as evident in psychiatric disorders such as avoidant personality disorder and obsessive-compulsive disorder. We therefore investigated the neural factors underlying individual imbalances in harm avoidance behavior. Our findings show that such imbalances can be predicted by the function and structure of an individual's striatum, a brain region that is critical for goal-directed decisionmaking. Moreover, the neural signals expressed in this region revealed key processes through which individuals learn to avoid harm. These findings highlight a neural basis for imbalanced harm avoidance behavior, extreme forms of which may contribute to psychiatric pathology.

Author contributions: E.E., P.D., and R.J.D. designed research; E.E. and T.U.H. performed research; E.E. analyzed data; and E.E., T.U.H., P.D., and R.J.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: e.eldar@ucl.ac.uk.

²P.D. and R.J.D. contributed equally to this work.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1519829113/-DCSupplemental.

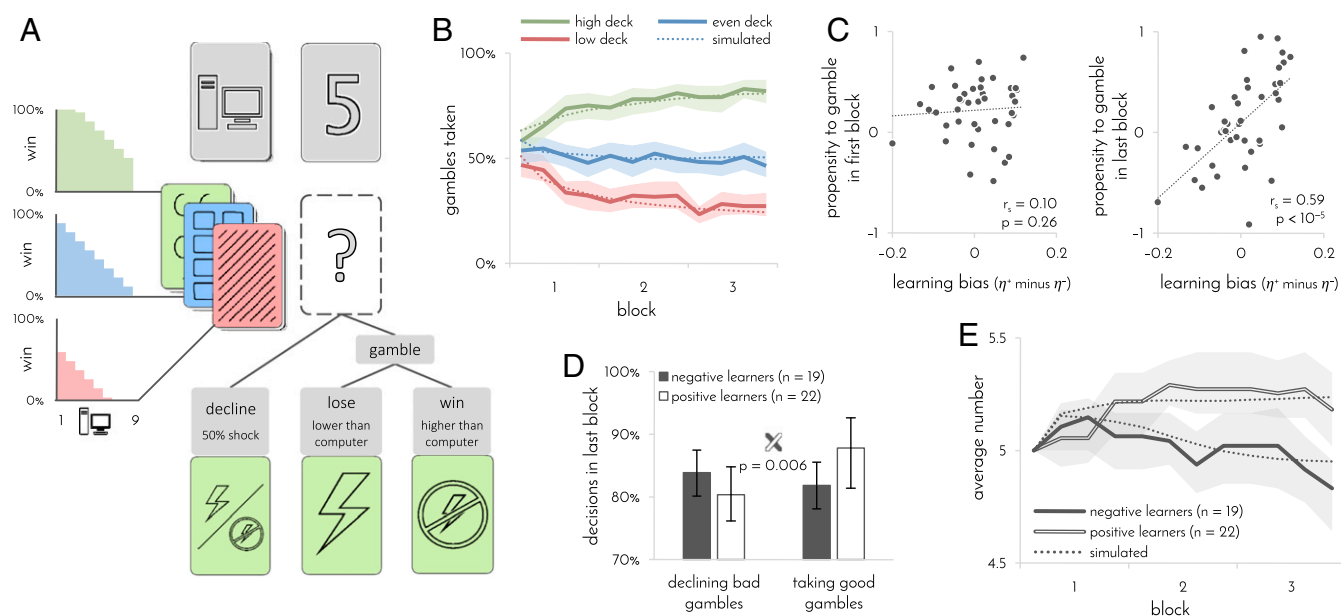


Fig. 1. Experimental design and learning performance. $n = 41$ participants. (A) Experimental design. On each trial, participants were presented with one of three possible decks and a number between 1 and 9 drawn by the computer. If participants decided to gamble, a shock was delivered only if the number that they drew was lower than the computer's number. Participants were only informed whether they won or lost the gamble, not which number they drew. Participants had to learn by trial and error how likely gambles were to succeed with each of the three decks. One deck contained a uniform distribution of numbers between 1 and 9 (even deck), one deck contained more 1's (low deck), making gambles 30% less likely to succeed, and one deck contained more 9's (high deck), making gambles 30% more likely to succeed. Opting to decline the gamble resulted in a 50% probability of shock regardless of which numbers were drawn by the computer. (B) Gambles taken with each deck as a function of time. Percentages were computed separately for each set of 15 contiguous trials (4 sets/60 trials per block). (C) Participants' propensity to gamble in first (Left) and last (Right) blocks of trials as a function of learning bias. Propensity to gamble was computed by regressing out the effects of deck and computer's number on participant's choices using logistic regression. The numbers 1 and -1 correspond to always and never gamble, respectively. Learning bias was inferred from a participant's choices using the learning model. (D) Proportion of bad gambles that were declined and good gambles that were taken in last block of trials as a function of learning bias. Participants with a positive learning bias (positive learners) declined fewer bad gambles and took more good gambles than participants with a negative learning bias (negative learners). Gambles were defined as good or bad based on probability of winning (good: $>50\%$; bad $<50\%$). Error bars: 95% bootstrap CI. (E) Average number drawn by computer as a function of time and learning bias. To maintain participants at a 50% gambling rate, numbers increased for positive learners and decreased for negative learners. In B and E, dotted line indicates simulated task performance of learning model. Shaded areas: 95% bootstrap CI.

had to learn by trial and error how likely a gamble was to be successful with each deck.

In principle, participants can acquire information about the decks from both successful and unsuccessful gambles. Indeed, we observed from their behavior that as they gained experience with the three decks, their willingness to gamble with each deck differed (Fig. 1B). A more in-depth analysis indicated that they did not learn from the two types of outcomes to the same degree. Thus, the learning algorithm that best explained participants' choices included two different learning rates, one for learning from shock outcomes (η^-) and one for learning from no-shock outcomes (η^+ ; log Bayes factor compared with algorithm with a single learning rate = 27.7). In this algorithm, the two learning rates determine the degree to which the two types of outcomes impact on subsequent expectations of gambling with each of the decks, and these expectations in combination with the numbers drawn by the computer determine whether subsequent gambles are taken or declined. The algorithm also accounts for each participant's baseline propensity to take gambles. Learning in the favored algorithm is weighted by associability (10, 11), and the choices made for each deck tend to persist (7) (see *SI Appendix* for details of all learning algorithms and a validation of the model comparison procedure).

Further analysis showed that the difference between the two learning rates captured significant interindividual variation (log Bayes factor of algorithm with two learning rates per participant compared with algorithm with one average learning rate per participant and a group parameter for difference between the learning rates = 20.1). On this basis, we next computed each

participant's learning bias as the difference between the two learning rates that best fitted the participant's choices (η^+ minus η^-). This bias reflects the degree to which a participant learned what gambles to take (because they resulted in no-shock outcomes) rather than what gambles to avoid (because they resulted in shocks). A positive learning bias ($\eta^+ > \eta^-$) entails that a propensity to gamble will emerge as the participant is learning from outcomes, whereas a negative learning bias ($\eta^+ < \eta^-$) should engender a propensity to decline a gamble (9). A comparison of participant's raw propensity to gamble at the beginning and end of the experiment confirmed this expectation (Fig. 1C). Consequently, by the end of the experiment, participants with a positive learning bias came to take more good gambles, but also decline fewer bad gambles than participants with a negative learning bias (Fig. 1D).

To ensure that all participants nevertheless gambled at a similar rate and, thus, received roughly equal amounts of information about the decks, we increased or decreased the numbers the computer drew according to the participants' own gambling rate, while ensuring that the three decks are always matched against similar computer's numbers. Thus, because of their propensity to gamble, positive learners ended up playing against high numbers, whereas negative learners ended up playing against low numbers [Fig. 1E; mean difference 0.36, confidence interval (CI) 0.15–0.6, $P = 0.001$, bootstrap test; as a result, positive learners received a slightly higher number of shocks (mean 74.5, CI 72.7–76.5) than negative learners (mean 72.0, CI 70.3–73.3)]. In sum, participants who learned more from painful outcomes developed a propensity to avoid gambling, whereas

participants who learned more from success in preventing pain developed a propensity to gamble.

Striatal Response to Pain Underlies Learning Bias. Behaviorally, learning in our task was best explained as driven by outcome prediction errors (*SI Appendix, Figs. S1 and S2*). Such prediction errors are known to be expressed in the striatum (7–9, 12, 13), and, thus, we focused on examining participants' striatal activity in response to task outcomes. BOLD responses in the striatum (and elsewhere) were overwhelmingly stronger in response to shock than to no-shock outcomes (*SI Appendix, Fig. S3*), likely reflecting sensory and affective processing of the shocks. Thus, to specifically test for prediction error signals, we examined the response to each type of outcome (shock and no-shock) separately, testing whether these responses were modulated by expectation.

We found that BOLD responses to no-shock outcomes in the head of the caudate nucleus were negatively modulated by expectation (derived by using the learning algorithm that best fitted participants' behavior; Fig. 2*A*), consistent with an appetitive prediction error. This signal was not explained away by the effect of associability [$P < 0.05$, familywise error (FWE) corrected]. In contrast, when accounting for prediction errors, we did not find evidence for an associability signal in the striatum or elsewhere ($P > 0.05$ FWE corrected). Moreover, when decomposing the prediction error (PE) signal into the two components that affect a participant's expectation, the one reflecting previous experience with the current deck and the other reflecting the number drawn by the computer, we found a significant effect within the identified striatal region in response to both PE components {Fig. 2*B*; peak Montreal Neurological Institute (MNI) coordinates: deck expectation [$-10\ 6\ 6$], $t_{40} = -4.0$, $P = 0.04$; number [$-8\ 6\ 2$], $t_{40} = 4.2$, $P = 0.03$; P values are FWE corrected across the striatum}.

Whereas responses to no-shock outcomes were similarly modulated by expectations in positive and negative learners, the response to shock outcomes was qualitatively different in the two groups. No prediction error signal was evident in striatal responses to shocks across both groups of participants ($P > 0.05$, FWE small-volume corrected). However, examining the two groups separately revealed that the BOLD response to shocks was positively modulated by expectation in negative learners and negatively modulated by expectation in positive learners (Fig. 2*C*; responses measured in the same area where the response to no-shock outcomes was modulated by expectations; correlation between learning bias and average general linear model (GLM) coefficient r values = 0.4, $P = 0.005$, permutation test; partial correlation accounting for number of shocks received: $\rho_s = 0.38$, $P = 0.007$, permutation test; a similar difference between positive and negative learners was evident in the caudate in a voxel-based analysis across the striatum—peak MNI coordinates [$-10\ -2\ 18$], $P = 0.03$ small-volume corrected, permutation test—as well as when averaging across the whole caudate, $P = 0.05$, Bonferroni corrected for the three subregions of the striatum). In other words, in response to shocks, striatal activity in negative learners was consistent with an aversively signed prediction error, as seen in previous studies of pain prediction (14, 15). By contrast, in positive learners, activity was consistent with an appetitively signed prediction error, as in previous work on pain-relief prediction (16). Importantly, this difference between the groups was not associated with a difference in the mean response to shock compared with no-shock outcomes ($P = 0.60$, bootstrap test, GLM; *SI Appendix, Fig. S3 C and D*). These results suggest that, across both types of outcomes, positive learners express an appetitively signed prediction error, whereas negative learners express the absolute value of the prediction error.

If striatal prediction errors indeed drive avoidance learning, then the modulation of striatal responses to outcomes by expectation should predict whether a participant subsequently developed a propensity to gamble. Modulation consistent with an appetitive prediction error (signaling which gambles should

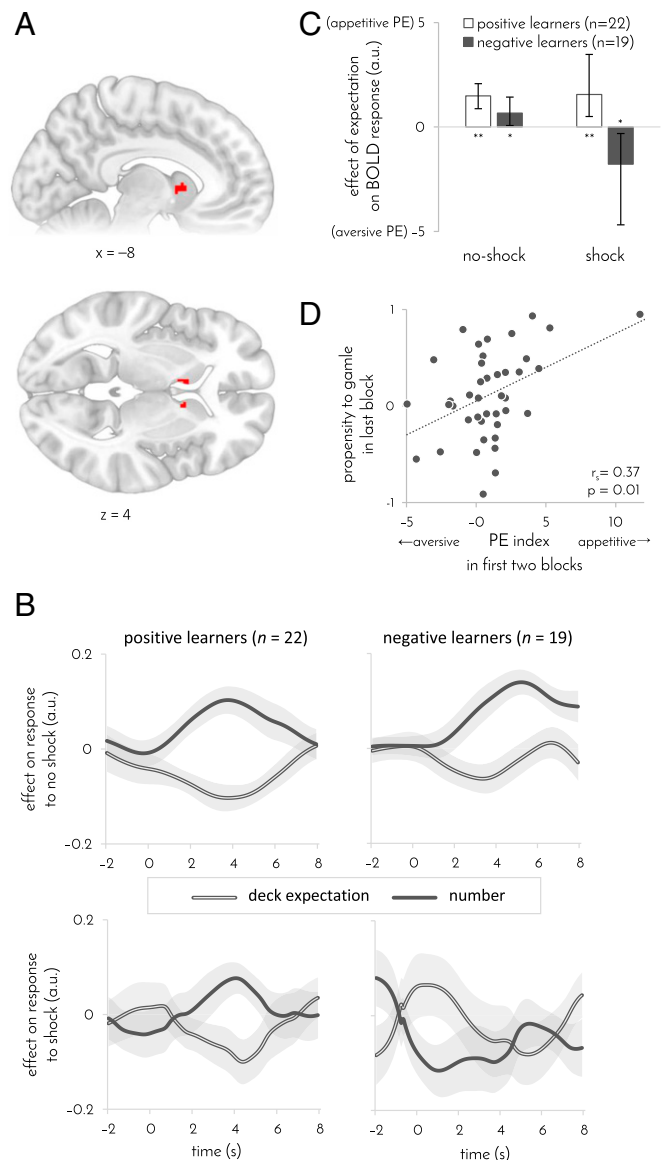


Fig. 2. PE signaling in the striatum. $n = 41$ participants. (A) Striatal area where BOLD response to no-shock outcomes was modulated by expectations. Map is masked and FWE corrected for volume of striatum ($P < 0.05$, GLM). Extent: 22 voxels. Peak MNI coordinates: [$-6\ 10\ 2$] $t_{40} = 4.5$, corrected $P = 0.01$; [$10\ 8\ 4$] $t_{40} = 4.7$, corrected $P = 0.007$. x and z denote Montreal Neurological Institute (MNI) coordinates. (B) The two components of the prediction error. Effects of deck-based expectation and the number drawn by the computer on BOLD response to shock (*Bottom*) and no shock (*Top*) outcomes in the striatal ROI identified in A. Results are shown separately for positive (*Left*) and negative (*Right*) learners. A higher deck expectation and a lower number indicate lower expectation of a shock, and thus, an appetitive PE is consistent with a rise in the effect of number and a dip in the effect of deck expectation on the response to the outcomes. The converse pattern is consistent with an aversive PE. Time 0 indicates outcome onset. Shaded area: SEM. (C) Effect of expectations on BOLD response to no-shock and shock outcomes as a function learning bias. Positive values indicate an effect that is consistent with an appetitive PE. * $P < 0.05$, ** $P < 0.005$, error bars: 95% bootstrap CI. (D) Propensity to gamble in the last task block as a function of a participant's PE index, computed as the average effect of expectation on striatal response across both types of outcomes. An appetitive PE index predicted subsequent propensity to gamble (computed as in Fig. 1C). GLM coefficients in C and D were taken from the striatal area identified in A.

area where BOLD responses to shocks were modulated by expectation (Fig. 4C; regression of individual expectation GLM coefficients on group gray matter-density predictive coefficients: mean 0.44, CI 0.22–0.77, $P = 10^{-5}$, bootstrap test). That is, a more positive learning bias was predicted by higher gray matter density where responses to shocks were consistent with a (appetitive or aversive) prediction error signal, and lower gray matter density in areas that showed no such signal.

Amygdala Specifically Involved in Learning from Pain. A large body of animal and human work implicates the amygdala and periaqueductal gray (PAG) in learning from aversive outcomes and, in particular, in generating aversive prediction errors (19–24). BOLD responses to outcomes in these areas suggests both areas were involved in learning in our task, albeit in different ways. Responses to shock and no-shock outcomes in the PAG were modulated by expectations in the same way as in our striatal ROI (SI Appendix, Fig. S3E; although we note that the ability of fMRI to discern PAG signals from neighboring structures is limited; ref. 25). In contrast, the amygdala showed no significant modulation in the response to no-shock outcomes, but its response to shocks was consistent with an aversive prediction error across the whole group (extent: 1 voxel, MNI coordinates [28 –6 –18], $t_{40} = 2.8$, $P = 0.04$ corrected for the volume of the amygdala). This latter effect was particularly pronounced in participants with a negative learning bias (Fig. 5A), mirroring the modulation of striatal and PAG activity in these same participants.

Interestingly, a positive learning bias was associated with a stronger response to shock than to no-shock outcomes in right posterior insula (extent: three voxels, peak MNI coordinates [46 –32 20], $P = 0.04$ FWE-corrected for all voxels that significantly responded to shocks, shown in SI Appendix, Fig. S3A). The insula is thought to feed information about salient, painful, and aversive events to amygdala and striatum (26–28). Therefore, we next examined functional connectivity between this area and the striatal and amygdala regions in which response to outcomes was modulated by expectation. Positive learners showed significant functional connectivity between the insula and striatal regions, whereas negative learners showed significant functional connectivity between the insula and amygdala regions (correlation of learning bias with difference between striatal and amygdala functional connectivity: $r_s = 0.42$, $P = 0.005$; Fig. 5B). Taken together, these results are consistent with previous suggestions that the amygdala is exclusively involved in learning from increases in aversive outcomes, whereas the striatum and PAG also partake in learning from decreases in aversive outcomes (21, 23, 29).

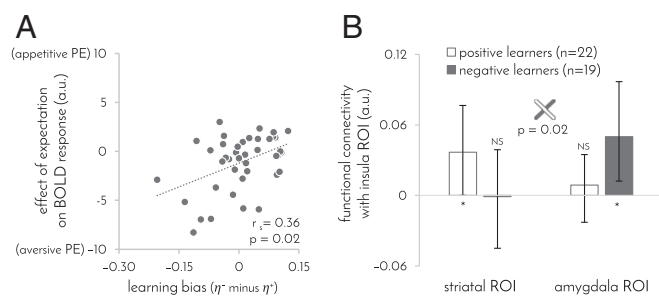


Fig. 5. Individual learning biases outside of the striatum. (A) Effect of expectation on BOLD response to shocks as a function of learning bias in the amygdala region where this effect was significant (MNI coordinates [28 –6 –18]). Responses were most consistent with an aversive prediction error in participants who mostly learned from shock outcomes. $n = 41$ participants. (B) Functional connectivity between striatal (Fig. 2A) and amygdala (A) ROIs, and the insula area in which the response to shocks correlated with learning bias (MNI coordinates [46 –32 20]). Error bars: 95% bootstrap CI, $*P \leq 0.05$, NS: $P > 0.1$.

Discussion

We demonstrate that the two ways through which one can learn to avoid harm are used to different degrees by different individuals. In negative learners—those who primarily learned from being shocked and, thus, developed a propensity to avoid gambles—dorsal striatal and amygdalar responses to shocks were consistent with an aversive prediction error. In contrast, in positive learners—those who primarily learned from their success in preventing shocks and, thus, developed a propensity to gamble—the dorsal striatum’s response to shocks was consistent with an appetitive prediction error. This difference in striatal responses to outcomes anticipated observed differences in learned behavior, and in the neural responses to stimuli about which participants were learning. Participants’ learning bias was also predicted by the structure of their striatum, indicating that learning biases in our task reflected, at least in part, stable individual traits. Together, the findings reveal neural underpinnings of an elementary behavioral trait that predicts whether an individual learns predominantly what to do to prevent harm or what to avoid doing.

A multitude of animal and human studies implicate the striatum in learning about aversive outcomes (12, 13, 30–34). Striatal areas, including the caudate region identified in our study, have been shown to represent prediction error signals in both classical conditioning (14–16) and instrumental pain avoidance learning (35, 36). However, individual differences in the expression of these prediction error signals have not been studied to our knowledge. The present study was designed to assess the degree to which participants learn what actions have painful outcomes compared with what actions help avoid pain. The latter type of learning, defined in the animal literature as active avoidance learning, is particularly interesting, because it specifies that the very absence of a shock is reinforcing (37). Indeed, our results show that striatal response to outcomes in participants who were biased in favor of active avoidance learning mimicked striatal responses typically seen in studies of reward learning, with no-shock outcomes cast as reward and shock outcomes as reward omission. In contrast, in participants biased in favor of passive avoidance learning (i.e., learning what gambles should be avoided), striatal response to painful outcomes was consistent with an aversive prediction error, as seen in fear conditioning (14).

Thus, our results show that striatal response to pain is consistent with an appetitive prediction error in some individuals and with an aversive prediction error in others. In contrast, striatal responses to successful prevention of pain seem broadly consistent with an appetitive prediction error. Put another way, while some individuals represent an appetitively signed prediction error in response to both types of outcomes, others represent an unsigned prediction error (38) (aversively signed in response to pain and appetitively signed in response to pain prevention). That said, we note that in our experiment, the response to pain was generally stronger than the response to pain prevention, which is inconsistent with a purely appetitive prediction error, although this stronger response could reflect affective and sensory processing of the shocks.

Our findings concur with a view of the striatum as involved in processing both appetitive and aversive outcomes (26, 39). The amygdala, in contrast, was involved in our study solely in processing aversive outcomes, but not their omission. This involvement of the amygdala was most evident in participants who primarily learned from shock outcomes, and underscoring by greater functional connectivity with the insula, a region with an established role in processing salient and aversive outcomes (40). These results strongly concur with previous work, indicating that learning from aversive outcomes engages the amygdala, whereas learning from success in avoiding aversive outcomes involves inhibition of the amygdala and activation of the striatum (29, 41–43). That said, we note that animal studies that were able to examine subregions of the striatum and amygdala with greater spatial

resolution reveal a more complex picture (44, 45). Finally, whereas functional connectivity between medial prefrontal cortex, amygdala, and striatum has been shown to mediate avoidance learning (43), it did not mediate learning biases in our task, suggesting that such connectivity might be equally involved in learning from increases and decreases in aversive outcomes.

Several reports have linked variations in the structure of the striatum to individual differences in healthy and pathological decision-making behavior (46, 47) and to the expression of certain pain disorders (48, 49). Of particular relevance to the present study is the observation that obsessive-compulsive disorder (OCD), which features compulsive harm-avoidance behavior, is associated with higher gray matter density in the putamen (50). In our study, higher gray matter density in the putamen (and lower gray matter density in the caudate) predicted better learning from shocks and poorer learning from success in avoiding shocks. It is possible that such insensitivity to safety signals might engender excessively persistent harm avoidance behaviors, which in healthy individuals normally terminate when safety is attained. Thus, this finding raises the interesting possibility that failure to adjust to success in harm avoidance may contribute to compulsivity in OCD.

- American Psychiatric Association (2013) *Diagnostic and Statistical Manual of Mental Disorders* (Am Psychiatr Assoc, Washington, DC), 5th Ed.
- Bar-Eli M, Azar OH, Ritov I, Keidar-Levin Y, Schein G (2007) Action bias among elite soccer goalkeepers: The case of penalty kicks. *J Econ Psychol* 28:606–621.
- Abramowitz JS, Taylor S, McKay D (2009) Obsessive-compulsive disorder. *Lancet* 374(9688):491–499.
- Frank MJ, Seeberger LC, O'reilly RC (2004) By carrot or by stick: Cognitive reinforcement learning in parkinsonism. *Science* 306(5703):1940–1943.
- Frank MJ, Moustafa AA, Haughey HM, Curran T, Hutchison KE (2007) Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proc Natl Acad Sci USA* 104(41):16311–16316.
- Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD (2006) Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* 442(7106):1042–1045.
- Schönberg T, Daw ND, Joel D, O'Doherty JP (2007) Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *J Neurosci* 27(47):12860–12867.
- Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J Neurosci* 28(22):5623–5630.
- Niv Y, Edlund JA, Dayan P, O'Doherty JP (2012) Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *J Neurosci* 32(2):551–562.
- Li J, Schiller D, Schoenbaum G, Phelps EA, Daw ND (2011) Differential roles of human striatum and amygdala in associative learning. *Nat Neurosci* 14(10):1250–1252.
- Boll S, Gamer M, Gluth S, Finsterbusch J, Büchel C (2013) Separate amygdala subregions signal surprise and predictiveness during associative fear learning in humans. *Eur J Neurosci* 37(5):758–767.
- Cohen JY, Haesler S, Vong L, Lowell BB, Uchida N (2012) Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature* 482(7383):85–88.
- Delgado MR, Li J, Schiller D, Phelps EA (2008) The role of the striatum in aversive learning and aversive prediction errors. *Philos Trans R Soc Lond B Biol Sci* 363(1511):3787–3800.
- Seymour B, et al. (2004) Temporal difference models describe higher-order learning in humans. *Nature* 429(6992):664–667.
- Schiller D, Levy I, Niv Y, LeDoux JE, Phelps EA (2008) From fear to safety and back: Reversal of fear in the human brain. *J Neurosci* 28(45):11517–11525.
- Seymour B, et al. (2005) Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nat Neurosci* 8(9):1234–1240.
- Ashburner J, Friston KJ (2000) Voxel-based morphometry—the methods. *Neuroimage* 11(6 Pt 1):805–821.
- Helms G, Draganski B, Frackowiak R, Ashburner J, Weiskopf N (2009) Improved segmentation of deep brain grey matter structures using magnetization transfer (MT) parameter maps. *Neuroimage* 47(1):194–198.
- Phillips RG, LeDoux JE (1992) Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behav Neurosci* 106(2):274–285.
- Kim JJ, Rison RA, Fanselow MS (1993) Effects of amygdala, hippocampus, and periaqueductal gray lesions on short- and long-term contextual fear. *Behav Neurosci* 107(6):1093–1098.
- Cole S, McNally GP (2008) Complementary roles for amygdala and periaqueductal gray in temporal-difference fear learning. *Learn Mem* 16(1):1–7.
- Johansen JP, Tarpley JW, LeDoux JE, Blair HT (2010) Neural substrates for expectation-modulated fear learning in the amygdala and periaqueductal gray. *Nat Neurosci* 13(8):979–986.
- Yacubian J, et al. (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *J Neurosci* 26(37):9530–9537.
- Mobbs D, et al. (2007) When fear is near: Threat imminence elicits prefrontal-periaqueductal gray shifts in humans. *Science* 317(5841):1079–1083.

In conclusion, we describe for the first time to our knowledge individual biases in learning from actual painful outcomes on the one hand and from their prevention on the other. These biases are associated with qualitative differences in striatal prediction error signaling and predicted by differences in striatal structure. Further research should reveal how these functional and structural characteristics map onto psychiatric disorders that feature imbalanced harm avoidance behavior.

Materials and Methods

The experimental protocol was approved by the University of College London local research ethics committee, and informed consent was obtained from all participants. Electrical stimulation was individually titrated to induce a moderate subjective pain level. Participants performed the experiment while being scanned in a Siemens Trio 3T MRI scanner. See *SI Appendix* for further details of the experiment, modeling, and functional MRI procedures.

ACKNOWLEDGMENTS. This work was funded by Wellcome Trust's Cambridge-University College London Mental Health and Neurosciences Network Grant 095844/Z/11/Z (to E.E. and R.J.D.), Wellcome Trust Investigator Award 098362/Z/12/Z (to R.J.D.), the Gatsby Charitable Foundation (P.D.), and Swiss National Science Foundation Grant 151641 (to T.U.H.).

- Linnman C, Moulton EA, Barmettler G, Becerra L, Borsook D (2012) Neuroimaging of the periaqueductal gray: State of the field. *Neuroimage* 60(1):505–522.
- Navratilova E, Porreca F (2014) Reward and motivation in pain and pain relief. *Nat Neurosci* 17(10):1304–1312.
- Bushnell MC, Ceko M, Low LA (2013) Cognitive and emotional control of pain and its disruption in chronic pain. *Nat Rev Neurosci* 14(7):502–511.
- Leong JK, Pestilli F, Wu CC, Samanez-Larkin GR, Knutson B (2016) White-matter tract connecting anterior insula to nucleus accumbens correlates with reduced preference for positively skewed gambles. *Neuron* 89(1):63–69.
- Rogan MT, Leon KS, Perez DL, Kandel ER (2005) Distinct neural signatures for safety and danger in the amygdala and striatum of the mouse. *Neuron* 46(2):309–320.
- Salamone JD (1994) The involvement of nucleus accumbens dopamine in appetitive and aversive motivation. *Behav Brain Res* 61(2):117–133.
- Pezze MA, Feldon J (2004) Mesolimbic dopaminergic pathways in fear conditioning. *Prog Neurobiol* 74(5):301–320.
- McNally GP, Westbrook RF (2006) Predicting danger: The nature, consequences, and neural mechanisms of predictive fear learning. *Learn Mem* 13(3):245–253.
- Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. *Science* 315(5811):515–518.
- Bolstad I, et al. (2013) Aversive event anticipation affects connectivity between the ventral striatum and the orbitofrontal cortex in an fMRI avoidance task. *PLoS One* 8(6):e68494.
- Seymour B, Daw ND, Roiser JP, Dayan P, Dolan R (2012) Serotonin selectively modulates reward value in human decision-making. *J Neurosci* 32(17):5833–5842.
- Roy M, et al. (2014) Representation of aversive prediction errors in the human periaqueductal gray. *Nat Neurosci* 17(11):1607–1612.
- Mowrer OH, Lamoreaux RR (1946) Fear as an intervening variable in avoidance conditioning. *J Comp Psychol* 39:29–50.
- Bromberg-Martin ES, Matsumoto M, Hikosaka O (2010) Dopamine in motivational control: Rewarding, aversive, and alerting. *Neuron* 68(5):815–834.
- Ramirez F, Moscarello JM, LeDoux JE, Sears RM (2015) Active avoidance requires a serial basal amygdala to nucleus accumbens shell circuit. *J Neurosci* 35(8):3470–3477.
- Menon V, Uddin LQ (2010) Saliency, switching, attention and control: A network model of insula function. *Brain Struct Funct* 214(5–6):655–667.
- Moscarello JM, LeDoux JE (2013) Active avoidance learning requires prefrontal suppression of amygdala-mediated defensive reactions. *J Neurosci* 33(9):3815–3823.
- Choi JS, Cain CK, LeDoux JE (2010) The role of amygdala nuclei in the expression of auditory signaled two-way active avoidance in rats. *Learn Mem* 17(3):139–147.
- Collins KA, Mendelsohn A, Cain CK, Schiller D (2014) Taking action in the face of threat: Neural synchronization predicts adaptive coping. *J Neurosci* 34(44):14733–14738.
- Badrinarayan A, et al. (2012) Aversive stimuli differentially modulate real-time dopamine transmission dynamics within the nucleus accumbens core and shell. *J Neurosci* 32(45):15779–15790.
- Martinez RC, et al. (2013) Active vs. reactive threat responding is associated with differential c-Fos expression in specific regions of amygdala and prefrontal cortex. *Learn Mem* 20(8):446–452.
- van den Bos W, Rodriguez CA, Schweitzer JB, McClure SM (2014) Connectivity strength of dissociable striatal tracts predict individual differences in temporal discounting. *J Neurosci* 34(31):10298–10310.
- Kreitzer AC, Malenka RC (2008) Striatal plasticity and basal ganglia circuit function. *Neuron* 60(4):543–554.
- Apkarian AV, Hashmi JA, Baliki MN (2011) Pain and the brain: Specificity and plasticity of the brain in clinical chronic pain. *Pain* 152(3, Suppl):S49–S64.
- Schmidt-Wilcke T, et al. (2007) Striatal grey matter increase in patients suffering from fibromyalgia—a voxel-based morphometry study. *Brain* 132(Suppl 1):S109–S116.
- Rotge JY, et al. (2010) Gray matter alterations in obsessive-compulsive disorder: An atomic likelihood estimation meta-analysis. *Neuropsychopharmacology* 35(3):686–691.