

A computational and neural model of momentary subjective well-being

Robb B. Rutledge^{a,b,1}, Nikolina Skandali^a, Peter Dayan^c, and Raymond J. Dolan^{a,b}

^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 WC1N 3AR, United Kingdom

Edited by Wolfram Schultz, University of Cambridge, Cambridge, United Kingdom, and accepted by the Editorial Board July 2, 2014 (received for review April 30, 2014)

The subjective well-being or happiness of individuals is an important metric for societies. Although happiness is influenced by life circumstances and population demographics such as wealth, we know little about how the cumulative influence of daily life events are aggregated into subjective feelings. Using computational modeling, we show that emotional reactivity in the form of momentary happiness in response to outcomes of a probabilistic reward task is explained not by current task earnings, but by the combined influence of recent reward expectations and prediction errors arising from those expectations. The robustness of this account was evident in a large-scale replication involving 18,420 participants. Using functional MRI, we show that the very same influences account for task-dependent striatal activity in a manner akin to the influences underpinning changes in happiness.

reward prediction error | dopamine | striatum | insula

Philosophers from Aristotle to Bentham have argued for the central importance of subjective well-being in human conscious experience. Bentham suggested that “it is the greatest happiness of the greatest number that is the measure of right and wrong” (1). This dictum informs the policies of many nations who deploy population measures of well-being in pursuit of this goal (2). However, happiness is a difficult concept to define (3–5) and the complexity of the relationship between happiness and wealth (6–8) suggests that there is no simple happiness–reward relationship. Here, we provide an analysis of one of the foundations on which happiness is assumed to be built, namely the subjective response to rewards. We focus on rewards that are external quantifiable objects (e.g., money) that might elicit affective and motivational responses (9).

To address the relationship between reward and ongoing happiness, it is essential to be able to measure happiness reliably and to influence it on an appropriate time scale. Experience sampling is an established methodology that measures phenomenological states as subjects engage in daily life. By repeatedly asking participants to report on their subjective emotional state, these feelings can be related to antecedent life events including rewards (10–13). Momentary measures of happiness or hedonic well-being reveal emotional reactivity to recent events and thus differ from overall life satisfaction, although it is possible that life satisfaction relates to the temporal integral of momentary happiness over a longer time scale (12).

Here we asked subjects to perform a probabilistic reward task in which they chose between certain and risky monetary options while being asked after every few trials to report, “How happy are you right now?” We expected this task to elicit rapid changes in affective state, and we therefore used a more frequent variant of experience sampling adapted to laboratory and functional MRI (fMRI) settings. Importantly, the experiential sampling question makes no reference to past events and concerns the overall subjective emotional state rather than the cue-elicited emotional responses to reward-related stimuli that have been the focus of previous studies (14–16).

Much is now known about how the brain responds to rewards. For example, midbrain dopamine neurons represent reward prediction error (RPE) signals in animals (17–19) and humans (20). Neuroimaging studies report correlates of RPEs in the ventral striatum, an area that is a target for dopamine projections, in tasks from reinforcement learning (21, 22) to gambling (23). Many studies have also related subjective feelings about discrete events to neural activity (24–26). However, it remains unknown how these events cumulatively influence happiness.

We modeled behavioral data using a computational model inspired by models of dopamine function. Here we show that momentary subjective well-being is explained not by task earnings but by the cumulative influence of recent reward expectations and prediction errors resulting from those expectations. We note that the temporal difference errors that dopamine neurons are thought to represent are closely related to these quantities. Our model explained momentary subjective well-being better than a model that accounts for the influence of rewards but does not include a role for expectations. Furthermore, we replicated these behavioral findings in two laboratory-based behavioral experiments as well as a large-scale smartphone-based experiment. Using fMRI we probed the relationship between reward-related task events, neural responses to those events, and subjective well-being. Task-dependent neural activity in the ventral striatum, a major projection site for dopamine neurons, correlated with subsequent reports of subjective well-being, consistent with this area playing a role in changes in happiness.

We scanned 26 subjects while they made choices between certain and risky monetary options (Fig. 14). Chosen gambles

Significance

A common question in the social science of well-being asks, “How happy do you feel on a scale of 0 to 10?” Responses are often related to life circumstances, including wealth. By asking people about their feelings as they go about their lives, ongoing happiness and life events have been linked, but the neural mechanisms underlying this relationship are unknown. To investigate it, we presented subjects with a decision-making task involving monetary gains and losses and repeatedly asked them to report their momentary happiness. We built a computational model in which happiness reports were construed as an emotional reactivity to recent rewards and expectations. Using functional MRI, we demonstrated that neural signals during task events account for changes in happiness.

Author contributions: R.B.R., N.S., P.D., and R.J.D. designed research; R.B.R. and N.S. performed research; R.B.R. analyzed data; and R.B.R., P.D., and R.J.D. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. W.S. is a guest editor invited by the Editorial Board.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: robb.rutledge@ucl.ac.uk.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1407535111/-DCSupplemental.

were resolved after a brief delay. After every two to three choice trials, subjects were asked to report, “How happy are you right now?” by moving a slider. Responses were converted to a 0–100 scale (*SI Methods*). Our first goal was to establish whether there was a relationship between happiness ratings and prior rewards. Subjects earned £28.51 ± 7.60 (mean ± SD) in 150 trials, a significant increase in wealth from an initial £20 endowment [$t(25) = 5.7, P < 0.0001$]. Prizes up to £3 were sufficient to elicit changes in happiness with root-mean-square differences between successive ratings of 17 ± 8 (mean ± SD, range 6–43). However, self-reported overall happiness did not increase between the beginning and the end of the experiment [initial happiness: 60 ± 18, final happiness: 54 ± 20 (mean ± SD), $t(25) = -1.0, P = 0.33$; Fig. 1*B*]. The relationship between task earnings and the difference between initial and final happiness was not significant ($P = 0.16, r^2 = 0.079$).

Computational Model of Momentary Subjective Well-Being

We next examined the relationship between chosen certain rewards (CRs), the expected values (EVs) of chosen gambles, and RPEs (the difference between experienced and predicted rewards) and happiness. Note that these quantities including EVs and RPEs are linked to dopamine activity (17) and we hypothesized that these dopamine-related quantities might explain momentary happiness. We considered influences that decay exponentially in time (*SI Methods*):

$$\text{Happiness}(t) = w_0 + w_1 \sum_{j=1}^t \gamma^{t-j} CR_j + w_2 \sum_{j=1}^t \gamma^{t-j} EV_j + w_3 \sum_{j=1}^t \gamma^{t-j} RPE_j,$$

where t is the trial number, w_0 is a constant term, other weights w capture the influence of different event types, $0 \leq \gamma \leq 1$ is a forgetting factor that makes events in more recent trials more

influential than those in earlier trials, CR_j is the CR if chosen instead of a gamble on trial j , EV_j is the EV of a gamble (average reward for the gamble) if chosen on trial j , and RPE_j is the RPE on trial j contingent on choice of the gamble. If the CR was chosen, then $EV_j = 0$ and $RPE_j = 0$; if the gamble was chosen, then $CR_j = 0$. Parameters were fit to happiness ratings in individual subjects. We found that CR, EV, and RPE weights were on average positive [all $t(25) > 4.6, P < 0.0001$] with EV weights lower than RPE weights [$t(25) = 4.3, P < 0.001$; Fig. 2*A*]. The forgetting factor γ was 0.61 ± 0.30 (mean ± SD). This model explained moment-to-moment fluctuations in happiness well with $r^2 = 0.47 \pm 0.21$ (mean ± SD; Fig. 1) and, when judged according to complexity, explained this reactive happiness better than a range of alternative models, including models without exponential constraints, parameters for unchosen options, and utility-based models (*SI Methods*, Tables S1 and S2, and Fig. S1).

A prediction arising out of our model is that the final happiness rating should depend only on the final task events and not on earlier events. Consider the effect of just receiving £1 versus receiving £1 five trials ago. For a subject with the group average forgetting factor of 0.61, the latter would have only 8% of the impact of the former. For subjects with larger (0.8) or smaller (0.4) forgetting factors, this relative impact would be 33% or 1%, respectively. For most subjects, rewards received more than 10 trials in the past should have little effect on current happiness (Fig. S1). Therefore, any relationship between task earnings and change in happiness from the initial to the final rating should be accounted for by the final trials of the task. We used only the final 10 trials to predict the final happiness rating based on parameters estimated from a model fit to the first 140 trials. Residual errors of the predicted final happiness were uncorrelated with task earnings ($P = 0.81, r^2 = 0.003$), suggesting that our model accounts for any relation between task earnings and happiness.

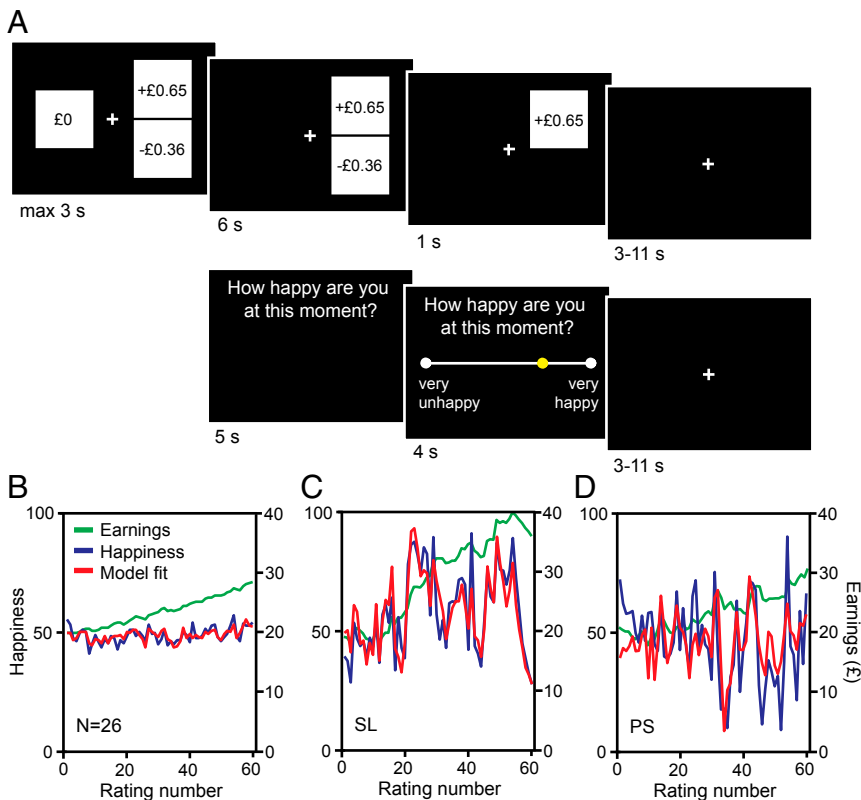


Fig. 1. Effect of previous rewards and expectations on happiness ratings. (A) Experimental design. In each trial, subjects chose between a certain option and a gamble. Chosen gambles were resolved after a 6-s-delay period. Every two to three trials, subjects were asked to indicate “How happy are you at this moment?” by using button presses to move a cursor. (B–D) Cumulative task earnings and happiness ratings across subjects ($n = 26$) in B and in example subjects in C and D. Happiness model fits are displayed for the model in Fig. 2*A* [$r^2 = 0.47 \pm 0.21$ (mean ± SD); example subjects $r^2 = 0.79$ in C and $r^2 = 0.41$ in D].

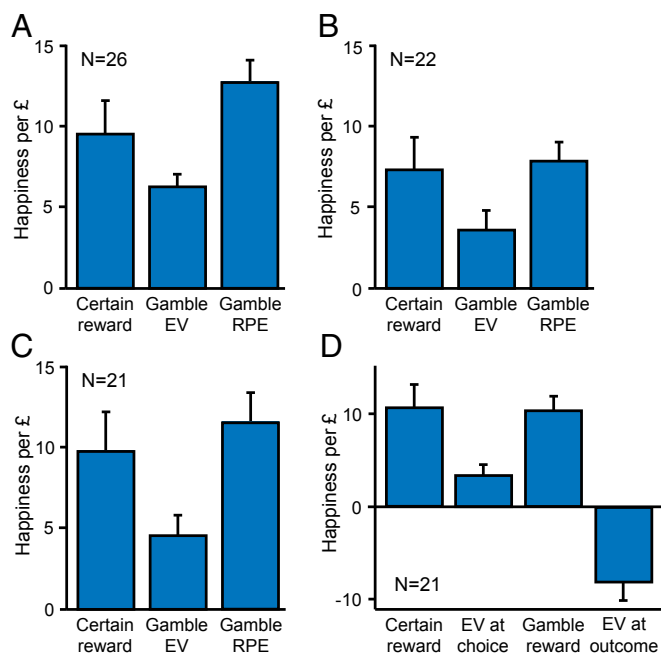


Fig. 2. Computational model fits for three experiments. (A) The computational model that best explained happiness in the fMRI experiment ($n = 26$) had positive weights for previous CRs, gamble EVs, and gamble RPEs. (B) A behavioral experiment (“current earnings always shown”; $n = 22$) in which the current level of wealth was displayed at all times during the experiment, including during happiness ratings, replicated behavioral findings from the fMRI experiment. (C) An additional behavioral experiment (“only some gamble outcomes shown”; $n = 21$) similarly replicated previous findings. In this experiment, gamble choices had a 50% probability of ending with the text “outcome added to total” instead of the outcome being revealed. (D) This experimental design allowed the separation of expectation effects related to choices and outcomes. When the RPE term (reward minus EV) was split into separate GR and gamble EV terms, happiness ratings were positively correlated with GRs and negatively correlated with gamble EV at outcome.

We ran three additional behavioral experiments to validate our model (Fig. S2). Because subjects may be poor at estimating their current earnings, we conducted a behavioral experiment (“current earnings always shown”) in which current task earnings were always displayed, including at the time when happiness ratings were made (SI Methods). If subjects are not reporting their momentary happiness but instead their belief about their current success or earnings, we should see a strong relationship between ratings and earnings in this experiment. Instead, we replicated our previous findings with positive CR, EV, and RPE weights, and EV weights lower than RPE weights [all $t(21) > 3.0$, $P < 0.01$, $n = 22$; Fig. 2B]. The relationship between earnings and change in happiness was stronger than in the scanning experiment ($P = 0.042$, $r^2 = 0.19$). However, when we again used only the final 10 trials of the task to predict the final happiness rating, the residual errors of the predicted final happiness were uncorrelated with task earnings ($P = 0.17$, $r^2 = 0.094$), suggesting that even when subjects always know their exact earnings, the model still accounts for any relation between task earnings and happiness.

To verify a role for both reward expectations and RPEs in happiness, we also conducted an additional behavioral experiment in which the actual outcomes were only presented in randomly selected trials (“only some gamble outcomes shown”). Otherwise, the text “outcome added to total” was displayed when the outcome would normally be revealed (SI Methods). We again replicated our previous findings showing that CR, EV, and RPE weights are all positive and EV weights are lower than RPE weights

[all $t(20) > 3.5$, $P < 0.005$, $n = 21$; Fig. 2C]. This model fits the data better (median $r^2 = 0.43$) than an alternative model with a gamble reward (GR) outcome term (median $r^2 = 0.39$) instead of an RPE term (SI Methods). There was again no relationship between task earnings and change in happiness ($P = 0.24$, $r^2 = 0.07$). When we fit a model with separate EV weights depending on whether the gamble outcome was revealed, we found that expectations had a positive effect on happiness even when outcomes were not revealed [$t(20) = 2.8$, $P = 0.011$]. This task version allowed us to dissociate the effects of expectations at choice and outcome as well as to apply a more stringent test for a relationship between a signal and RPEs (27) in which the RPE term is split into its separate components: rewards and expectations (SI Methods). If happiness is positively affected by RPEs, then because RPE is equal to the reward minus EV, the weight for EV should be negative and subjects with larger negative EV weights should have larger positive reward weights to balance the two RPE components. As predicted, happiness was positively modulated by reward [$t(20) = 6.6$, $P < 0.0001$] and negatively modulated by EV [$t(20) = -4.3$, $P < 0.001$] and weights were anticorrelated ($r = -0.58$, $P = 0.006$; Fig. 2D).

Finally, laboratory experiments are necessarily based on relatively small numbers of subjects, raising issues of generalizability and demand characteristics. We were able to address these potential shortcomings by using a smartphone-based platform (The Great Brain Experiment, www.thegreatbrainexperiment.com) for iOS (Apple) and Android (Google) systems (SI Methods) that enabled us to run a 30-trial 12-rating version of the task. Here our sample comprised 18,420 anonymous unpaid participants who made over 200,000 happiness ratings. We divided the data into 92 subsets of 200 consecutive participants and fit our model in individual participants. CR, EV, and RPE weights were positive in all 92 data subsets [$t(199) > 2.0$, $P < 0.05$; Fig. 3]. EV weights were lower than RPE weights in all but one data subset [$t(199) > 2.0$, $P < 0.05$]. To test whether the model still applied when participants had minimal familiarity with the experimental context, we analyzed the first happiness rating preceded by a choice trial from each participant. CR, EV, and RPE weights were again positive with EV weights lower than RPE weights for a single happiness rating from each of 18,420 participants (all $P < 0.005$; SI Methods and Fig. 3C and Fig. S1C). Consistent with our previous results, earnings increased on average over time but happiness did not (Fig. S2C). Because the cursor always started at the midpoint on the rating scale in all experiments, this starting point might act as an anchor to counteract increases in happiness due to task earnings. In this case, happiness should increase in subjects who both increased their earnings and had an average happiness below the midpoint, but there was only a modest increase in these subjects [$n = 2,211$, initial happiness: 42 ± 18 , final happiness: 43 ± 19 (mean \pm SD)], suggesting that this potential influence does not explain why happiness does not increase with earnings (Fig. S3). We also verified the out-of-sample validity of our model by using parameter weights from the fMRI experiment to predict happiness ratings in the other experiments (median $r^2 > 0.23$ for the three experiments; Table S3).

Momentary Subjective Well-Being in Striatum and Insula

To test the relationship between neural activity in the fMRI experiment and the current level of happiness, we regressed blood-oxygen-level-dependent (BOLD) activity at task events from trials preceding happiness ratings (option and outcome onsets) on those subsequent ratings. BOLD activity in the ventral striatum was significantly correlated with z-scored future happiness ratings (Fig. 4A; left coordinates $-9, 8, -8$; $t(23) = 5.1$; right coordinates $18, 8, -5$, $t(23) = 3.9$; both $P < 0.05$, small-volume corrected). We then tested whether this activity was related to parameters of our behavioral model, regressing event-related activity in this region of interest (ROI) on parametric task

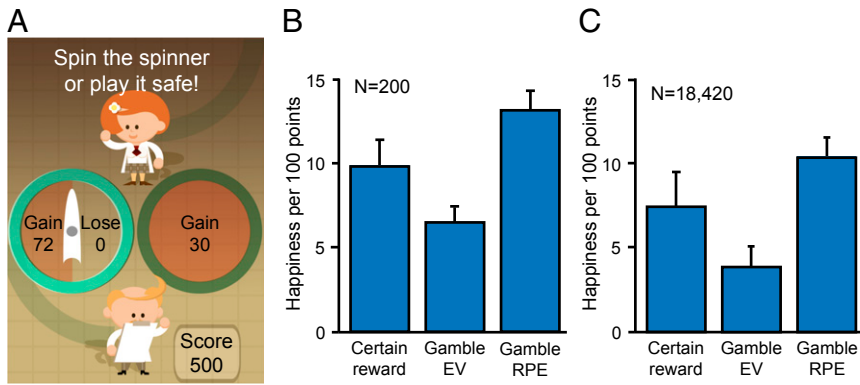


Fig. 3. Smartphone-based large-scale replication. (A) A screenshot from the smartphone experiment. In each trial, participants chose between a certain option and a gamble. Here the choice is between a certain 30 points and a gamble to gain 72 points or 0 points. Every two to three trials, participants were asked to indicate, “How happy are you at this moment?” (B) The computational model that best explained happiness in the first 200 participants had positive CR, gamble EV, and gamble RPE weights. Error bars represent SEM. (C) This model also explained the happiness rating after the first two to three trials from each participant, with similar model fits for a single happiness rating from each participant ($n = 18,420$). Error bars represent SE computed from the covariance matrix of the single model fit.

variables (*SI Methods*). The striatal weights associated with these factors were all significantly positive (all $P < 0.001$; Fig. 4B), as was the case in our behavioral model (Fig. 2A). As for the behavioral analysis, we again applied the higher standard of splitting the RPE term into separate components (27) and verified that striatal BOLD activity was positively modulated by EV at choice time, positively modulated by rewards at outcome time, and negatively modulated by EV at outcome time (all $P < 0.05$). We found similar results when we repeated the analysis in a fully independent ventral striatum ROI derived from the literature (Fig. S4). A test of whether striatal activity mediates the relationship between past events and happiness was not significant ($P > 0.2$) for either ventral striatum ROI (*SI Methods* and Fig. S5).

We next tested for neural correlates of the subjective well-being state itself, at the time when subjects rated “How happy are you at this moment?” (5 s before subjects could initiate a motor response) by regressing BOLD activity on parametric happiness ratings. BOLD activity in the ventral striatum ROI was uncorrelated with happiness ratings at this time ($P = 0.97$). However, right anterior insula BOLD activity was positively correlated with z-scored happiness ratings (Fig. 5A; coordinates 42, 5, -14; $t(23) = 4.3$, $P < 0.01$, small-volume corrected). This finding of a relationship between a subjective state and activity in right anterior insula is strikingly consistent with prior evidence that this area supports interoceptive (28, 29) and emotional awareness (30, 31). We also note evidence that gray matter volume in this area is greater in people reporting higher eudaimonic well-being (32). To explore this relationship further, we also measured life happiness in our subjects by asking before each experiment, “Taken all together, how happy are you with your life these days?” (Fig. S6). We tested whether neural responses varied with life happiness and found similar insula parameter estimates in subjects with high or low life happiness (Fig. 5B) and no relationship between life happiness and parameter estimates ($r = -0.20$, $P = 0.36$). This dissociation highlights the neural as well as psychological differences between different forms of happiness.

Discussion

Conscious emotional feelings, such as momentary happiness, are core to the ebb and flow of human mental experience. Our computational model suggests momentary happiness is a state that reflects not how well things are going but instead whether things are going better than expected. This includes positive and negative expectations, even in the absence of outcomes. Our results are in accord with well-being studies suggesting that the ongoing state of happiness varies around a hedonic set point (33). Indeed, an average forgetting factor of 0.61 in our fMRI experiment indicates that rewards received more than 10 trials ago have essentially no influence on current happiness in the

task. Although we used a stationary environment, value estimates in a dynamic environment could be continually updated, allowing agents to adapt to persistent changes in reward rates. This prediction would be consistent with the “hedonic treadmill,” the finding that individuals adapt to long-lasting changes in life circumstances (34, 35).

Across four separate studies we obtained qualitatively similar parameter estimates for model fits, including for a large-scale smartphone-based replication. Participants in the smartphone experiment were unpaid and anonymous, thus minimizing demand characteristics that may exist in the laboratory. However, even in this context we saw little change in the results, even when considering only a single happiness rating from each of the 18,420 participants, suggesting that the intrinsic rewards in smartphone games are sufficient to affect momentary happiness in the absence of monetary incentives.

The observation that recent rewards affected ongoing reactive happiness ratings is to be expected. More surprising, and a key finding across each of our studies, was the important role for expectations in determining happiness. This role balanced the influence of reward in a manner consistent with prediction errors playing a key determining role in momentary happiness. Because RPEs capture the difference between rewards and expectations, the effect of an outcome on happiness depends on what the alternatives are. For example, a £0 prize decreases happiness if the alternative was winning £2, but increases happiness if the alternative was losing £2. Expectations also affect happiness before outcomes are revealed and we show that choosing a gamble with a positive EV makes subjects happier even if the outcome is never revealed. In fact in the real world, rewards associated with life decisions are often not realized for a long time (e.g., jobs, marriage) and our results suggest that expectations

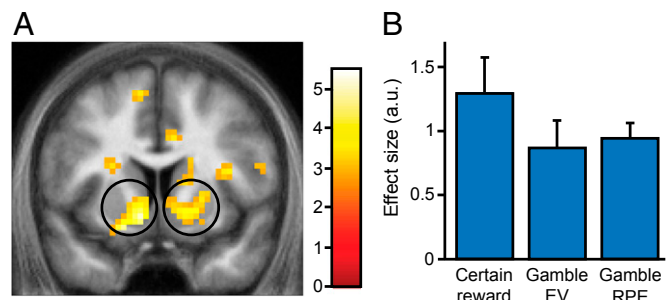


Fig. 4. Relationship between happiness and neural responses during preceding events. (A) Striatal activity during task events preceding subjective state ratings correlated with later self-reported happiness ($P < 0.05$, small-volume corrected). (B) Neural responses in ventral striatum were explained by the same parametric task variables as the variables that explained happiness. Error bars represent SEM.

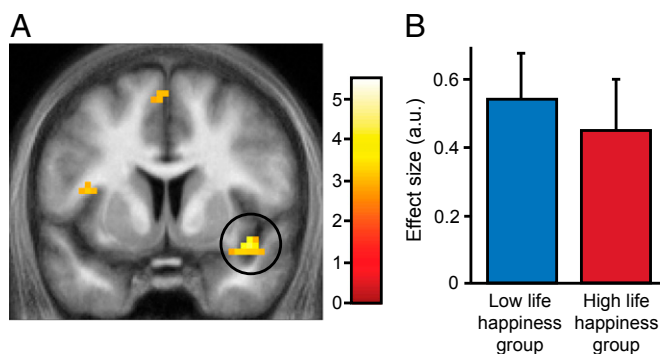


Fig. 5. Effect of the happiness question on neural activity in the right anterior insula. (A) In the right anterior insula, neural activity at the time of the happiness question presentation correlated with how happy subjects reported being ($P < 0.01$, small-volume corrected). (B) Parameter estimates were similar for subjects with low or high life happiness. Error bars represent SEM.

related to those decisions, both good and bad, do have an impact on happiness. In all our experiments, we found that EV weights were significantly lower than RPE weights, indicating that the overall effect of expectations on happiness is negative: a positive weight for EV in choices combines with a larger negative weight for EV in outcomes (because RPE is equal to reward minus EV) for an overall negative expectation effect. As a result, positive expectations effectively reduce the overall emotional impact of trials with positive outcomes and negative expectations effectively reduce the overall emotional impact of trials with negative outcomes. A role for expectations, independent of outcomes, also ensures that subjects are happier on average after choosing gambles with positive rather than negative EV.

The well-described peak-end rule (36) suggests that remembered feelings depend most on how experiences were at the peak and end, explaining for example why painful medical procedures are remembered as less unpleasant when extended to end with a less painful period (37, 38). Our model does not give extra weight to peak events because task outcomes are relatively homogeneous. However, in common with the rule, we found that recent events have relatively greater impact on mood, extending findings of previous studies to the “experiencing” and not just the “remembering” self. Pleasant and unpleasant feelings during medical procedures could be modeled using experience-sampling techniques and the role we highlight for expectations could potentially be used to leverage better outcomes. Our results also have potential policy implications. Lowering expectations increase the probability of positive outcomes (something routinely observed). However, lower expectations reduce well-being before an outcome arrives, limiting the beneficial scope of this manipulation. One intriguing notion would be to use a sufficiently negative expectation to create an overall positive emotional impact from a negative event. For example, news of a 1-h flight delay preceded by news that there is a 50% probability of a 6-h delay should, by our model, have a net positive impact for the average passenger. However, floating the extreme possibility might well have other negative consequences for the airline.

Our key finding is that happiness is related to quantities associated with temporal difference errors that phasic dopamine release is thought to represent, errors that signal changes in long-term expected reward (17–19, 39–41). We also found that happiness relates to BOLD activity in the striatum, a prominent target for ascending dopamine projections. At the very least, this hints at a link between dopamine and emotional state, consistent with suggestions that this neuromodulator plays a role in mood regulation in healthy and depressed subjects (42). The potential importance of this link is bolstered by observations that stress

engenders depression-like behavior via a modulation of striatal dopamine responses (43). A dopamine manipulation alone does not impact on overall mood (44) but, based on our model, we suggest that dopamine may act to modulate changes in emotional state in response to discrete events. However, one caveat is that we did not find that striatal activity significantly mediates the relationship between RPEs and happiness in our task. Although our results suggest that the ventral striatum is the best candidate region for mediating this relationship, this mediation may be sensitive to task demands and may be absent in tasks like ours where RPEs do not modify behavior. Striatal RPEs are notably absent in subjects that fail to learn in a reinforcement-learning task (45). Therefore, we predict that striatal activity in reinforcement-learning tasks will mediate the relationship between RPEs and happiness only in subjects who learn the reward associations.

Aberrant responses to daily life events are a defining characteristic of mood disorders. Our findings show that conscious emotional states can be precisely manipulated and characterized using computational models in a similar manner to studies of conscious perception (46). Our approach offers a rich quantitative means of relating emotional state to brain and behavior and in doing so provides a framework for the development of model-based assays of mood disorders that can be exploited so as to probe the underlying neurobiological mechanisms.

Methods

Subjects. Twenty-six healthy right-handed subjects took part in the fMRI experiment (age range 20–40 y, seven male). Twenty-one of these 26 subjects agreed on invitation to participate in an additional behavioral experiment (median 53 d apart; range, 3–162 d). Eleven of the 26 subjects agreed on invitation to participate in a second behavioral experiment (14–17 mo later). An additional 11 subjects were recruited for this second behavioral experiment (age range 20–34, four male). Subjects were endowed with £20 at the beginning of each experiment and paid according to performance. Two subjects were excluded from fMRI analysis due to excessive head movement. In a smartphone-based experiment we tested 18,420 unpaid participants (age 18 y and over, 8,557 male). All subjects gave informed consent and the Research Ethics Committee of University College London approved all studies.

Laboratory-Based Experimental Tasks. During each of the 150 trials of the tasks, subjects chose between a certain option and a gamble, with equal probabilities of two outcomes (Fig. 1A and *SI Methods*). There were three trial types: mixed trials (a certain amount of £0 or a gamble with a gain and loss amount), gain trials (a certain gain or a gamble with £0 and a larger gain), and loss trials (a certain loss or a gamble with £0 and a larger loss). Gamble choices remained on the screen for 6 s before gamble outcomes were revealed for 1 s. Subjects were presented with the question, “How happy are you at this moment?” after every two to three trials. After a 5-s delay period, a rating line appeared with endpoints labeled “very unhappy” and “very happy.” Subjects had 4 s to move the cursor along the scale with button presses, making a total of 63 happiness ratings. Current earnings were displayed after each of the three blocks of 50 trials. In the only-some-gamble-outcomes-shown behavioral experiment, although all choices counted for real money, only some gamble outcomes were revealed, enabling us to dissociate expectation effects at choice and outcome (*SI Methods*). After half of gamble choices, the delay period would end with the text “outcome added to total.” In the current-earnings-always-shown behavioral experiment, the current task earnings were displayed at all times, including during happiness ratings. Before each experiment, before the task instructions, we measured life happiness by asking subjects, “Taken all together, how happy are you with your life these days?”

Smartphone-Based Experimental Task. Researchers at the Wellcome Trust Centre for Neuroimaging at University College London worked with White Bat Games to develop The Great Brain Experiment, available as a free download on iOS and Android systems. One of these games was based on the task we used for the fMRI experiment. Subjects started the game with 500 points and made 30 choices in each play. In each trial, subjects chose between a certain option and a gamble. Chosen gambles, represented as spinners, were resolved after a brief delay. Subjects were presented with the question, “How happy are you at this moment?” after every two to three trials.

Subjects completed 30 choice trials and answered the happiness question 12 times in each play.

Happiness Computational Models. We modeled moment-to-moment happiness for all ratings preceded by choices using models that assume an exponential decay of previous event influences and terms for CRs, gamble EVs, and RPEs. Models were fit using nonlinear least squares using the optimization toolbox in MATLAB (MathWorks, Inc.). To verify the appropriateness of our model, we tested a wide variety of alternative models (*SI Methods* and *Tables S1* and *S2*) including utility-based models following established procedures for estimating utilities (47). We evaluated the models using Bayesian model comparison techniques (48, 49).

fMRI Data Acquisition and Analysis. All scanning was performed on a 3-Tesla Siemens Allegra scanner with a Siemens head coil with an echo-planar sequence (*SI Methods*). We used standard preprocessing in SPM8 (Wellcome Trust Centre for Neuroimaging) and estimated three general linear models for each subject which included regressors related to option and gamble values and happiness ratings (*SI Methods*). Statistical significance was

determined at the group level using a random-effects analysis. All analyses used a voxel-wise significance threshold of $P < 0.001$ and a corrected significance threshold of $P < 0.05$ based on a family-wise error cluster-level small-volume correction centered on coordinates from previous studies (*SI Methods*). We used the Multilevel Mediation and Moderation Toolbox (50, 51) to perform the mediation analysis (*SI Methods* and *Fig. S5*).

ACKNOWLEDGMENTS. We thank Helen Barron, Tim Behrens, Molly Crockett, Benedetto De Martino, Thomas FitzGerald, Mona Garvert, Marc Guitart-Masip, Laurence Hunt, Zeb Kurth-Nelson, Stephanie Lazzaro, and Nicholas Wright for helpful comments. We also thank Rick Adams, Harriet Brown, Peter Smittenaar, Peter Zeidman, and Neil Millstone for developing the smartphone app; Ric Davis, Chris Freemantle, and Rachael Maddock for supporting data collection; and Dan Jackson (University College London) and Craig Brierley and Chloe Sheppard (Wellcome Trust). This work was supported by the Max Planck Society (R.B.R. and R.J.D.), the Gatsby Charitable Foundation (P.D.), and Wellcome Trust Grant 078865/Z/05/Z (to R.J.D.). The Wellcome Trust Centre for Neuroimaging is supported by core funding from Wellcome Trust Grant 091593/Z/10/Z. App development was funded by Wellcome Trust Engaging Science: Brain Awareness Week Award 101252/Z/13/Z.

- Bentham J (1776) *A Fragment on Government* (T. Payne, London).
- Oswald AJ, Wu S (2010) Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science* 327(5965):576–579.
- Russell JA (1980) A circumplex model of affect. *J Pers Soc Psychol* 39(6):1161–1178.
- Diener E, Suh EM, Lucas RE, Smith HL (1999) Subjective well-being: Three decades of progress. *Psychol Bull* 125(2):276–302.
- Barrett LF, Mesquita B, Ochsner KN, Gross JJ (2007) The experience of emotion. *Annu Rev Psychol* 58:373–403.
- Layard R (2005) *Happiness: Lessons from a New Science* (Allen Lane, London).
- Kahneman D, Deaton A (2010) High income improves evaluation of life but not emotional well-being. *Proc Natl Acad Sci USA* 107(38):16489–16493.
- Easterlin RA, McVey LA, Switek M, Sawangfa O, Zweig JS (2010) The happiness-income paradox revisited. *Proc Natl Acad Sci USA* 107(52):22463–22468.
- Berridge KC, Robinson TE (1998) What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Res Rev* 28(3):309–369.
- Csikszentmihalyi M, Larson R (1987) Validity and reliability of the Experience-Sampling Method. *J Nerv Ment Dis* 175(9):526–536.
- Reis HT, Gable SL (2000) Event-sampling and other methods for studying everyday experience. *Handbook of Research Methods in Social and Personality Psychology*, eds Reis HT, Judd CM (Cambridge Univ Press, New York), pp 190–222.
- Kahneman D, Krueger AB, Schkade DA, Schwarz AA (2004) A survey method for characterizing daily life experience: The day reconstruction method. *Science* 306(5702):1776–1780.
- Killingsworth MA, Gilbert DT (2010) A wandering mind is an unhappy mind. *Science* 330(6006):932.
- Knutson B, Adams CM, Fong GW, Hommer D (2001) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J Neurosci* 21(16):RC159.
- Knutson B, Fong GW, Bennett SM, Adams CM, Hommer D (2003) A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: Characterization with rapid event-related fMRI. *Neuroimage* 18(2):263–272.
- Samanez-Larkin GR, et al. (2007) Anticipation of monetary gain but not loss in healthy older adults. *Nat Neurosci* 10(6):787–791.
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275(5306):1593–1599.
- Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47(1):129–141.
- 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.
- Zaghloul KA, et al. (2009) Human substantia nigra neurons encode unexpected financial rewards. *Science* 323(5920):1496–1499.
- O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. *Neuron* 38(2):329–337.
- 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.
- Rutledge RB, Dean M, Caplin A, Glimcher PW (2010) Testing the reward prediction error hypothesis with an axiomatic model. *J Neurosci* 30(40):13525–13536.
- O'Doherty JP, Buchanan TW, Seymour B, Dolan RJ (2006) Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* 49(1):157–166.
- Rolls ET, Grabenhorst F, Parris BA (2008) Warm pleasant feelings in the brain. *Neuroimage* 41(4):1504–1513.
- Xiang T, Lohrenz T, Montague PR (2013) Computational substrates of norms and their violations during social exchange. *J Neurosci* 33(3):1099–1108.
- Behrens TEJ, Hunt LT, Woolrich MW, Rushworth MFS (2008) Associative learning of social value. *Nature* 456(7219):245–249.
- Critchley HD, Wiens S, Rotshtein P, Öhman A, Dolan RJ (2004) Neural systems supporting interoceptive awareness. *Nat Neurosci* 7(2):189–195.
- Damasio A, Carvalho GB (2013) The nature of feelings: Evolutionary and neurobiological origins. *Nat Rev Neurosci* 14(2):143–152.
- Craig AD (2009) How do you feel—now? The anterior insula and human awareness. *Nat Rev Neurosci* 10(1):59–70.
- Singer T, Critchley HD, Preusschoff K (2009) A common role of insula in feelings, empathy and uncertainty. *Trends Cogn Sci* 13(8):334–340.
- Lewis GJ, Kanai R, Rees G, Bates TC (2014) Neural correlates of the 'good life': Eudaimonic well-being is associated with insular cortex volume. *Soc Cogn Affect Neurosci* 9(5):615–618.
- Lykken D, Tellegen A (1996) Happiness is a stochastic phenomenon. *Psychol Sci* 7(3):186–189.
- Brickman P, Campbell DT (1971) Hedonic relativism and planning the good society. *Adaptation-Level Theory*, ed Appley MH (Academic, New York), pp 287–305.
- Frederick S, Loewenstein G (1999) Hedonic adaptation. *Well-Being: The Foundations of Hedonic Psychology*, eds Kahneman D, Diener E, Schwarz N (Russell Sage Foundation, New York), pp 302–329.
- Kahneman D (2000) Experienced utility and objective happiness: A moment-based approach. *Choices, Values, and Frames*, eds Kahneman D, Tversky A (Cambridge Univ Press, New York), pp 673–692.
- Kahneman D, Fredrickson BL, Schreiber CA, Redelmeier DA (1993) When more pain is preferred to less: Adding a better end. *Psychol Sci* 4(6):401–405.
- Redelmeier DA, Katz J, Kahneman D (2003) Memories of colonoscopy: A randomized trial. *Pain* 104(1–2):187–194.
- Sutton RS, Barto AG (1998) *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).
- Montague PR, Dayan P, Sejnowski TJ (1996) A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J Neurosci* 16(5):1936–1947.
- Hart AS, Rutledge RB, Glimcher PW, Phillips PEM (2014) Phasic dopamine release in the rat nucleus accumbens symmetrically encodes a reward prediction error term. *J Neurosci* 34(3):698–704.
- Nestler EJ, Carlezon WA, Jr (2006) The mesolimbic dopamine reward circuit in depression. *Biol Psychiatry* 59(12):1151–1159.
- Lemos JC, et al. (2012) Severe stress switches CRF action in the nucleus accumbens from appetitive to aversive. *Nature* 490(7420):402–406.
- Liggins J, Pihl RO, Benkelfat C, Leyton M (2012) The dopamine augments L-DOPA does not affect positive mood in healthy human volunteers. *PLoS ONE* 7(1):e28370.
- 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.
- Koch C (2004) *The Quest for Consciousness: A Neurobiological Approach* (Roberts and Co, Englewood, CO).
- Sokol-Hessner P, et al. (2009) Thinking like a trader selectively reduces individuals' loss aversion. *Proc Natl Acad Sci USA* 106(13):5035–5040.
- Schwarz G (1978) Estimating the dimension of a model. *Ann Stat* 6(2):461–464.
- Burnham KP, Anderson DR (1998) *Model Selection and Inference* (Springer, New York).
- Wager TD, Davidson ML, Hughes BL, Lindquist MA, Ochsner KN (2008) Prefrontal-subcortical pathways mediating successful emotion regulation. *Neuron* 59(6):1037–1050.
- Atlas LY, Bolger N, Lindquist MA, Wager TD (2010) Brain mediators of predictive cue effects on perceived pain. *J Neurosci* 30(39):12964–12977.