Different Affective Learning Systems Contribute to the Accumulation of Assets and Debt

Brian Knutson[1], Camelia M. Kuhnen[2], Gregory R. Samanez-Larkin[1]

[1] Department of Psychology, Stanford University

[2] Kellogg School of Management, Northwestern University

Abstract

In an experimental setting that combines a financial investment task, functional brain imaging and credit report data regarding our participants' finances, we find that individuals who learn better about gains have more real-life assets while individuals who learn better about losses have less debt, and document that brain areas related to emotion processing are responsible for incorporating financial information into choice. The results are robust to the inclusion of cognitive (i.e., memory, cognitive flexibility, numeracy) and demographic (i.e., age, sex, ethnicity, education) controls. The sensitivity of the medial prefrontal cortex to expected value during gain learning and the sensitivity of the anterior insula to expected value during loss learning are predictors of individual differences in performance on the financial choice task. Moreover, within individuals, learning from gains and learning from losses are not correlated. These findings suggest that distinct systems in the emotion-related areas of the brain guide gain and loss learning and influence real-life financial outcomes.

Why do some individuals accumulate assets while others gather debt? Environmental factors related to family socioeconomic status and inheritance undoubtedly play critical roles in determining life financial outcomes (Webley et al., 2001). But some who are born into poverty eventually amass riches, while others who are born into wealth squander their endowments. Do life financial outcomes depend solely upon chance circumstances, or might individual differences exert a subtle yet persistent influence over time?

Individuals reliably differ in both cognitive and emotional capacities (Bouchard, 1994). While some evidence suggests that individual differences in cognitive capacities may influence financial preferences and outcomes (Agarwal et al., 2007; Burks et al., 2009; Mischel et al., 1989), minimal research has focused on the influence of individual differences in emotional capacities. An emerging body of human research indicates that individuals' anticipatory affect — or emotional states prior to significant outcomes — can influence subsequent choice. Specifically, neural activity associated with positive arousal (i.e., positive and aroused feelings like "excitement") in the nucleus accumbens (NAcc) and medial prefrontal cortex (MPFC) precedes acceptance of risky gambles and purchase of products. On the other hand, neural activity associated with negative arousal (e.g., negative and aroused feelings like "anxiety") in the anterior insula conversely precedes rejection of risky gambles and refusal to purchase products (Knutson and Greer, 2008; Kuhnen and Knutson, 2005; Preuschoff et al., 2006). Contrary to the traditional assumption that only one system processes value (e.g., as implied by reinforcement learning theory), these findings suggest that distinct neurophysiological circuits process potential gains and losses.

Recent neuroimaging evidence has also linked activity in these anticipatory affective circuits to learning. For instance, whereas activation of the NAcc and associated ventral tegmental area have been associated with learning about gains (Adcock et al., 2006; McClure et al., 2003; O'Doherty et al., 2003; Pessiglione et al., 2006; Wittmann et al., 2005), activation of the anterior insula and amygdala have been associated with learning about losses (Buchel et al., 1998; Gottfried et al., 2003; LaBar et al., 1998; Samanez-Larkin et al., 2008). While this collected evidence is suggestive, no studies have explicitly examined whether individual differences in the recruitment of different anticipatory affective systems translate into faster learning about gains and losses.

Surprisingly, neuroimaging research on affective learning has largely ignored individual differences in performance. Although a number of studies have implicated prefrontal and striatal regions in the representation of reward predictions (or expected value) and reward prediction errors (Berns et al., 2001; Breiter et al., 2001; Knutson, et al 2005; McClure et al., 2003; O'Doherty et al., 2003; 2004; Rodriguez, et al 2006), very few studies have demonstrated a link between activation in these regions and individual differences in successful learning (Pessiglione et al 2008; Samanez-Larkin et al 2008; Schonberg et al 2007; Tobler et al, 2006). No prior studies have directly linked continuous measures of anticipatory and consummatory neural activation during affective learning with individual differences in gain and loss learning performance.

It is possible that the cumulative influence of individual differences in the neural sensitivity of these regions and associated affective learning biases may influence different life financial outcomes. Instead of integrating assets and debt into a combined measure of personal wealth over time, individuals may instead maintain separate "mental accounts" (Thaler, 1980), based on a fundamental distinction between gains and losses (Kahneman and Tversky, 1979). Thus, individual differences in learning about gains and losses might not only bias immediate choices but also long-term financial outcomes.

This study aimed to assess individual differences in affective learning (with a combination of neuroimaging, computational, behavioral, and self-report measures) and to determine their relationship to significant life financial outcomes (i.e., assets and debt). Based on an anticipatory affect model (Knutson and Greer, 2008), we predicted that neurobehavioral variables related to gain learning would promote acquisition of assets, while neurobehavioral variables related to loss learning would instead promote avoidance of debt. To test these predictions, we validated behavioral measures of gain and loss learning with brain imaging data on a subset of subjects, validated self-reported measures of assets and debt with objective credit report data on a subset of subjects, and assessed and controlled for potential demographic and cognitive confounds.

Behavioral measures of gain learning and loss learning were elicited with probabilistic learning tasks that included separate gain and loss conditions (Figure 1). Subjects' overall selection of the high versus low probability gain cue indexed "gain correct choices," while their overall selection of the low versus high probability loss cue indexed "loss correct choices." Supporting the optimality of these correct choice measures, gain correct choices were robustly associated with optimal gain learning, while loss correct choices were robustly but separately associated with optimal loss learning (estimated with a Bayesian model; see Supplementary Procedure 1 and Table S1). Since correct choices and optimal learning indices yielded similar results, simpler percent correct choice indices were used in primary analyses.

Results and Discussion

Regression models in Table 1 tested the key predictions that gain learning might account for accumulated assets and loss learning might account for accumulated debt. Supporting this conjecture, we found that gain and loss correct choices ($r=.09$, n.s.) were uncorrelated within subject (Figure 2) suggesting that they represent different abilities and therefore may influence different aspects of individuals' economic behaviors and outcomes.

We included both gain and loss learning as determinants in models predicting assets and debt. Since assets and debt were only moderately positively correlated ($r=.21$, $p<.05$; suggesting partial independence), assets were included in models that accounted for debt and vice-versa. The first reduced regression model revealed a significant positive association of gain correct choices (but not loss correct choices) with assets, and a full model (including cognitive and demographic variables) yielded similar results. Of the potential confounds, only age and ethnicity (i.e., Hispanic) were significantly associated with assets. Conversely, the second reduced regression model revealed a significant negative association of loss correct choices (but not gain correct choices) with debt, and a full model yielded similar results. Of the potential confounds, only ethnicity (i.e., Hispanic) was significantly associated with debt.

To examine whether distinct neural systems promoted gain correct choices versus loss correct choices, functional magnetic resonance imaging (FMRI) data was acquired in approximately half of the sample (n=40) during the task. One participant's data had to be discarded due to excessive head motion, leaving 39 subjects in the brain imaging sample. Regression analyses showed that individuals whose MPFC activation closely tracked gain predictions made more gain correct choices, while individuals whose insular activation more closely tracked loss predictions made more loss correct choices (Table 2). Controlling for age did not significantly change these associations. Therefore, we document brain activation patterns that can predict individual differences in the ability to learn information about losses and the ability to learn about gains.

Moreover, as shown in Table 3 we found that the brain's sensitivity to gain information and sensitivity to loss information are not related within individuals, further supporting the result in Figure 2 that gain learning and loss learning are different abilities, and are not correlated within individuals.

These results extend prior evidence for a relationship between striatal and medial prefrontal regions during gain learning and insular regions during loss learning (Pessiglione et al 2006), by providing direct evidence for a relationship between learning-related neural signals and individual differences in both gain and loss learning performance. Although prior studies have documented differences between successful and unsuccessful learners in striatal (Schonberg et al 2007) and orbitofrontal activation

(Tobler et al 2006) in supraliminal learning tasks and a relationship between reward earnings and expected value (but not prediction error) signals in extrastriate regions in a subliminal task (Pessiglione et al 2008), to the best of our knowledge no prior studies have directly linked continuous measures of anticipatory and consummatory neural activation during affective learning with individual differences in gain and loss learning performance.

One surprising result is that neural representations of expected value at anticipation were more highly correlated with learning performance than neural representations of prediction error at outcome. We acknowledge that there is a relationship between expected value and prediction errors in the context of the reinforcement learning model used here. Although expected values are based on the integration of prior prediction errors, the temporal proximity of representations of expected value at anticipation may explain why this signal exerted a stronger influence on behavioral outcomes.

To verify that the self-reported life financial outcomes reflected objective financial status, credit reports were obtained in approximately half of the sample (n=37), including measures of overall credit and percent credit used. The available credit amount was associated with self-reported assets but not debt, while the percent of credit used was more robustly associated with self-reported debt than assets, suggesting that self-reported assets and debt reflected distinct aspects of objective finances (Table 4 and Supplementary Procedure 2).

Together, these findings not only implicate distinct anticipatory affective systems in gain learning versus loss learning, but further illustrate selective associations of gain learning with increased assets and of loss learning with decreased debt. Cognitive and demographic variables could not account for these associations. Further, affective learning variables were not significantly associated with a more traditional unitary finance measure of wealth — the debt to asset ratio – in unreported regressions. Instead of implicating a single mechanism for value assessment, these findings are consistent with an anticipatory affect model, in which distinct systems assess gain and loss (Knutson and Greer, 2008). Moreover, the findings imply that anticipatory affect may not only influence immediately subsequent learning, but also may have a systematic and cumulative influence on significant life financial outcomes (Figure 3).

This initial demonstration that individual differences in emotional capacities may influence life financial outcomes supports a causal account in which gain learning promotes asset acquisition while loss learning promotes debt avoidance. Although these associations are predicted, specific, and robust, they are not causal. According to one alternative interpretation, other individual difference variables related to cognition and demographics might play more prominent roles in life financial outcomes. Although these variables may play a role in life financial outcomes, they did not reduce the predicted associations in the present sample after being entered into the regression model (Table 1). An alternative reverse causality account might hold that greater assets increase gain learning, while higher debt increases loss learning. Based on the economic notion of diminishing marginal returns, however, it seems unlikely that increased assets would increase (rather than decrease) individuals' sensitivity to gains (Bernoulli, 1738 / 1954). Additionally, gain learning and loss learning measures showed moderate test-retest reliability (r~.50) and did not significantly change as a function of repeated testing in an independent sample, consistent with stability over time (Supplementary Procedure 3). Only future longitudinal studies will be able to determine whether gain learning and loss learning causally influence future life financial outcomes.

These findings uniquely span multiple levels of analysis and timescales, offering a number of advances over previous research. Specifically, we recruited a community sample with significant assets and debt rather than a sample of convenience, demonstrated selective and robust functional dissociations, validated behavioral measures of gain and loss learning with neuroimaging data, validated measures of life financial outcomes with credit report data, and assessed and controlled for other potentially important individual difference variables.

In conclusion, beyond external social and economic forces, these findings suggest that individual differences may have a systematic influence on life financial outcomes. Specifically, affective learning capacities may exert a previously undocumented influence in which sensitivity to gains promotes approach towards financial opportunities, while sensitivity to losses instead promotes avoidance of financial threats. In addition to influencing single choices, these mechanisms may promote selection of different choice environments that highlight either the presence of financial opportunities or the absence of threats. Although affective learning can facilitate dynamic adjustments to environmental events, it may also exacerbate choice biases over time. Fortunately, identification of these biases may resolve targets for intervention – either on the part of individuals or their advisors.

Experimental Procedures

Subjects. A survey research firm initially contacted individuals who were representative of San Francisco Bay Area residents. 82 of the contacts (mean age=55, SD=18, range=20–85; 55 male) participated. Subjects received fixed compensation of $20 per hour, as well as cash equivalent to their total earnings in the task. Subjects were also informed that they could lose money on the task and that any losses would be deducted from their total earnings. To associate distinct neural circuits with correct gain versus loss choices, functional magnetic resonance imaging (FMRI) data was acquired on approximately half of the sample (n=40) as they engaged in the experimental task. To validate self-reported assets and debt, credit reports were acquired for nearly half of the sample (n=37).

Monetary Incentive Learning (MIL) Task. The MIL task was modified from conventional reinforcement learning tasks to separately assess learning about gains and losses (Kim et al., 2006; Pessiglione et al., 2006; Samanez-Larkin et al., 2007) . Subjects saw and chose between one of three pairs of fractal cues (gain acquisition, loss avoidance, or neutral) in each run of 12 trials per condition for a total of 36 trials. After choosing one of the cues from each pair, subjects saw the outcome associated with their choice. On average, one of the cues yielded a better outcome, while the other yielded a worse outcome. Specifically, in gain pairs, the better cue had a higher probability of returning gains (66% +$1.00 and 33% +$0.00) than the worse cue (33% +1.00 and 66% +$0.00), while in loss pairs, the better cue had a higher probability of returning nonlosses (66% –$0.00 and 33% –$1.00) than the worse cue (33% –$0.00 and 66% –$1.00), and in the neutral condition, choice of either cue had no impact on outcomes (100% $0.00). Within each pair, cues appeared randomly and with equal frequency on the left or right side of the screen. Cue pairings with better or worse outcomes were randomly assigned by the computer at the start of testing and counterbalanced across subjects. To minimize memory-related interference, different cue pairs were used for practice and experimental sessions. Subjects were explicitly informed about cue probabilities before the practice session and instructed to try to maximize their earnings throughout the experiment.

Indices of gain learning and loss learning performance were calculated by counting the percentage of choices that matched the "correct" (optimal) choice in gain and loss conditions, excluding the first trial in each session. Optimality of gain and loss choices were defined as the fraction of trials when the subject made the correct Bayesian choice, excluding trials in which either asset had equal chance of being optimal (Supplement Procedure 1 and Table S1).

Life Financial Outcomes. Assets and debt were assessed by self-report and validated with credit report information on a subset of subjects (n=37). Assets were assessed with the question: "What are your approximate current assets? (i.e., portion of home owned, bank accounts, investments, belongings)" using a 16-category ordinal response scale ranging from <+$500.00 in the lowest category to >+$1,500,000.00 in the highest. Debt were assessed with the question: "What are your approximate current debt? (i.e., outstanding home loans, outstanding car loans, outstanding student loans, credit card debt, medical debt)" using a 16 category ordinal response scale ranging from <–\$500.00 in the lowest category to >–\$1,500,000.00 in the highest. Other demographic data was also obtained in this questionnaire (e.g., age, sex, ethnicity, education). In addition to overall credit score, measures of available credit amount and the percent of credit used were extracted from credit reports and used to validate and distinguish self-reported assets and debt.

Cognitive and Demographic Variables. Selected neuropsychological tests were administered to assess potential cognitive confounds. The WAIS-III Digit Span Test assessed working memory capacity by requiring subjects to repeat numerical strings forward and backwards (Wechsler, 1997). The Trail Making Test (TMT) assessed cognitive flexibility by requiring subjects to first connect circled numbers in a sequential order, and next to connect a series of numbers and letters in an alternating order (Reitan, 1993). Finally, a numeracy inventory (11 items) assessed quantitative skills with basic number problems (Lipkus et al., 2001).

Brain image acquisition and analysis. FMRI (1.5 T General Electric magnetic resonance scanner with a standard quadrature head coil) was acquired as subjects played the MIL task using standard functional scanning parameters (24 4-mm thick slices with in-plane resolution 3.75 X 3.75 mm in plane resolution and no gap; extending axially from the midpons to the top of the skull; T2*-sensitive spiral in-out pulse sequence with repetition time 2 sec, echo 40 msec, and flip 90 degrees). High-resolution structural scans were subsequently acquired to facilitate localization and coregistration of functional data

(T1-sensitive spoiled gradient recalled acquisition pulse sequence with repetition time 100 msec, echo 7 msec, and flip 90 degrees). Standard preprocessing was applied to FMRI data prior to multiple regression analysis (i.e., voxel time series were concatenated across runs, sinc interpolated to correct for nonsimultaneous slice acquisition within each volume, corrected for three dimensional motion, high-pass filtered to remove frequencies > 90 s, and converted to percent signal change with respect to each voxel's mean activation over the entire experiment). Visual inspection of motion correction estimates confirmed that no subject's head moved more than 2 mm in any dimension from one volume acquisition to the next.

Whole brain analyses of FMRI data first regressed voxel-based activity against computationally-derived models of various aspects of reinforcement learning (Cox, 1996; Pessiglione et al., 2006). Consistent with the notion that distinct systems evaluate anticipated gain and loss, multiple regression coefficients modeled four parameters related to neural activity during anticipation of gain (i.e., gain expected value; GEV) and anticipation of loss (i.e., loss expected value; LEV), as well as in response to gain outcomes (i.e., gain prediction error; GPE) and loss outcomes (i.e., loss prediction error; LPE) (Knutson et al., 2005; Yacubian et al., 2006). A standard reinforcement learning algorithm was fit to individual subject choices to derive these regressors (O'Doherty et al., 2004; Sutton and Barto, 1998). For each cue in a stimulus pair (e.g., A and B), the model estimates the expected values of choosing A (Qa) and B (Qb), or the reward expected by selecting the cue. Q values were initialized at 0, and the value of the chosen stimulus (e.g., A) was updated after every trial according to the rule $Qb_{t+1} = Qa_t + \alpha*\delta_t$. δ_t

represents the prediction error, $R_t - Qa_t$, or the difference between the expected (i.e., Qa_t) and actual outcome (i.e., R_t) at trial t. The probability of selecting each action was estimated using the softmax rule, which calculates the probability of selecting a cue based on the values (e.g., for choosing A, $Pa_t = \exp(Qa_t/\beta)/(\exp(Qa_t/\beta) + \exp(Qb_t/\beta))$. The constants α (learning rate) and β (temperature) were adjusted to maximize the probability of the observed choices under the model. Best-fitting values for gain learning were $\alpha = 0.34$ (95% CI: 0.25–0.43) and $\beta = 0.12$ (95% CI: 0.06–0.17), and for loss learning were $\alpha = 0.39$ (95% CI: 0.27–0.51) and $\beta = 0.22$ (95% CI: 0.16–0.28), comparable to previous findings (Pessiglione et al., 2006). The learning model was fitted with a single set of parameters across all subjects, since individual fits yielded less consistent results. GEV and LEV were modeled during anticipation and referred to the subsequently chosen cue, whereas GPE and LPE were modeled during feedback and referred to the outcome of cue choice. Subjects who underwent FMRI played a double-length version of the task (24 trials per condition). For these subjects, both behavioral modeling and neuroimaging regression analyses were conducted using all 24 trials for each condition. Measures of behavioral performance for the entire sample only included the first half of each session (12 trials per condition), however, to equate the percentage-based performance measures across both scanned and unscanned subjects.

For individual difference analyses, we averaged and extracted coefficient data from volumes of interest in the MPFC, NAcc and anterior insula specified on the basis of 6 mm diameter spheres centered on foci identified in previous studies (MPFC and NAcc) or volumes described in the Talairach atlas (insula anterior to A=0) for each individual.

After examining zero-order correlations, anticipatory model coefficients from each volume of interest were regressed against gain correct choices and loss correct choices in two separate regression models.

Acknowledgments

We thank Dennis Chan, Stephanie Greer, Robert Ruelas, and Andrew Trujillo for assistance with data acquisition; Darrell Worthy for help with computational modeling; and Samuel M. McClure, Jeanne L. Tsai, and members of the Spanlab for feedback on prior drafts. This research was supported by a grant from the FINRA Investor Education Foundation as well as U.S. National Institute on Aging grants AG030778 to B.K. and AG032804 to G.R.S.L.

Author Contributions

B.K., G.R.S.L., and C.M.K. designed the study. G.R.S.L. collected the data. B.K., G.R.S.L. and C.M.K. analyzed the data. B.K., G.R.S.L. and C.M.K. prepared the manuscript.

Table 1. Predictors of self-declared assets and debt (* p<.05, ** p<.01, ***p<.001; predicted associations in **bold**)

	Assets	*Assets*	*Debt*	*Debt*
Gain correct choices	3.30	2.06	1.74	0.80
	2.85**	**2.16***	1.05	0.46
Loss correct choices	–2.66	–1.62	–5.92	–7.83
	–1.30	–0.90	**–2.14***	**–2.59***
Debt	0.10	0.13		
	1.25	2.01*		
Assets			0.19	0.42
			1.25	2.01*
Working memory		0.12		–0.08
		0.95		–0.37
Cognitive flexibility		–0.01		–0.02
		–0.15		–0.57
Numeracy		0.09		0.33
		0.35		0.72
Age		0.13		–0.09
		5.28***		–1.89*
Education		0.16		0.33
		1.07		1.30
Sex = male		–0.34		0.85
		–0.52		0.75
Constant	11.97	0.64	7.80	4.67
	7.22***	0.18	2.79***	0.76
Ethnicity FEs included		YES		YES
Adjusted R^2	.11	.50	.07	.16
Observations	82	80	82	80

Table 2. Individual differences: MPFC sensitivity to gain information predicts gain learning and Insula sensitivity to loss information predicts loss learning. Learning performance is defined according to either Bayesian updating (columns 2 and 4) or to Reinforcement Learning (columns 3 and 5).

	$GainCorrect$ $(Bayes)$	$GainCorrect$	$LossCorrect$ $(Bayes)$	$LossCorrect$
$\beta\_MPFC^{ANT}\_EVOP\_GAIN$	0.08	0.07		
	(2.34)**	(1.87)*		
$\beta\_INSULA^{ANT}\_EVOP\_LOSS$			0.03	0.06
			(1.30)	(2.46)**
Constant	0.65	0.65	0.69	0.63
	(13.74)***	(12.17)***	(19.79)***	(19.36)***
Adj. R^2	0.11	0.06	0.02	0.12
Observations	39	39	39	39

Table 3. Brain sensitivity to gain information and sensitivity to loss information are not related within individuals.

	β_MPFC^{ANT} $\_EVOP\_GAIN$	β_MPFC^{ANT} $\_EVOP\_GAIN$
$\beta\_MPFC^{ANT}\_EVOP\_LOSS$	-0.01	
	(-0.13)	
$\beta\_INSULA^{ANT}\_EVOP\_LOSS$		-0.08
		(-0.46)
Constant	0.08	0.07
	(0.33)	(0.30)
Adj. R^2	-0.03	-0.02
Observations	39	39

Table 4. Self-declared assets and debt correlate with distinct credit report variables. (top entry: coefficient (s.e.m.); bottom entry: t-statistic; * $p<.05$, ** $p<.01$, two-tailed)

	Assets	*Debt*
Credit amount	4.79×10^{-6} 2.37*	6.64×10^{-6} 3.39**
Percent credit used	–2.12 –1.58	7.18 5.56***
Constant	13.62 17.39***	2.40 3.17**
Adj. R^2	.11	.63
Observations	37	37

Figure 1. Monetary Incentive Learning task gain (top) and loss (bottom) trial structure.

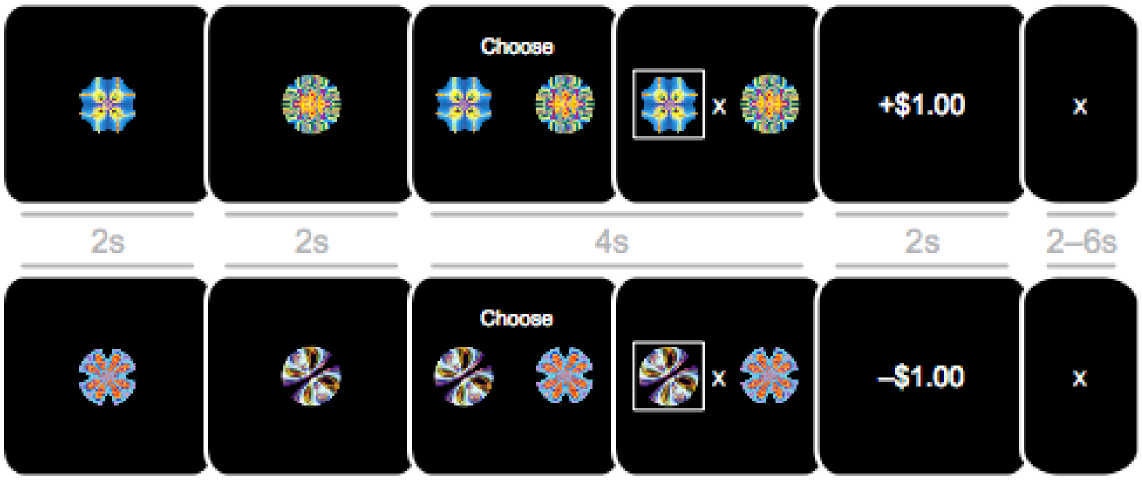

Figure 2. Learning about gains and learning about losses are not correlated within individuals.

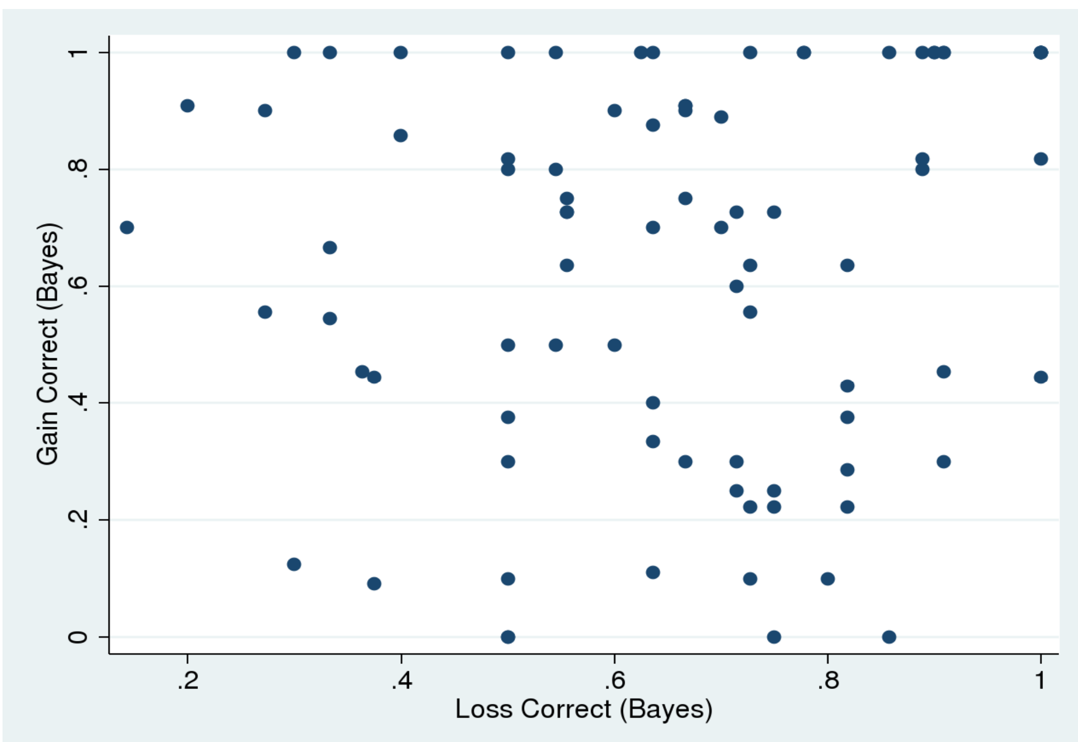

Figure 3. Individual differences in gain and loss learning account for assets and debt. Anticipatory activity in the MPFC correlated with gain expected value (A), while anticipatory activity in the insula correlated with loss expected value (B). Gain learning correlated with self-reported assets (green; C), while loss learning correlated with self-reported debt (red; D). Self-reported assets correlated with credit amount (E), while self-reported debt correlated with percentage credit used (obtained from credit reports; F). Panels C-F are residual plots with the trendline depicting the correlation between residuals.

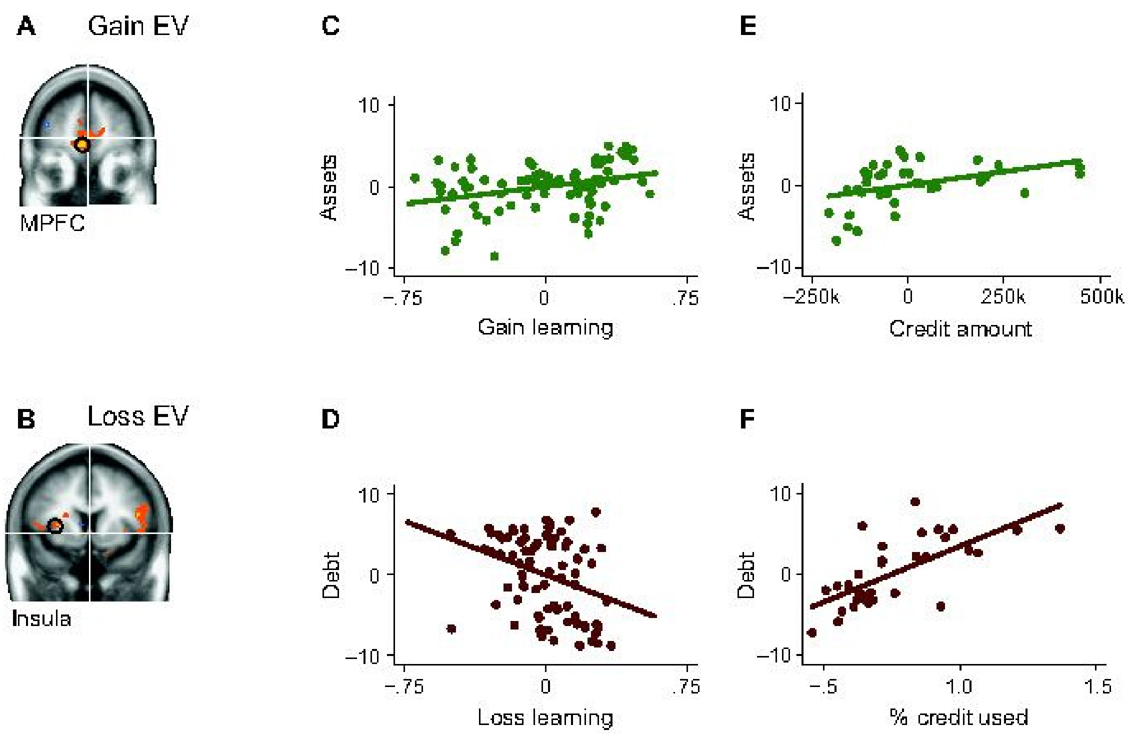

References

Adcock, R.A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., and Gabrieli, J.D.E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. Neuron *50*, 507-517.

Agarwal, S., Driscoll, J.C., Gabaix, X., and Laibson, D. (2007). The age of reason: Financial decisions over the lifecycle. (Harvard University).

Bernoulli, D. (1738 / 1954). Exposition of a new theory on the measurement of risk. Econometrica *22*, 23-36.

Bouchard, T.J. (1994). Genes, environment, and personality. Science *264*, 1700-1701.

Buchel, C., Morris, J., Dolan, R.J., and Friston, K.J. (1998). Brain systems mediating aversive conditioning: An event-related FMRI study. Neuron *20*, 947-957.

Burks, S.V., Carpenter, J.P., Goette, L., and Rustichini, A. (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. Proceedings of the National Academy of Science *106*, 7745-7750.

Cox, R.W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance images. Computers in Biomedical Research *29*, 162-173.

Gottfried, J.A., O'Doherty, J., and Dolan, R.J. (2003). Appetitive and aversive olfactory learning in humans studied using event-related functional magnetic resonance imaging. Journal of Neuroscience *22*, 10829-10837.

Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica *47*, 263-291.

Kim, H., Shimojo, S., and O'Doherty, J.P. (2006). Is avoiding an aversive outcome rewarding? Neural substrates of avoidance learning in the human brain. Public Library of Science Biology *4*, e233.

Knutson, B., and Greer, S.M. (2008). Anticipatory affect: Neural correlates and consequences for choice. Philosophical Transactions of the Royal Society B *363*, 3771-3786.

Knutson, B., Taylor, J., Kaufman, M., Peterson, R., and Glover, G. (2005). Distributed neural representation of expected value. Journal of Neuroscience *25*, 4806-4812.

Kuhnen, C.M., and Knutson, B. (2005). The neural basis of financial risk-taking. Neuron *47*, 763-770.

LaBar, K.S., Gatenby, J.C., Gore, J.C., LeDoux, J.E., and Phelps, E.A. (1998). Human amygdala activation during conditioned fear acquisition and extinction: A mixed-trial FMRI study. Neuron *20*, 937-945.

Lipkus, I.M., Samsa, G., and Rimer, B.K. (2001). General performance on a numeracy scale among highly educated samples. Medical Decision Making *21*, 37-44.

McClure, S.M., Berns, G.S., and Montague, P.R. (2003). Temporal prediction errors in a passive learning task activate human striatum. Neuron *38*, 338-346.

Mischel, W., Shoda, Y., and Rodriguez, M.L. (1989). Delay of gratification in children. Science *244*, 933-938.

O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., and Dolan, R.J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science *304*, 452-454.

O'Doherty, J.P., Dayan, P., Friston, K.J., Critchley, H.D., and Dolan, R.J. (2003). Temporal difference models and reward-related learning in the human brain. Neuron *38*, 329-337.
Pessiglione, M., Seymour, B., Flandin, G., Dolan, R.J., and Frith, C.D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. Nature *442*, 1042-1045.
Preuschoff, K., Bossaerts, P., and Quartz, S.R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. Neuron *51*, 381-390.
Reitan, R.M. (1993). The Halstead-Reitan neuropsychological test battery: Theory and clinical interpretation, 2nd edn (New York: Neuropsychology Press).
Samanez-Larkin, G.R., Gibbs, S.E.B., Khanna, K., Nielsen, L., Carstensen, L.L., and Knutson, B. (2007). Anticipation of monetary gain but not loss in healthy older adults. Nature Neuroscience *10*, 787-791.
Samanez-Larkin, G.R., Hollon, N.G., Carstensen, L.L., and Knutson, B. (2008). Individual differences in insular sensitivity during loss anticipation predict avoidance learning. Psychological Science *19*, 320-323.
Sutton, R.S., and Barto, A.G. (1998). Reinforcement learning (Cambridge, MA: MIT Press).
Thaler, R. (1980). Toward a positive theory of consumer choice. Journal of Economic Behavior and Organization *1*, 39-60.
Webley, P., Burgoyne, C.B., Lea, S.E.G., and Young, B.M. (2001). The economic psychology of everyday life (Hove, UK: Psychology Press).
Wechsler, D. (1997). Wechsler Adult Intelligence Scale, 3rd edn (San Antonio, TX: Psychological Corporation).
Wittmann, B.C., Schott, B.H., Guderian, S., Frey, J.U., Heinze, H.J., and Duzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. Neuron *45*, 459-467.
Yacubian, J., Glascher, J., Schroeder, K., Sommer, T., Braus, D.F., and Buchel, C. (2006). Dissociable systems for gain-and loss-related value predictions and errors of prediction in the human brain. Journal of Neuroscience *26*, 9530-9537.

Supplement

Supplementary Procedure 1. Validation of gain and loss correct choice indices with optimal choice estimates

This study utilized subjects' percentage "correct" choices in gain and loss conditions (excluding the first trial) as the primary predictor of life financial outcomes. Based on information that each individual receives during the task, one can define measures of optimal choice for that individual as the fraction of trials in the gain condition or loss condition when subjects made the correct Bayesian choice, excluding trials in which either option had an equal chance of being optimal. As shown in Table S1 gain optimal choices but not loss optimal choices were associated with gain correct choices. Loss optimal choices but not gain optimal choices were associated with loss correct choices.

Table S1. Reinforcement learning correct choices and Bayesian optimal choices are highly correlated. (*p<.05, **p<.01, ***p<.001, two-tailed)

	Gain Correct	*Loss Correct*
Gain Correct (Bayes)	1.01	0.01
	22.19***	0.15
Loss Correct (Bayes)	0.01	0.66
	0.07	8.23***
Constant	–0.04	0.17
	–0.72	3.02**
Adj. R^2	.85***	.44***
Observations	82	82

Supplementary Procedure 2. Validation of self-reported assets and debts with credit report data

Based on information derived from credit reports in approximately half of the sample (n=37), self-reported assets and debts were regressed against credit amount and percent credit used, respectively. Self-reported assets and debt were also regressed against overall credit score (i.e., FICO score retrieved from Experian). Results indicated that while credit amount was associated with assets, percent credit used was more robustly associated with debts. Experian FICO score was also correlated with both assets and debts in opposite directions. Together, these results suggest that self-reported assets and debts are associated with more objective indices of financial status, and that self-reported assets may relate to different aspects of financial status than do self-reported debts.

Supplementary Procedure 3. Test-retest reliability of gain and loss choice indices (independent sample)

To examine test-retest reliability of the gain and loss correct choice indices, 30 subjects (22 female) repeatedly participated in an online version of the Monetary Incentive Learning Task. The version featured three sessions of 12 trials of gain learning and three sessions of 12 trials of loss learning, with all six sessions mixed in a randomized order. Subjects played for the same amounts of cash as in the MIL task described in the main manuscript, but were informed that they would actually receive their earnings if their name was drawn at the end of the month (10% chance). Subjects played the exact same task twice, with a minimum of three weeks between repeated plays. Subjects' number of correct choices (minus the first trial) were computed for each of the three gain learning and loss learning sessions and then averaged across sessions within condition. Next, gain and loss averages were correlated across the first and second administration of the task. The test-retest reliability of gain correct choices ($r=.49$, $p<.001$, two-tailed) and of loss correct choices ($r=.49$, $p<.001$, two-tailed) was significant and moderate, suggesting reasonable temporal stability across repeated administrations. Interestingly, performance did not significantly change from the first to the second administration.