

A neurocomputational model for intrinsic reward

Benjamin Chew^{1,2}, Bastien Blain^{1,2}, Raymond J Dolan^{1,2}, Robb B Rutledge^{1,2}

¹ Max Planck University College London Centre for Computational Psychiatry and Ageing Research, London WC1B 5EH, United Kingdom.

² Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3BG, United Kingdom.

SUMMARY

Standard economic indicators provide an incomplete picture of what we value both as individuals and as a society. Furthermore, canonical macroeconomic measures, such as GDP, do not account for non-market activities (e.g., cooking, childcare) that nevertheless impact well-being. Here, we introduce a computational tool that measures the subjective reward value of experiences (e.g., playing a musical instrument without errors). We go on to validate this tool with neural data, using fMRI to measure neural activity in subjects performing a reinforcement learning task that incorporated periodic ratings of subjective affective state. Learning performance determined level of payment (i.e., extrinsic reward). Crucially, the task also incorporated a skilled performance component (i.e., intrinsic reward) which did not influence payment. Both extrinsic and intrinsic rewards influenced affective dynamics, and their relative influence could be captured in our computational model. Individuals for whom intrinsic rewards had a greater influence on affective state than extrinsic rewards had greater ventromedial prefrontal cortex (vmPFC) activity for intrinsic than extrinsic rewards. Thus, we show that computational modelling of affective dynamics can index the subjective value of intrinsic relative to extrinsic rewards, a 'computational hedonometer' that reflects both behavior and neural activity that quantifies the subjective reward value of experience.

INTRODUCTION

A key index of quality of life is subjective well-being which reflects “how people experience and evaluate their lives and specific domains and activities in their lives” (1). Individuals with higher subjective well-being display lower mortality rates (2, 3) and have a lesser risk of disease (4). In the workplace, employees who report higher subjective well-being have higher productivity without loss of output quality (5), reduced rates of absenteeism (6), and are rated more positively by their supervisors (7). On this basis, maximising subjective well-being should be of prime interest not only to individuals but also to companies and governments, as well as a target for health and economic policies (8).

A problem arises when it comes to designing effective measures likely to increase well-being. When contemplating the future, people exhibit biases in *affective forecasting* when making predictions about what it would feel like to experience specific events, consistently misjudging how future events will impact their affective state and leading them to perform actions that may be detrimental to maximization of their subjective well-being (9, 10). In particular, people overestimate both the intensities and durations of their hedonic responses to future events, and this bias is referred to as an impact bias (11, 12). Furthermore, the value of tangible goods can be quantified by prices or willingness-to-pay (13), but the value of intangible goods and experiences that are intrinsically rewarding (e.g., hobbies, recreational sports) are more difficult to define or elicit accurately due to biases (14, 15), while the predictive validity of implicit measures is unclear (16, 17).

Neuroscience-informed methods can provide a means to evaluate the subjective value of an intrinsic reward (e.g., the experience of mastering a musical composition for its own sake), allowing extrinsic and intrinsic rewards to be compared using a common scale of objectively measured neural activity (18). We hypothesized that extrinsic and intrinsic rewards would both influence affective states, and the extent of their relative influences should be reflected in regional brain activity. Recent studies (19–21) demonstrate that experience sampling during reward-based tasks can link affective and motivational responses to extrinsic reward. Here we extend this approach to investigate how affective state is influenced by the history of intrinsic rewards.

We developed a reinforcement learning task incorporating both an explicit reward component and a skilled performance component, where the latter did not affect payment (Figure 1A). On each trial, subjects selected one of two options, one of which was on average more rewarding than the other and then navigated a cursor past a series of barriers (see Experimental Procedures). We hypothesized that the experience of successful skilled performance, a source of intrinsic reward, would influence the momentary happiness of subjects in a manner that is quantitatively akin to the impacts of extrinsic rewards and that this would also be evident at the level of neural activity.

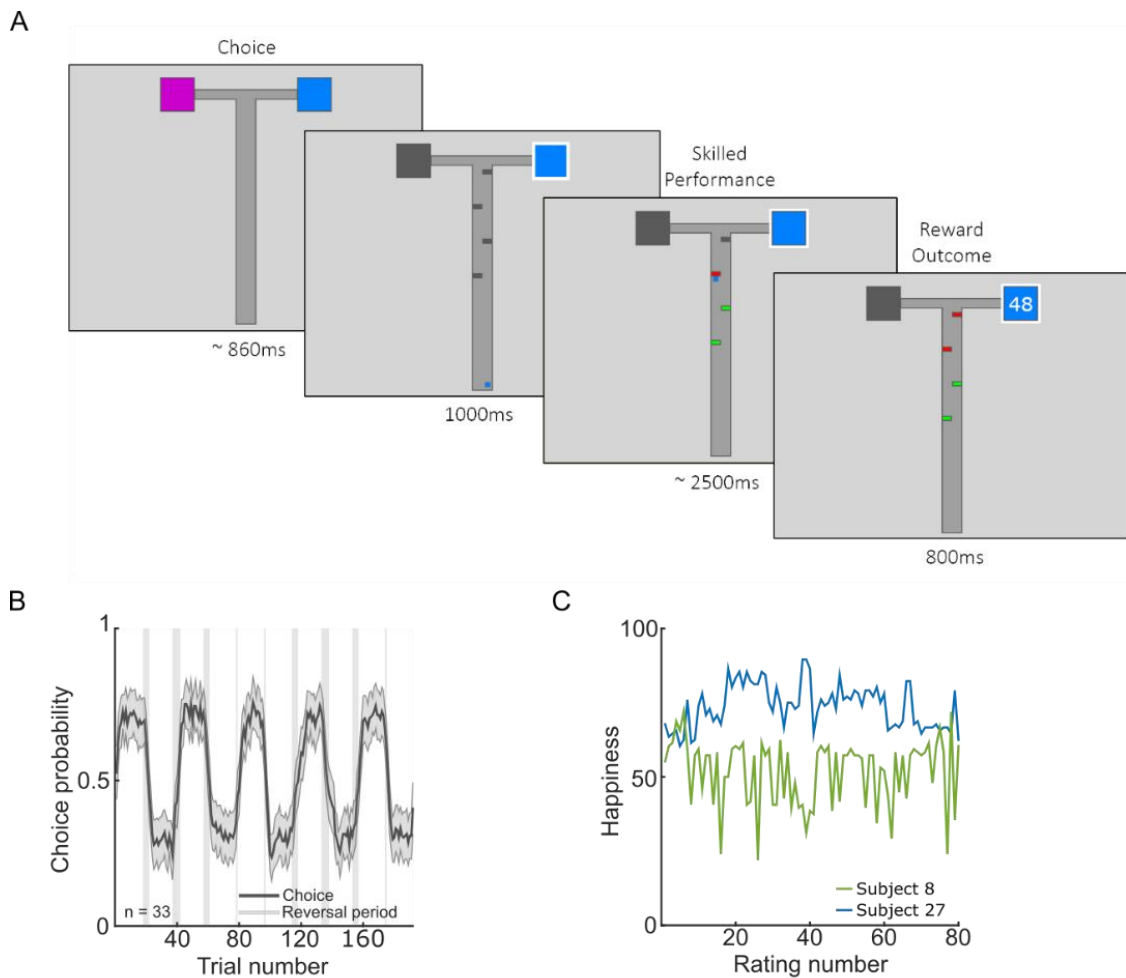


Figure 1. Extrinsic and Intrinsic Reward Paradigm

(A) Subjects ($n = 33$) experienced both extrinsic and intrinsic rewards on each trial. A trial starts with subjects selecting from one or two available options each associated with an implicit extrinsic reward. One option leads to the larger reward (mean 50, SD 10) whereas the other leads to a lower reward (mean 25, SD 10) with a reversal every 19-23 trials. Four barriers then appear along the path to the outcome and a cursor appears at the bottom of the screen which automatically advances after a 1s delay. Subjects press left and right keys to navigate around barriers, constituting a form of skilled performance that can be intrinsically rewarding. Successfully avoiding a barrier turns it green whereas contact with a barrier turns it red. There is no financial penalty for contact with barriers nor financial benefit for avoiding them. Earnings depend only on the outcome delivered at the end of the trial. After every 2-3 trials, subjects report their current happiness by moving a cursor on a rating line.

(B) Probability of choice to the initial high-reward option averaged across subjects ($n = 33$) in black. Shaded areas correspond to SEM. Grey vertical bands represent intervals where probability reversals could occur.

(C) Happiness ratings across the task as reported by two example subjects.

RESULTS

Subjects completed two trial blocks while being scanned with fMRI. We first asked whether subjects could learn the reward contingencies (Figure 1B) and found that they could, making $84.8 \pm 5.6\%$ (mean \pm SD) of choices to the current high-reward option. Subjects were not penalized for contact with barriers, and thus actual performance was non-instrumental to the receipt of eventual monetary reward. We observed no correlation between earnings and how often subjects successfully avoided barriers (Spearman $\rho = 0.21$, $p = 0.24$). During debriefing, all 33 subjects reported that they believed there was no association between successful skilled performance and earnings.

Reports of affective state for example subjects are included in Figure 1C. On average, subjects reported being happier after receiving outcomes from the high- compared to low-reward option ($p < 0.001$), consistent with previous research (19, 20). On average, subjects reported also being happier when they navigated through the barriers without collisions compared to when they contacted at least one barrier ($p < 0.001$), suggesting that intrinsic rewards contribute to an influence on subjective affective state.

We next z-scored ratings within subjects so that those with greater variance in ratings did not have a disproportionate impact on our analysis. Consistent with analyses using non-normalized ratings, subjects reported greater average happiness after receiving high compared to low rewards ($t_{32} = 8.4$, $p < 0.001$, Figure 2A). Subjects also reported being happier after navigating through the maze without contacting any barriers compared to when they collided with at least one barrier ($t_{32} = 6.4$, $P < 0.001$, Figure 2A), consistent with an impact of intrinsic rewards. There was considerable variation across subjects in terms of how much extrinsic rewards and skilled performance contributed to momentary happiness (Figure 2B).

Computational model of affective dynamics

We next employed a previously established methodology (19, 20) to quantify the extent to which rewards impacted on the affective state of our participants. We considered influences that decay exponentially in time:

$$\text{Happiness}(t) = w_0 + w_1 \sum_{j=1}^t \gamma^{t-j} \text{Reward}_j$$

where t and j are trial numbers, w_0 is a baseline mood parameter, w_1 captures the influence of reward which is the z-scored reward outcome of the selected option on each trial, and $0 \leq \gamma \leq 1$ represents a forgetting factor that reduces the impact of distal relative to recent events. If this parameter is equal to 0, only the most recent reward outcome influences happiness. Parameters were first fit to non-normalized happiness ratings in each individual subject. The mean r^2 was 0.20 and the forgetting factor was 0.47 ± 0.32 (mean \pm SD). Consistent with previous findings (19, 20), happiness was significantly associated with the history of reward with positive happiness reward parameters on average ($p < 0.001$).

Likewise, consistent with previous findings during risky decision making (22), we found that baseline mood parameters, estimated while accounting for mood dynamics due to reward history, were negatively correlated with symptom severity assessed using

the Beck Depression Inventory (BDI-II) (23) (Spearman $\rho = -0.36$, $p = 0.04$). This result shows that happiness ratings during reinforcement learning, as we found previously during risky decision making (22), are tethered to the severity of depressive symptoms such that an affective set point, to which happiness returns to over time, is lower in individuals with a greater symptom load.

We also found baseline mood parameters were negatively correlated with apathy as measured by Apathy Evaluation Scale (AES) (24) (Spearman $\rho = -0.36$, $p = 0.04$) and behavioral apathy as assessed by the Apathy Motivation Index (AMI) (25) (Spearman $\rho = -0.34$, $p = 0.05$). We found no correlation between baseline mood parameters and the average staircased cursor speed (Spearman $\rho = 0.21$, $p = 0.24$), suggesting that the speed of the cursor was not significantly associated with persistent affective state.

We z-scored happiness ratings to better evaluate the contributions of extrinsic and intrinsic reward to affective state, thereby preventing individuals with greater rating variance from disproportionately affecting on the analyses. We expanded the model to include now an additional term that accounts also for influences pertaining to skilled performance:

$$\text{Happiness}(t) = w_1 \sum_{j=1}^t \gamma^{t-j} \text{Reward}_j + w_2 \sum_{j=1}^t \gamma^{t-j} \text{Performance}_j$$

where t and j are trial numbers, w_1 and w_2 capture the influence of task events related to reward and performance, respectively, and $0 \leq \gamma \leq 1$ represents a forgetting factor that reduces the impact of distal relative to recent events. Reward is the z-scored outcome of the selected option on each trial, and performance is the z-scored result of whether a barrier was contacted on each trial, assigning a 1 when no barriers were contacted and 0 if at least one barrier was contacted. This simple model explained a substantial amount of variance in happiness with $r^2 = 0.27 \pm 0.15$ (mean \pm SD, Figure 2C). Weights for both performance ($t_{32} = 5.79$, $p < 0.001$, Figure 2D) and reward ($t_{32} = 8.27$, $p < 0.001$, Figure 2D) were positive on average. The forgetting factor γ was 0.44 ± 0.31 (mean \pm SD), indicating that happiness depended on the past 4-5 trials on average.

In previous studies we found expectations of reward exerted a substantial influence on happiness (19, 20). In the current study, we used high- and low-reward distributions with minimal overlap to maximize learning accuracy. We also employed a staircase to keep skilled performance stable and at a similar level across individuals. These features render the current design unsuitable for quantifying the impact of expectations on happiness. Instead, we chose a design that maximized our power for quantifying individual differences in the relative subjective values of extrinsic and intrinsic rewards.

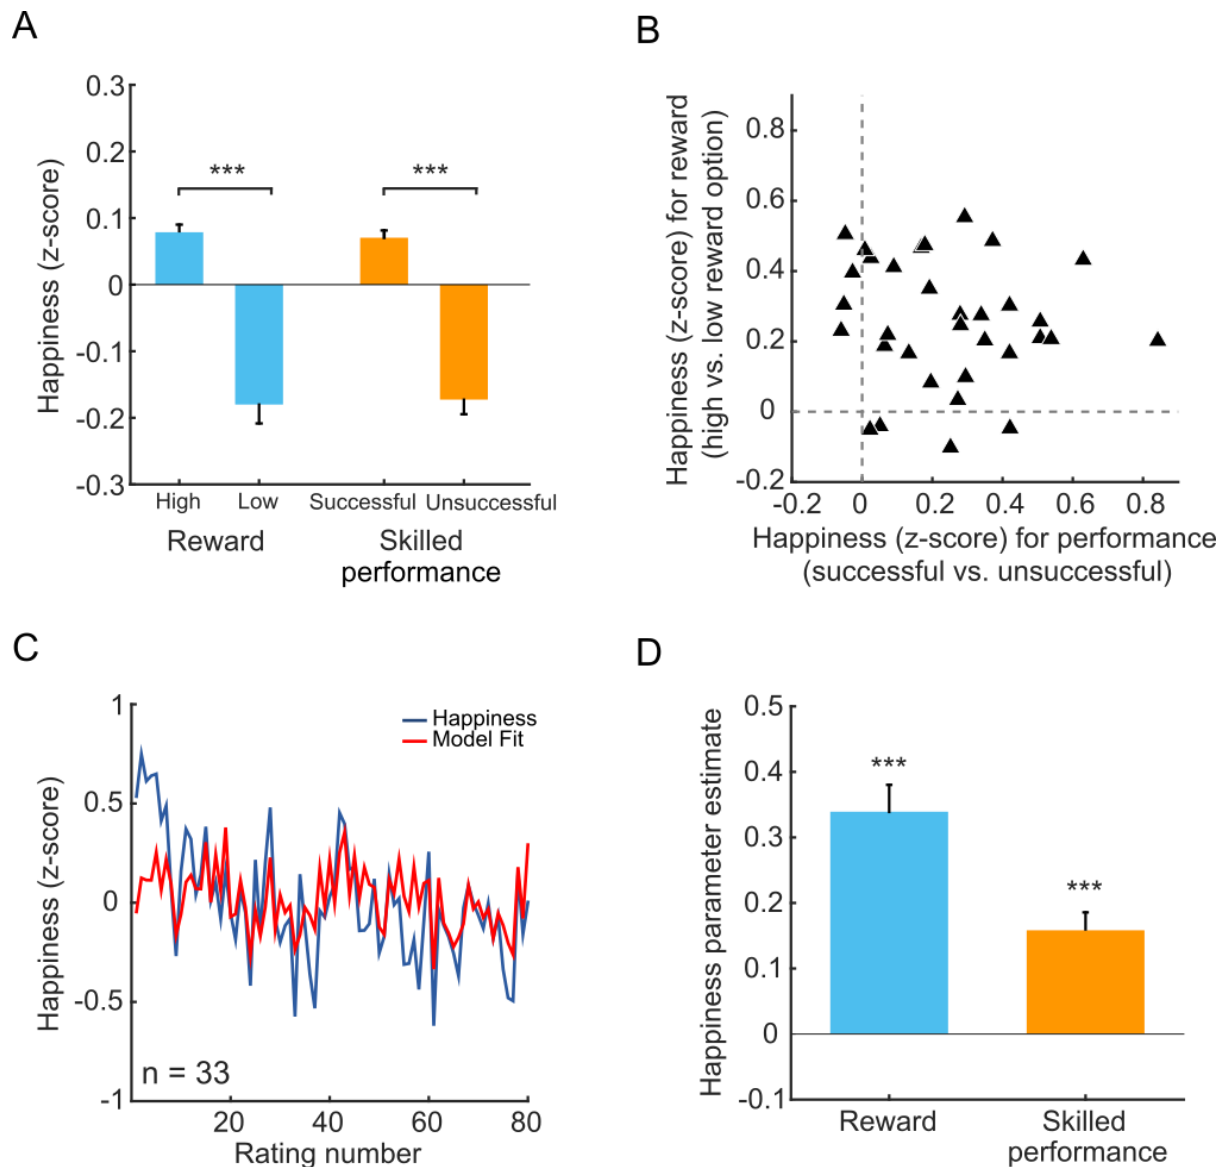


Figure 2. Computational modelling of affective dynamics

(A) Subjects were happier when they received a reward from high- compared to low-reward options ($t_{32} = 8.4$, $p < 0.001$, in blue). Subjects were happier on average when they navigated through the barriers without contacting with them, compared to when they contacted at least one barrier ($t_{32} = 6.4$, $p < 0.001$, in orange). *** $p < 0.001$.

(B) The majority of subjects (29 of 33) were happier after receiving a reward from a high-compared to low-reward option. The majority of subjects (29 of 33) were happier after successful compared to unsuccessful performance. There was no relationship between happiness for reward outcomes and happiness for skilled performance ($r = -0.11$, $p = 0.55$).

(C) Average happiness across subjects and model fit is displayed for the computational model ($n = 33$, mean $r^2 = 0.27$).

(D) Happiness was significantly related to the history of extrinsic rewards in the form of points converted to money ($t_{32} = 8.3$, $p < 0.001$) and also to the history of skilled performance, a proxy for intrinsic rewards ($t_{32} = 5.8$, $p < 0.001$). *** $p < 0.001$.

Model comparison (Table 1) shows that a model with happiness parameters for past rewards and performance outperformed models containing individual terms for reward (mean $r^2 = 0.19$) and performance (mean $r^2 = 0.08$) alone. These results show that the happiness of subjects in this task is, on average, dependent on both receipt of explicit rewards (e.g., money) and the non-instrumental experience of skilled performance.

Model	Parameters	Mean r^2	BIC	Δ BIC
Reward	2	0.19	-6747	507
Performance	2	0.08	-5907	1347
Reward and Performance	3	0.27	-7254	0
Reward and Performance (separate γ)	4	0.28	-7131	123

Table 1. Model Comparison Results

Bayesian Information Criterion (BIC) scores are summed across 33 subjects. The winning model (lowest BIC) was the model with both reward and performance having the same forgetting factor rather than a model where the influence of past reward and performance differs in their forgetting factor. Δ BIC refers to the difference in BIC between each model and the winning model. Ratings are z-scored to prevent individuals with greater rating variance from disproportionately influencing model comparison.

We found considerable variation across individuals in how much reward outcomes contributed to affective dynamics, even though subjects on average learned reward contingencies to a similar degree (Figure 3A). Despite performance being held constant due to staircasing of cursor speed (successful performance: $69.1 \pm 2.4\%$, mean \pm SD, Figure 3B), there was considerable variation also across individuals in how much non-instrumental performance influenced affective state. Many subjects showed a negligible impact of successful performance on affective state, despite a similar level of successful performance. Furthermore, learning choice accuracy was not correlated with either happiness reward parameters (Spearman $\rho = 0.3$, $p = 0.09$) or successful skilled performance (Spearman $\rho = 0.02$, $p = 0.92$).

The median cursor speed and the happiness performance parameter were positively correlated (Spearman $\rho = 0.45$, $p = 0.009$), consistent with individuals who found successful performance more intrinsically rewarding also being more intrinsically motivated to improve performance. This pattern was present despite the success of the staircasing procedure: there was no correlation between percent successful skilled performance and the happiness performance parameter (Spearman $\rho = 0.16$, $p = 0.39$). Subjects with happiness performance parameters greater than the median had an average cursor movement duration of 1980ms compared to 2188ms for subjects with happiness performance parameters below the median.

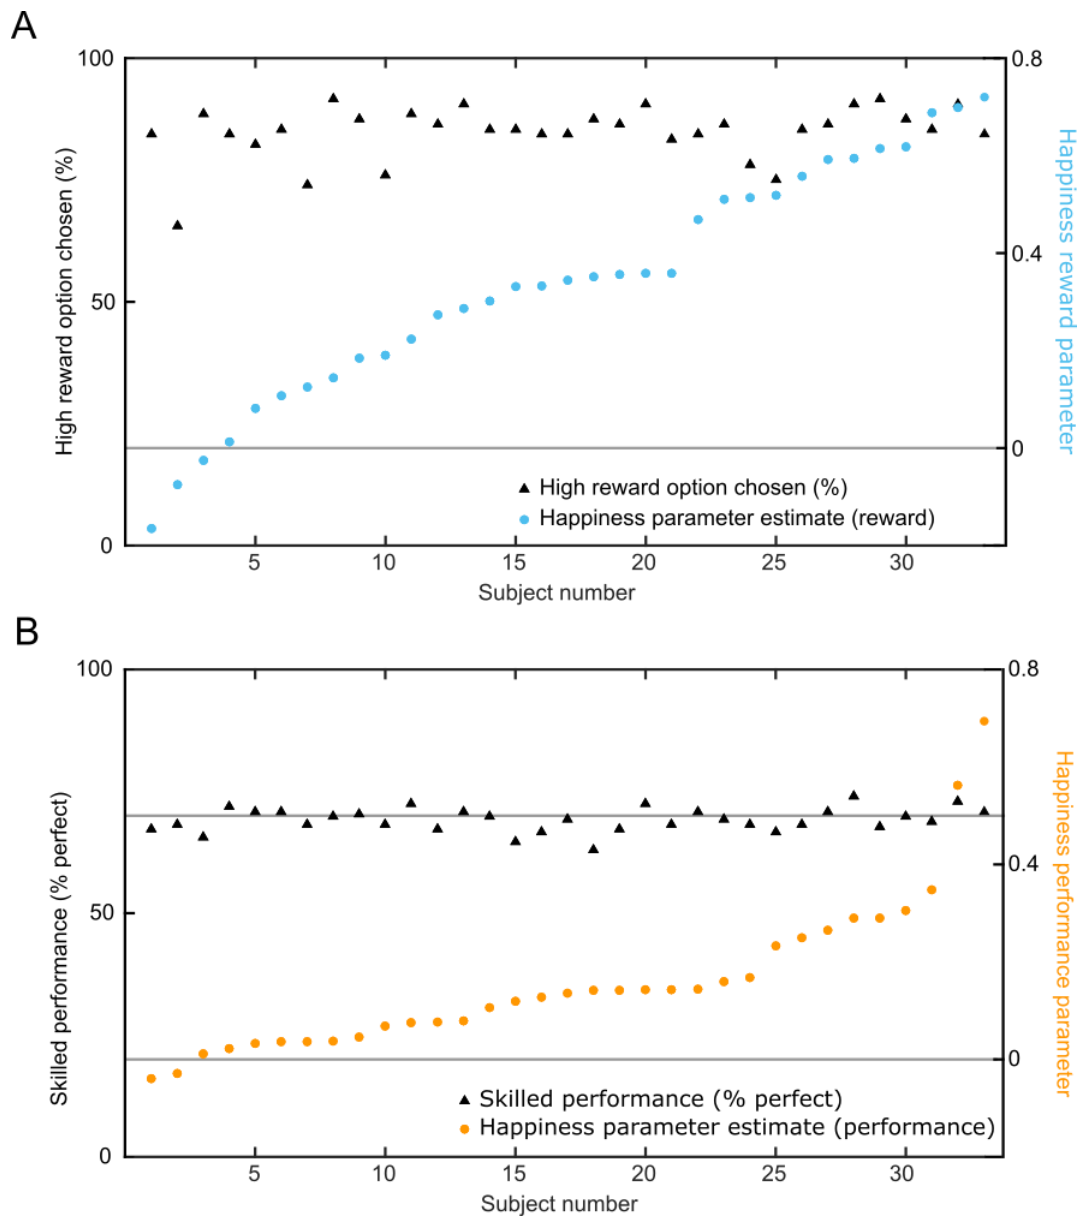


Figure 3. Computational model parameters and task behavior

(A and B) The contribution of reward to happiness varied across subjects despite a similar high choice accuracy across subjects. Despite titrating difficulty at the individual level to match performance across subjects around 70%, subjects displayed considerable variation in the degree to which performance impacted affective state as captured by the computational model.

Neural correlates of extrinsic and intrinsic rewards

Having established inter-individual variability in the impact of outcomes and performance on reported happiness, we next asked whether this variability was also predictive of neural responses to both rewards and performance. The experiment was separated into two scans and we first evaluated whether happiness model parameters were stable across scans. We found that both extrinsic ($\rho = 0.59$, $P < 0.001$) and intrinsic ($\rho = 0.48$, $P = 0.005$) reward computational parameters were positively correlated across the two scans. This was also true for difference between intrinsic and extrinsic reward parameters ($\rho = 0.40$, $p = 0.02$).

We regressed event-related activity on parametrically modulated task events to assess where brain activity areas relates to receipt of extrinsic and intrinsic rewards. We found an effect of reward magnitude at time of outcome in vmPFC (Figure 4A, top: -3, 38, -1; $t_{32} = 5.92$, $p < 0.05$ Family-Wise-Error (FWE) cluster-corrected at the whole-brain level), as well as an effect of successful skilled performance in an overlapping region of the vmPFC (Figure 4A, bottom: -3, 50, -1; $t_{32} = 4.24$, $p < 0.05$ FWE cluster-corrected).

The vmPFC is widely implicated in representation of subjective reward value. On this basis, we used an independent vmPFC mask from a meta-analysis of subjective value studies of extrinsic reward for further analysis (26). Within this ROI, we extracted weights for reward magnitude and skilled performance from each individual subject. We found that within this independent region-of-interest (ROI), BOLD activity was significantly associated with both reward magnitude ($t_{32} = 3.36$, $p = 0.002$) and skilled performance ($t_{32} = 2.90$, $p = 0.007$, Figure 4B).

Having established that neural responses in vmPFC are associated with both extrinsic and intrinsic rewards, we next examined whether neural responses were predicted by computational parameters estimated from individual affective dynamics. Across subjects, we found a positive relationship (Spearman $\rho = 0.52$, $p = 0.002$, Figure 4D) between the relative weights for extrinsic and intrinsic rewards in our happiness computational model and the relative effect sizes for neural responses in the vmPFC. For illustration purposes, we subdivided subjects into two groups comprising a group with higher happiness performance than reward parameters and a group with the opposite pattern. The group with higher performance than reward parameters showed greater vmPFC responses for skilled performance compared to the group with larger reward than performance parameters ($p = 0.003$, Figure 4C). These findings suggest that the pattern of momentary affective dynamics, reflecting the impact of both extrinsic and intrinsic rewards, and is mirrored at the level of vmPFC activity.

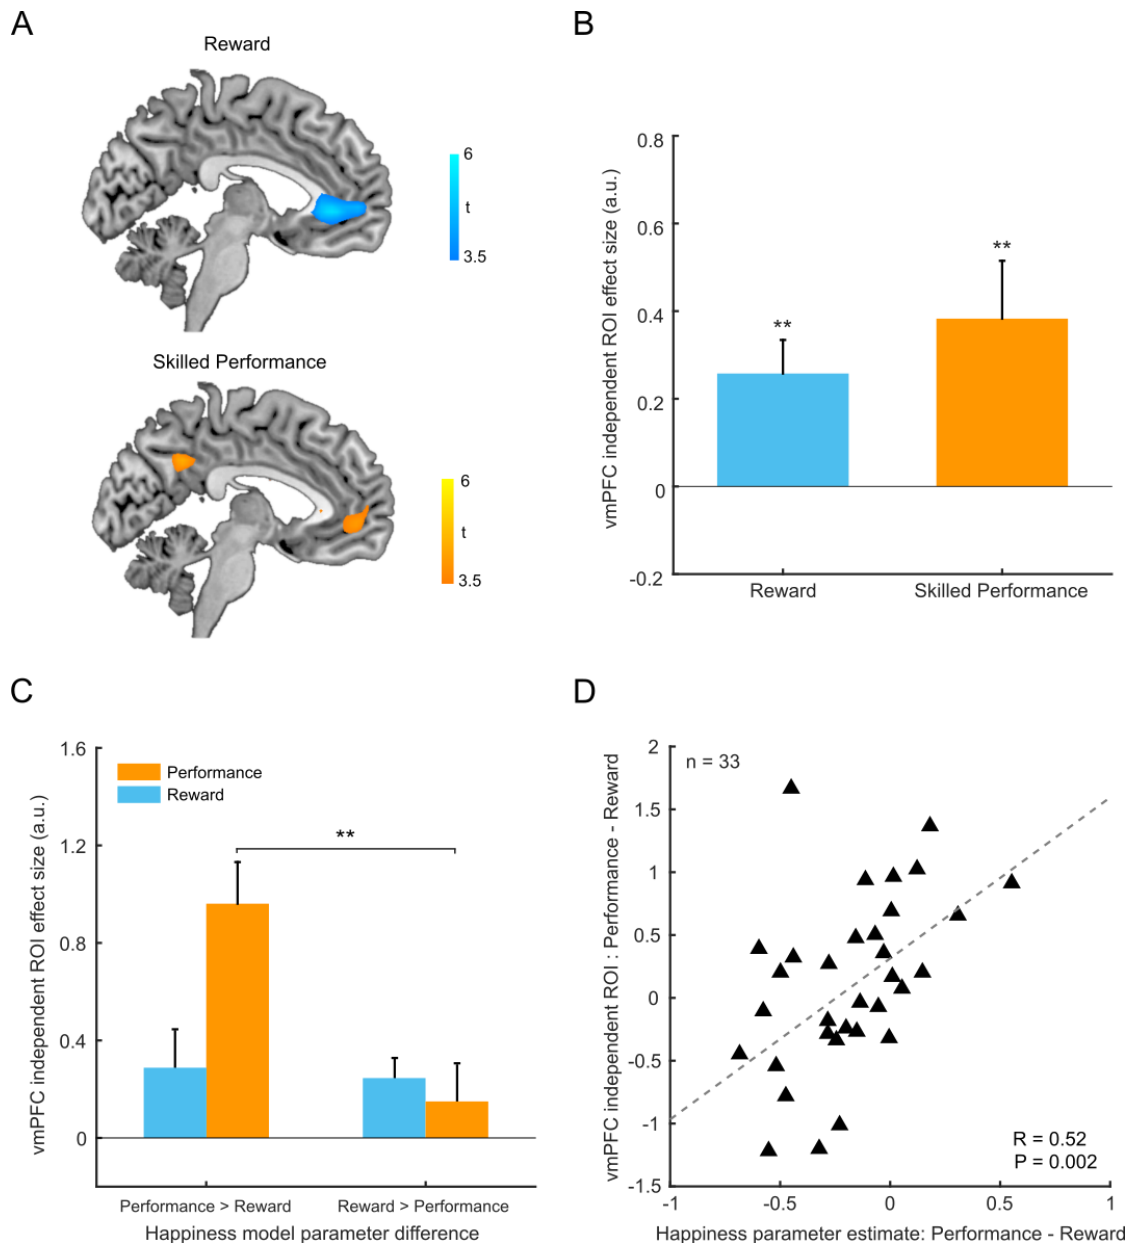


Figure 4. Relative affective impacts of reward and performance predict vmPFC activity

(A) *Top*. BOLD activity in vmPFC was parametrically modulated by reward magnitude (Peak: -3, 38, -1). *Bottom*. Bold activity in an overlapping region of vmPFC was modulated by trial-by-trial successful skilled performance (Peak: -3, 50, -1).

(B) An independent vmPFC ROI shows modulation by both reward magnitude and skilled performance (both $p < 0.01$).

(C) In the independent vmPFC ROI, subjects with higher performance than reward weights in the computational analysis of affective dynamics displayed stronger neural responses in the vmPFC for performance than subjects with higher reward than performance weights ($p < 0.01$).

(D) The difference between performance and reward weights in the happiness computational model predicted the difference in vmPFC neural responses for successful skilled performance relative to reward magnitude (Spearman $\rho = 0.52$, $p = 0.002$). * $p < 0.05$, ** $p < 0.01$.

DISCUSSION

Using experience sampling (27, 28) combined with functional neuroimaging, we show that intrinsic and extrinsic rewards contribute to affective dynamics (i.e., happiness). Recent studies demonstrate that computational approaches can quantify consistent relationships between subjective feelings and value-based decision-making (20, 21, 29, 30), including in relation to individual social preferences (31). Here, using the same computational approach, applied in the context of a reinforcement learning task, we show that momentary happiness is influenced by both objectively quantifiable rewards and by intrinsic rewards, where the latter involves experiences with no inherent worth.

The computational parameters we extract from affective dynamics enabled us to quantify, within a common value scale, the relative subjective reward value of intrinsic relative to extrinsic rewards. Our key finding here is our demonstration that vmPFC activity, a regions proposed to represent rewards in a common neural currency (32–34), is predicted by the relative weight of intrinsic and extrinsic reward extracted from affective dynamics.

While improvements in skilled performance can be enhanced by rewarding individuals for performance (35), holding performance constant across subjects allowed us to investigate how happiness varied independently of the level of skill individuals manifest in the task. We show that individuals, whose happiness was substantially influenced by intrinsic rewards, had increased vmPFC BOLD responses for successful versus unsuccessful skilled performance, relative to individuals whose happiness was influenced more by extrinsic rewards. The positive association between the happiness performance parameter and the staircased cursor speed also suggested that individuals who were either more proficient at the task or more intrinsically motivated to improve were also those whose momentary happiness was more strongly impacted by performance.

The vmPFC is known to represent the value of different types of goods, including food and juice (36, 37), money (38), aesthetic judgments (39, 40), and even perceived pleasantness (41). This suggests that vmPFC plays a central role in representing qualitatively different types of goods on a common scale, an operation that can facilitate making decisions between otherwise incommensurable goods (32–34). Our study builds on these prior results by now identifying an association between vmPFC BOLD activity and intrinsic rewards, here the experience of performing a skilled task without error. Whole-brain analysis showed that the representation of subjective intrinsic reward values involved an adjacent region in the vmPFC, anterior to the representation for extrinsic rewards, a finding that parallels a distinction between experienced and decision values previously mapped to anterior and posterior vmPFC, respectively (42).

The vmPFC has been demonstrated to play a role in affect with subjective emotional experiences elicited by images and pleasurable music leading to changes in both vmPFC BOLD activity and regional cerebral blood flow (43–45). Damage to the vmPFC can lead to aberrant emotional responses (46–48) and maladaptive decision-making in environments where emotional regulation may be useful (49, 50). Numerous studies suggest that subjective reward values are represented by vmPFC neural activity. Unfortunately, the constraints and expense of neuroimaging makes it impractical as an every-day tool for assessing individual values for non-market activities. The strong association between neural responses for intrinsic and extrinsic

rewards and computational parameters extracted from affective dynamics suggests that computational models, combined with experience sampling, provide a valid measure for the subjective reward value of experience.

Humans exhibit biases when it comes to predicting how future events are likely to impact on their emotional states, and are prone to making sub-optimal decisions by misjudging the hedonic consequences of options (9, 10, 15). Increasing subjective well-being is widely believed to be an appropriate societal goal (51), but these biases pose a difficulty for enacting policies that are likely to be successful. Additional factors such as social desirability bias (14) can decrease the reliability of self-reported values when an individual's assessment of a hypothetical experience or good, such as the availability of public parks, differs from prevailing social norms. An advantage of our method is that it can be applied to any repeatable experience without a need to probe people explicitly about the content of those experiences, reducing biases associated with social desirability. For example, asking participants to report a negative social emotion such as envy is likely to be unreliable or problematic in some populations. Using our approach, 'envy' parameters can be extracted from affective dynamics in relation to social inequality that predict monetary allocations in an independent dictator game (31). Similarly, affective dynamics reflect depressive symptoms (22) and show consistent relationships to reward in the lab and outside the lab in anonymous participants who did not interact with an experimenter (20). Our computational approach combined with experience sampling can be widely applied to more naturalistic settings such as a corporate workplace, identifying factors important for employee well-being.

Over a century ago, Francis Edgeworth described an idealized instrument, which he called a hedonometer, for 'continually registering the height of pleasure experienced by an individual' (52). Here, we introduce a kind of 'computational hedonometer' that has a distinct advantage over Edgeworth's hypothetical hedonometer in that it mathematically quantifies the relative contributions of different factors to an affective state, including the relative values of intrinsic and extrinsic rewards. We validate our computational tool using objective neural measurements, suggesting that computational parameters can capture implicit values for abstract goods and experiences that may be otherwise challenging to accurately quantify. The combination of computational modelling and experience sampling provides a useful tool that can be used to design and evaluate policies to increase subjective well-being.

EXPERIMENTAL PROCEDURES

Participants

37 healthy young adults (age: 25.8 ± 4.7 , mean \pm SD; 8 males, 29 females) were recruited through the University College London (UCL) Psychology Subject Database. Subjects were screened to ensure no history of neurological or psychiatric disorders. Four subjects were excluded due to excessive head movement during scanning, leaving a total of 33 subjects (age: 26.1 ± 4.9 ; 8 males, 25 females). The study was approved by the UCL research ethics committee, and all subjects gave written informed consent.

Study Design

Subjects completed the experiment at the Wellcome Centre for Human Neuroimaging at UCL in an appointment that lasted approximately 90 minutes. Stimuli were presented in MATLAB (MathWorks, Inc.) using Cogent 2000. The layout of each trial resembled a T-Maze (53). On each trial, subjects selected a blue or magenta box, one of which resulted in 50 points on average and the other which resulted in 25 points on average. The standard deviation of points received for each box was 10. Points assigned based on draws from Gaussian distributions. Every 19-23 trials, a reversal occurred where the box that previously contained the higher number of points on average now contained a lower number of points and vice versa. On half of the trials, subjects were afforded a free choice. For the remaining half, subjects were only presented with a single option. After a choice was made, the chosen option was indicated and four barriers appeared on the screen along with a small cursor at the bottom of the screen. Following a 1s delay, the cursor automatically advanced along the path to the outcome. Subjects were able to control the horizontal position of the cursor to avoid colliding with barriers. If they passed a barrier without colliding with it, the barrier turned green. Contact with a barrier turned it red and provided immediate feedback about performance. Crucially, the subjects' final payment depended only on the number of points accumulated across the experiment and not their ability to quickly navigate past barriers. After the cursor had entered the chosen box, the outcome was displayed for 800ms after a 1.5s delay. Total cumulative points were displayed on the top right of the screen throughout the experiment. Subjects were presented with the question, "How happy are you at this moment?" after every 2-3 trials. After a 1s delay period, a rating line appeared with a cursor at the midpoint and subjects had 4s to move a cursor along the scale with button presses. The left end of the line was labelled "very unhappy" and the right end of the line was labelled "very happy".

Staircase Procedure

To ensure that differences in affective responses were not due to skill-related differences in how often each subject collided with barriers, we used a standard staircase procedure called the Parametric Estimation by Sequential Testing (PEST) (54). This procedure calibrated the speed at which the cursor moved for every subject such that they did not contact the barriers on approximately 70% of trials. This calibration was carried out over 60 trials prior to the start of the task in the scanner. Continuation of the procedure during the task allowed small adjustments (e.g., fatigue) to maintain consistent successful skill performance.

Questionnaire Measures

Subjects were administered the Beck Depression Inventory (BDI-II) (23), Apathy Evaluation Scale (AES) (24), and Apathy Motivation Index (AMI) (25).

Image Acquisition

MRI scanning took place at the Wellcome Centre for Human Neuroimaging at UCL using a Siemens Prisma 3-Tesla scanner equipped with a 64-channel head coil. Functional images were acquired with a gradient echo T2*-weighted echo-planar sequence with whole-brain coverage. Each volume consisted of 48 slices with 3mm isotropic voxels [repetition time (TR): 3.36s; echo time (TE): 30ms; slice tilt: 0°] in ascending order. A field map [double-echo FLASH, TE1 = 10ms, TE2 = 12.46ms] with 3mm isotropic voxels (whole-brain coverage) was also acquired for each subject to correct the functional images for any inhomogeneity in magnetic field strength. Subsequently, the first 6 volumes of each run were discarded to allow for T1 saturation effects. Structural images were T1-weighted (1 x 1 x 1 mm resolution) images acquired using a MPRAGE sequence.

ACKNOWLEDGMENTS

We thank Tobias Hauser, Matilde Vaghi, and Rachel Bedder for helpful comments. B.C. is a predoctoral fellow of the International Max Planck Research School on Computational Methods in Psychiatry and Ageing Research. The participating institutions are the Max Planck Institute for Human Development and the University College London (UCL). B.C. is also supported by a scholarship from the Singapore Institute of Management. R.J.D. holds a Wellcome Trust Investigator Award (098362/Z/12/Z). R.B.R. is supported by a Medical Research Council Career Development Award (MR/N02401X/1) and a 2018 NARSAD Young Investigator Grant (27674) from the Brain and Behavior Research Foundation, P&S Fund. The Max Planck UCL Centre is a joint initiative supported by UCL and the Max Planck Society. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z).

REFERENCES

1. A. J. Oswald, S. Wu, Objective confirmation of subjective measures of human well-being: Evidence from the U.S.A. *Science* **327**, 576–579 (2010).
2. Y. Chida, A. Steptoe, Positive psychological well-being and mortality: A quantitative review of prospective observational studies. *Psychosomatic Medicine* **70**, 741–756 (2008).
3. A. Steptoe, A. Deaton, A. A. Stone, Subjective wellbeing, health, and ageing. *The Lancet* **385**, 640–648 (2015).
4. K. W. Davidson, E. Mostofsky, W. Whang, Don't worry, be happy: Positive affect and reduced 10-year incident coronary heart disease: the Canadian Nova Scotia Health Survey. *European Heart Journal* **31**, 1065–1070 (2010).
5. A. J. Oswald, E. Proto, D. Sgroi, Happiness and productivity. *Journal of Labor Economics* **33**, 789–822 (2015).
6. L. H. Pelled, K. R. Xin, Down and out: An investigation of the relationship between mood and employee withdrawal behavior. *Journal of Management* **25**, 875–895 (1999).
7. S. J. Peterson, F. Luthans, B. J. Avolio, F. O. Walumbwa, Z. Zhang, Psychological capital and employee performance: A latent growth modeling approach. *Personnel Psychology* **64**, 427–450 (2011).
8. P. Dolan, M. P. White, How can measures of subjective well-being be used to inform public policy? *Perspectives on Psychological Science* **2**, 71–85 (2007).
9. T. Meyvis, R. K. Ratner, J. Levav, Why don't we learn to accurately forecast feelings? How misremembering our predictions blinds us to past forecasting errors. *Journal of Experimental Psychology. General* **139**, 579–589 (2010).
10. T. D. Wilson, D. T. Gilbert, Affective forecasting: Knowing what to want. *Current Directions in Psychological Science* **14**, 131–134 (2005).
11. D. T. Gilbert, T. D. Wilson, Propection: Experiencing the future. *Science* **317**, 1351–1354 (2007).
12. C. K. Morewedge, E. C. Buechel, Motivated underpinnings of the impact bias in affective forecasts. *Emotion* **13**, 1023–1029 (2013).
13. H. Plassmann, J. O'Doherty, A. Rangel, Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience* **27**, 9984–9988 (2007).
14. T. F. Van de Mortel, Faking it: Social desirability response bias in self-report research. *The Australian Journal of Advanced Nursing* **25**, 40 (2008).
15. E. K. Nisbet, J. M. Zelenski, Underestimating nearby nature: Affective forecasting errors obscure the happy path to sustainability. *Psychological Science* **22**, 1101–1106 (2011).

16. D. Keatley, D. D. Clarke, M. S. Hagger, The predictive validity of implicit measures of self-determined motivation across health-related behaviours. *British Journal of Health Psychology* **18**, 2–17 (2013).
17. C. Levesque, K. J. Copeland, R. A. Sutcliffe, Conscious and nonconscious processes: Implications for self-determination theory. *Canadian Psychology* **49**, 218–224 (2008).
18. T. H. B. FitzGerald, B. Seymour, R. J. Dolan, The role of human orbitofrontal cortex in value comparison for incommensurable objects. *Journal of Neuroscience* **29**, 8388–8395 (2009).
19. R. B. Rutledge, N. Skandali, P. Dayan, R. J. Dolan, Dopaminergic modulation of decision making and subjective well-being. *Journal of Neuroscience* **35**, 9811–9822 (2015).
20. R. B. Rutledge, N. Skandali, P. Dayan, R. J. Dolan, A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences USA* **111**, 12252–12257 (2014).
21. F. Vinckier, L. Rigoux, D. Oudiette, M. Pessiglione, Neuro-computational account of how mood fluctuations arise and affect decision making. *Nature Communications* **9** (2018).
22. R. B. Rutledge, *et al.*, Association of neural and emotional impacts of reward prediction errors with major depression. *JAMA Psychiatry* **74**, 790 (2017).
23. A. T. Beck, R. A. Steer, R. Ball, W. Ranieri, Comparison of Beck Depression Inventories -IA and -II in psychiatric outpatients. *Journal of Personality Assessment* **67**, 588–597 (1996).
24. R. S. Marin, R. C. Biedrzycki, S. Firinciogullari, Reliability and validity of the apathy evaluation scale. *Psychiatry Research* **38**, 143–162 (1991).
25. Y.-S. Ang, P. Lockwood, M. A. J. Apps, K. Muhammed, M. Husain, Distinct subtypes of apathy revealed by the apathy motivation index. *PLOS ONE* **12**, e0169938 (2017).
26. O. Bartra, J. T. McGuire, J. W. Kable, The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage* **76**, 412–427 (2013).
27. D. Kahneman, A. B. Krueger, D. A. Schkade, N. Schwarz, A. A. Stone, A survey method for characterizing daily life experience: The day reconstruction method. *Science* **306**, 1776–1780 (2004).
28. H. T. Reis, S. L. Gable, “Event-sampling and other methods for studying everyday experience” in *Handbook of Research Methods in Social and Personality Psychology*, (Cambridge University Press, 2000), pp. 190–222.
29. E. Eldar, C. Roth, P. Dayan, R. J. Dolan, Decodability of reward learning signals predicts mood fluctuations. *Current Biology* **28**, 1433-1439.e7 (2018).
30. E. Eldar, R. B. Rutledge, R. J. Dolan, Y. Niv, Mood as representation of momentum. *Trends in Cognitive Sciences* **20**, 15–24 (2016).

31. R. B. Rutledge, A. O. de Berker, S. Espenhahn, P. Dayan, R. J. Dolan, The social contingency of momentary subjective well-being. *Nature Communications* **7**, 11825 (2016).
32. V. S. Chib, A. Rangel, S. Shimojo, J. P. O'Doherty, Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *Journal of Neuroscience* **29**, 12315–12320 (2009).
33. D. J. Levy, P. W. Glimcher, Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *Journal of Neuroscience* **31**, 14693–14707 (2011).
34. D. J. Levy, P. W. Glimcher, The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology* **22**, 1027–1038 (2012).
35. S. K. Sugawara, S. Tanaka, S. Okazaki, K. Watanabe, N. Sadato, Social rewards enhance offline improvements in motor skill. *PLOS ONE* **7**, e48174 (2012).
36. T. A. Hare, J. Malmaud, A. Rangel, Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *Journal of Neuroscience* **31**, 11077–11087 (2011).
37. C. Padoa-Schioppa, Orbitofrontal cortex and the computation of economic value. *Annals of the New York Academy of Sciences* **1121**, 232–253 (2007).
38. B. D. Martino, D. Kumaran, B. Seymour, R. J. Dolan, Frames, biases, and rational decision-making in the human brain. *Science* **313**, 684–687 (2006).
39. T. Jacobsen, R. I. Schubotz, L. Höfel, D. Y. v. Cramon, Brain correlates of aesthetic judgment of beauty. *NeuroImage* **29**, 276–285 (2006).
40. H. Kawabata, S. Zeki, Neural correlates of beauty. *Journal of Neurophysiology* **91**, 1699–1705 (2004).
41. H. Plassmann, J. O'Doherty, B. Shiv, A. Rangel, Marketing actions can modulate neural representations of experienced pleasantness. *Proceedings of the National Academy of Sciences USA* **105**, 1050–1054 (2008).
42. D. V. Smith, *et al.*, Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *Journal of Neuroscience* **30**, 2490–2495 (2010).
43. A. J. Blood, R. J. Zatorre, Intensely pleasurable responses to music correlate with activity in brain regions implicated in reward and emotion. *Proceedings of the National Academy of Sciences USA* **98**, 11818–11823 (2001).
44. A. Winecoff, *et al.*, Ventromedial prefrontal cortex encodes emotional value. *Journal of Neuroscience* **33**, 11032–11039 (2013).
45. D. H. Zald, D. L. Mattson, J. V. Pardo, Brain activity in ventromedial prefrontal cortex correlates with individual differences in negative affect. *Proceedings of the National Academy of Sciences USA* **99**, 2450–2454 (2002).
46. J. Hiser, M. Koenigs, The multifaceted role of ventromedial prefrontal cortex in emotion, decision-making, social cognition, and psychopathology. *Biological Psychiatry* **83**, 638–647 (2018).

47. M. Koenigs, *et al.*, Damage to the prefrontal cortex increases utilitarian moral judgements. *Nature* **446**, 908–911 (2007).
48. D. H. Zald, C. Andreotti, Neuropsychological assessment of the orbital and ventromedial prefrontal cortex. *Neuropsychologia* **48**, 3377–3391 (2010).
49. M. Grossman, *et al.*, The role of ventral medial prefrontal cortex in social decisions: Converging evidence from fMRI and frontotemporal lobar degeneration. *Neuropsychologia* **48**, 3505–3512 (2010).
50. J. Spaniol, F. Di Muro, E. Ciaramelli, Differential impact of ventromedial prefrontal cortex damage on “hot” and “cold” decisions under risk. *Cognitive, Affective, & Behavioral Neuroscience* **19**, 477–489 (2019).
51. OECD, *How's life? 2017: Measuring well-being* (OECD, 2017) https://doi.org/10.1787/how_life-2017-en (December 9, 2019).
52. F. Y. Edgeworth, *Mathematical psychics: An essay on the application of mathematics to the moral sciences* (Kegan Paul, 1881).
53. M. W. Howe, P. L. Tierney, S. G. Sandberg, P. E. M. Phillips, A. M. Graybiel, Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature* **500**, 575–579 (2013).
54. M. M. Taylor, C. D. Creelman, PEST: Efficient estimates on probability functions. *The Journal of the Acoustical Society of America* **41**, 782–787 (1967).