

Smokers' brains compute, but ignore, a fictive error signal in a sequential investment task

Pearl H Chiu¹⁻⁴, Terry M Lohrenz^{1,2,4} & P Read Montague¹⁻³

Addicted individuals pursue substances of abuse even in the clear presence of positive outcomes that may be foregone and negative outcomes that may occur. Computational models of addiction depict the addicted state as a feature of a valuation disease, where drug-induced reward prediction error signals steer decisions toward continued drug use. Related models admit the possibility that valuation and choice are also directed by 'fictive' outcomes (outcomes that have not been experienced) that possess their own detectable error signals. We hypothesize that, in addiction, anomalies in these fictive error signals contribute to the diminished influence of potential consequences. Using a simple investment game and functional magnetic resonance imaging in chronic cigarette smokers, we measured neural and behavioral responses to error signals derived from actual experience and from fictive outcomes. In nonsmokers, both fictive and experiential error signals predicted subjects' choices and possessed distinct neural correlates. In chronic smokers, choices were not guided by error signals derived from what might have happened, despite ongoing and robust neural correlates of these fictive errors. These data provide human neuroimaging support for computational models of addiction and suggest the addition of fictive learning signals to reinforcement learning accounts of drug dependence.

In healthy learners, positive prediction errors generated from actual experience indicate that a state is 'better than expected' and have been shown to bias actions, an observation interpreted as an attempt by reward-harvesting mechanisms to maximize the amount of reward obtained over time¹⁻⁶. The temporal difference (TD) model of reinforcement learning associates these experiential prediction errors with phasic changes in dopaminergic activity that tracks ongoing differences between expected and actually experienced rewards⁷⁻¹¹. In this framework, transient increases in dopamine induced by addictive drugs^{12,13} have been posited to generate invariably positive experiential reward-prediction errors, thus exaggerating the value of drug-induced states and reinforcing drug-seeking behavior¹.

However, dopamine increase during drug use is not sufficient for addiction¹⁴. One key characteristic of addicted individuals is that they pursue and consume subjective rewards even in the clear presence of positive outcomes that might be foregone or negative outcomes that might occur¹⁵ as a consequence of their actions. Such data strongly suggest that addicts have an impaired capacity to use possible outcomes to intervene on habitual drug-taking. One hypothesis is that chronic drug-takers have a diminished capacity in their nervous system for computing the appropriate control signals that guide behavior based on fictive outcomes. Specifically, error signals for such fictive outcomes may not be produced in the brains of addicts. An alternative hypothesis is that an addicted brain does produce error signals related to fictive outcomes, but the influence of these signals on actual behavioral choice is substantially diminished. Naturally, a myriad of intermediate

possibilities also exist, and the role of perturbed learning signals as antecedents or consequences of addiction remains an open question. In this paper, we sought to distinguish between the aforementioned two basic hypotheses by comparing nonsmokers and chronic smokers in an experiment designed to probe the impact of experiential and fictive errors. These error signals are those derived from actual experience and from outcomes that were not experienced, respectively. The possibility that these ideas can be tested has been suggested by recent work showing that fictive error signals guide behavior and can be tracked in healthy human brains using a simple investment game (see schematic representation in **Fig. 1a,b**)².

As was done previously², we used functional magnetic resonance imaging (fMRI) and a sequential gambling task that assesses ongoing reward expectations, rewards gained and the ensuing difference between fictive rewards that could have been gained and those that are actually gained. In this game, subjects view price histories from historical stock markets and choose percentages of an initial \$100 endowment to invest in each segment of ongoing market fluctuations (see **Fig. 1b** and Methods for complete task description). Each market comprised 20 sequential investment decisions, and all subjects played ten active markets. After each investment, the current portfolio amount and percentage market change are updated and shown to the subject (**Fig. 1b**).

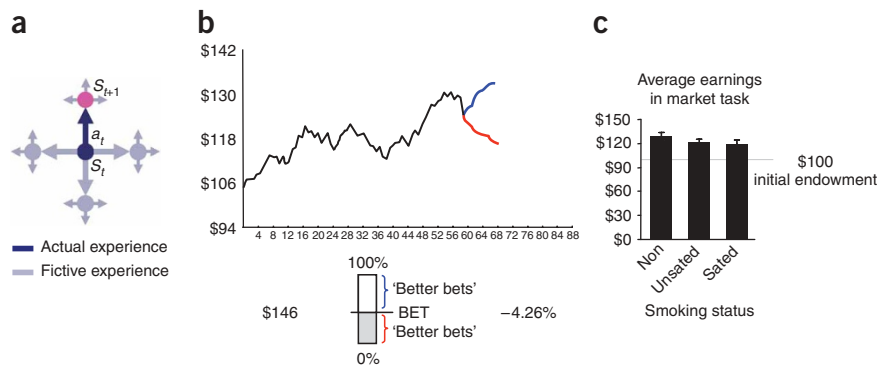
Thirty-one chronic cigarette smokers participated in the investment game on 2 separate days of fMRI scanning: a 'sated' day and an 'unsated' day (see Methods for complete participant information). For

¹Computational Psychiatry Unit, ²Department of Neuroscience, ³Menninger Department of Psychiatry & Behavioral Sciences, Baylor College of Medicine, One Baylor Plaza, Houston, Texas 77030, USA. ⁴These authors contributed equally to this work. Correspondence should be addressed to P.R.M. (rmontague@cpu.bcm.edu).

Received 9 November 2007; accepted 7 February 2008; published online 2 March 2008; doi:10.1038/nn2067

Figure 1 Schematic representation of fictive error and total earnings on the market task.

(a) Schematic representation of the state space dynamics of reward processing using actual and fictive experience. A learner starts in the state S_t (represented by the center dark circle), takes the action a_t , receives reward r_t and transitions to state S_{t+1} (represented by the top red circle). This is the learner's actual experience. The learner can gain additional valuable information about the environment from fictive experience, or what might have been experienced had the learner taken other actions (light arrows) and received alternate rewards \hat{r}_t (light circles). (b) Schematic representation of the market task and measurement of reward expectations, rewards gained and fictive outcomes that might have been. The y axis depicts the market price and the x axis represents time. Each investment decision or bet (b_t , indicated by the gray slider bar at bottom center) in the game is a proxy for the subjects' reward expectations at time t . On submitting an investment decision, the subsequent piece of market history, indicated by the black trace, is revealed; the incremental change relative to the prior market price is the market return (r_t). Fictive error over gains, generated from rewards that might have been earned on positive market fluctuations, is computed as the difference between the maximum possible gain (blue trace) and the actual gain, $f^+ = (1 \cdot r_t^+) - (b_t r_t^+)$. After each investment, both the current portfolio amount (here, \$146) and percentage market fluctuation (here, -4.26%) were updated and displayed to the subject. (c) Final monetary earnings in the investment game. All participants were given an initial endowment of \$100 to invest in the market game. The bar plots depict mean and s.e.m. of final monetary earnings in each group. Smokers and nonsmokers did not differ in their final portfolios.



the sated session, subjects were asked to smoke as usual throughout the day and also to smoke until satiated on arrival at the laboratory. For the unsated session, subjects were asked to refrain from smoking after midnight before the experimental session. A measure of exhaled carbon monoxide was obtained from all smokers as an objective proxy of satiety and blood nicotine saturation^{16,17}. The smokers' behavioral and neural data were compared both across sated and unsated sessions and relative to a group of 31 nonsmokers (a subset of participants from our previous work²); the groups are hereafter referred to as sated smokers, unsated smokers and nonsmokers, respectively.

The market investment game used historical stock markets and, as such, was designed to elicit strategic behavioral choices in a context in which employing specific strategies should not affect the final outcome

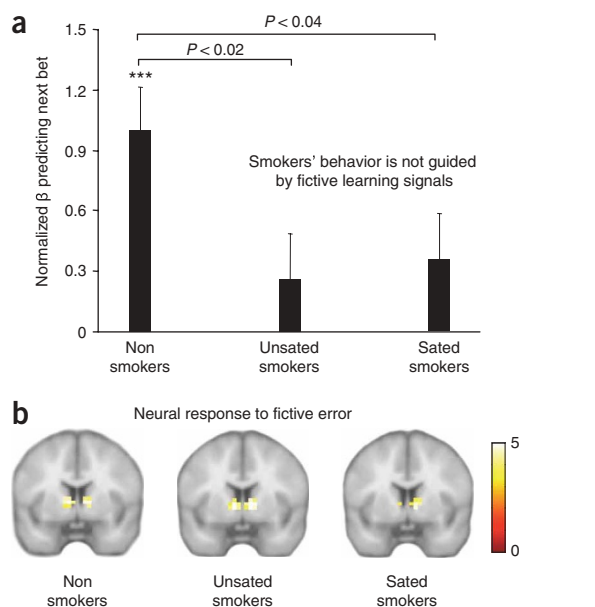
of the game. Implementing such a task enabled us to parse the contributions of fictive and experiential learning signals on behavioral choice, separate from the impact that differential monetary earnings might confer on neural and behavioral responses. As would be expected of naive investors in real-world markets, participants across groups did not differ in average total earnings in the game (Fig. 1c).

RESULTS

Responses to error signals derived from fictive outcomes

To assess the impact of fictive error signals on subjects' behavior, we quantified fictive error over gains (f^+ ; signals derived from what might have happened) as the difference between the maximum possible reward (equal to the positive market return, r_t^+) and the actual reward

Figure 2 Smokers' and nonsmokers' responses to error signals derived from fictive outcomes. Fictive error over gains was quantified as the ongoing difference between the maximum possible gain that could have occurred at time t and the actual gain at time t , $f^+ = (1 \cdot r_t^+) - (b_t r_t^+)$. Here, r_t^+ is the relative positive change in the market at time t , and b_t is the bet placed at time t . In nonsmokers, multiple regression analyses on key market and behavioral variables showed a negative coefficient on the actual gain ($b_t r_t^+$) and concurrently a positive coefficient on the maximum fictive gain ($1 \cdot r_t^+$), in agreement with previous work showing that the fictive error is a major predictor of next bets on this task (Supplementary Table 1)². (a) Smokers' behavior was not guided by fictive error signals. The behavioral regressions also showed that the negative correlation between actual gain and next bet, as seen in nonsmokers, reduced to no correlation in smokers, thus demonstrating a disrupted behavioral influence of the fictive error on behavioral choice for the smoker group. The bar plot depicts normalized mean beta coefficients and s.e.m. for the actual gain term predicting subjects' next bets (beta for $b_t r_t^+$, normalized to the nonsmoker coefficient, which carries a negative value; Supplementary Table 1). This gain term is the part of the fictive error that loses its influence on the next bet in smokers. Only the nonsmoker beta coefficient for the actual gain was significantly different from zero (***) indicates $P < 0.0001$), whereas the actual gain had no influence on the next bet for the unsated ($P > 0.24$) and sated smokers ($P > 0.11$). The complete regression table is presented in Supplementary Table 1. (b) In nonsmokers and both smoker groups, robust neural response to fictive error was seen in bilateral caudate. Thresholded SPM2 t -statistic maps of the neural fictive error regressor are shown here. In contrast with the absent behavioral influence of fictive error in smokers, the neural response to fictive error was seen in all three groups, regardless of smoking status ($P < 0.001$, uncorrected; cluster size ≥ 5 , random effects analysis; $n = 31$ nonsmokers, $n = 31$ unsated smokers and $n = 31$ sated smokers; $y = 4$). The complete activation table is presented in Supplementary Table 3.



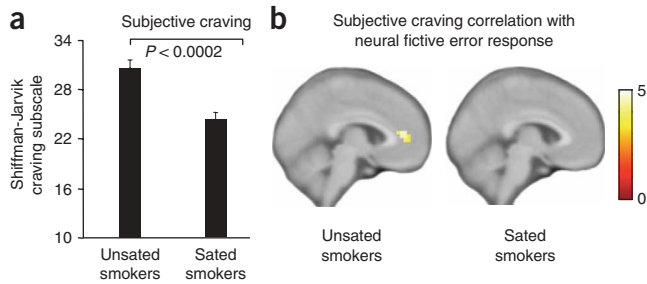


Figure 3 Subjective craving and correlation with neural response to fictive error in unsated and sated smokers. **(a)** Self-report measure of craving differentiated unsated from sated smokers. On arrival at the laboratory for both sessions, subjects completed a battery of self-report questionnaires, including the SJWS. The bar plot depicts means and s.e.m. from the Craving subscale of the SJWS. Unsated smokers reported significantly greater craving than sated smokers (unsated > sated, $P < 0.0002$). **(b)** Correlation of neural fictive error beta coefficients with subjective craving score. Unsated smokers showed a significant correlation in rACC of neural responses to fictive error with subjective craving ($P < 0.001$, uncorrected; cluster size ≥ 5 ; maximum: $x = -4$, $y = 40$, $z = 12$).

or gain at the market fluctuation at each time (t). In this case, gain is equal to the product of a player's bet (b_t) and the positive market return, r_t^+ , and the resulting fictive error is $f^+ = (1 \cdot r_t^+) - (b_t r_t^+)$. As such, the fictive error measures the ongoing difference between the maximum that could have been gained, $1 \cdot r_t^+$, and what was actually gained, $b_t \cdot r_t^+$, and is exactly the fictive error signal previously identified to strongly guide the bets of normal players². We focused on fictive errors in the gain domain of our market task to assess sensitivity to foregone positive outcomes, operationalized as what each subject could have earned, but did not earn, despite a positive market fluctuation.

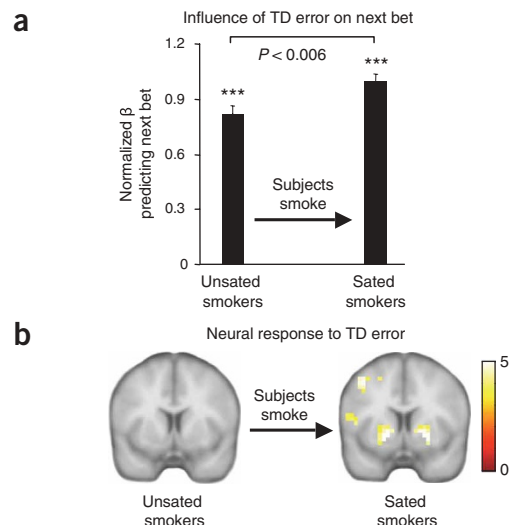
Multiple regression analyses of players' next bets (b_{t+1}) were carried out against the following key positive and negative^{2,18} market and behavioral variables: previous bet (b_t), previous market separated into positive and negative market returns (r_t^+ and r_t^- , respectively) and the respective interaction terms ($b_t r_t^+$ and $b_t r_t^-$). In all groups, the three first-order terms of previous bet, positive market and negative market (b_t , r_t^+ , r_t^-) significantly predicted subjects' next bets (b_{t+1} , all

$P < 0.0001$). However, a negative contribution from $b_t r_t^+$ was seen in nonsmokers only ($P < 0.0001$; see **Supplementary Table 1** online). Thus, as observed previously², fictive error over gains as defined above emerged as an important signal guiding the next bets of nonsmoking controls (**Fig. 2a** and **Supplementary Table 1**).

In marked contrast, the sated and unsated smokers' behavior was not influenced by these same fictive error signals. In chronic smokers, fictive errors over gains did not significantly guide players' next bets (**Fig. 2a** and **Supplementary Table 1**; $P > 0.1$ for both sated and unsated smokers relative to zero). Equally important, a robust neural response was observed to fictive errors in all groups, despite the absent influence of fictive errors on smokers' behavior. The neural signals were examined with fictive error regressors derived on the basis of each subject's unique behavior and fit to a standard hemodynamic response function that was time-locked to the presentation of each segment of market history (see **Supplementary Methods** online for a listing of all regressors used in the general linear model analyses¹⁹). Nonsmokers and both unsated and sated smokers all showed consistent neural responses in bilateral ventral caudate to fictive errors (**Fig. 2b**). This is exactly the region that was previously identified using the same f^+ regressor² and is also the region associated with the 'actor' role in actor-critic models of reinforcement learning⁶. Thus, in chronic smokers, the robust intact neural response to error signals derived from fictive outcomes is discordant with the absent impact of this error signal on behavioral choice.

Neural responses to the fictive error signal were further associated with subjective craving in anterior cingulate cortex. Specifically, linear regression analyses of craving (assessed using the Shiffman-Jarvik Withdrawal Scale; SJWS^{16,20}) with individual subjects' effect sizes of neural responses to fictive error identified a region of rostral anterior cingulate cortex (rACC) showing robust association with subjective craving only in unsated smokers (**Fig. 3**). rACC activation is commonly seen in response to stimuli that carry negative emotional salience and in tasks that elicit affective decision making or conflict^{21–25}. The association of subjective craving in unsated smokers with hemodynamic responses to fictive errors in the rACC further implicates this region in responses to undesirable outcomes derived from fictive experience and highlights that the dissociation between ventral caudate activation and smokers' behavioral response to fictive errors is not dependent on the state of nicotine satiety.

Figure 4 Responses to error signals derived from experienced outcomes. Individual subjects' experiential errors were computed as the ongoing difference between z-scored rewards already gained (\bar{g}_t) and z-scored rewards expected (\bar{b}_t), where the subjects' bets served as a proxy for expected reward ($TD_t = \bar{g}_t - \bar{b}_t$). Subjects' normalized next bets were regressed against normalized previous bet and the TD error. **(a)** Nicotine satiation increased behavioral influence of the TD error. The bar plot shows the normalized mean beta coefficients and s.e.m. of TD predicting next bets in the unsated and sated smokers (normalized to the sated smoker case). The behavior of both the unsated and sated smokers was strongly influenced by TD error (***) indicates beta coefficients with $P < 0.0001$ relative to zero). In addition, the influence of TD was significantly larger in the sated smokers (sated > unsated, $P < 0.006$). The influence of TD in nonsmokers was equivalent to that in the unsated smokers versus unsated smokers, $P = 0.74$). The complete regression tables are presented in **Supplementary Table 2**. **(b)** Nicotine satiation increased neural responses associated with TD error. Thresholded SPM2 t -statistic maps of the TD regressor in unsated and sated chronic smokers are shown. The neural response to TD error was robust in the sated smokers and diminished in the unsated group ($P < 0.001$, uncorrected; cluster size ≥ 5 , random effects analysis, $n = 31$ unsated smokers, $n = 31$ sated smokers, $y = 8$).



Responses to error signals derived from experienced outcomes

To assess the impact of nicotine saturation on prediction error signals derived from actually experienced reward, we used a model-free version of an experiential (TD-based) error signal and compared unsated and sated smokers' neural and behavioral responses to these signals. Specifically, individual subjects' experiential errors were computed as the ongoing difference between z -scored rewards already gained (\tilde{g}_t) and z -scored rewards expected (\tilde{b}_t), where the subjects' bets serve as a proxy for expected reward ($TD_t = \tilde{g}_t - \tilde{b}_t$). Individual TD regressors were then constructed for each subject and entered into linear multiple regression analyses predicting subjects' behavioral choices and into the general linear model examining neural responses associated with experiential prediction errors (see **Supplementary Methods** for details of general linear model analyses).

Both unsated and sated smokers' behavior was strongly influenced by experiential prediction error (**Fig. 4a**, see **Supplementary Table 2** online for multiple regression tables). Specifically, in both groups, subjects' subsequent bets increased with increasing magnitude of the experiential error (both $P < 0.0001$ versus zero; **Fig. 4a** and **Supplementary Table 2**). In addition, nicotine satiation conferred a magnified influence of this experiential learning signal in guiding smokers' next bets (sated $>$ unsated, $P < 0.006$; **Fig. 4a** and **Supplementary Table 2**).

Concomitant with the enhanced influence of experiential error on sated smokers' behavior, greater neural activation associated with TD error was seen in this group. That is, unsated smokers showed diminished neural responses to experiential errors in putamen, whereas sated smokers showed robust neural activation in bilateral putamen (**Fig. 4b**). Similar striatum activations to TD errors in decision-making tasks have been closely examined and consistently reported by our group and others (nonsmokers illustrated in **Supplementary Fig. 1** online)^{2–6}. Moreover, in ventral striatum, smokers' responses to TD errors (left putamen; **Supplementary Fig. 2** online) were not associated with subjective craving, but rather were parametrically associated with exhaled carbon monoxide levels; these levels are also known to correlate highly with plasma nicotine concentration^{16,17,26}.

DISCUSSION

From rewards gained to rewards that 'might have been', actual experience and fictive outcomes generate complementary learning signals that guide everyday behavior^{2,7,27}. Here, we found that chronic smokers, relative to nonsmokers, showed a reduced influence of abstractly framed learning signals on behavior without any accompanying loss of the associated neural signal. That is, although fictive error was indeed computed in this group (as indicated by the robust neural response in bilateral caudate), 'what might have been' did not emerge as a control signal that guides behavioral choice. In contrast, a marked behavioral impact of experiential TD-based learning signals, related directly to reward harvesting⁹, remained intact in smokers. Under nicotine satiation, neural signals associated with this experiential error exactly paralleled an enhanced impact of these errors in guiding behavior. Together, this model-based approach demonstrates that neural error signals derived from 'what might happen' remain intact in addicts, but, as exactly propounded by clinical criteria of addiction¹⁵, their influence on decision making is absent.

Our data further suggest that the source of the observed difference originates from the lost impact of the actual gain on smokers' next bets (**Fig. 2**). This is an important component of the findings presented here, as addicts are thought to have a diminished response to biological rewards. However, for the fictive error to have an impact on

nonsmokers' behavior, two relationships must concurrently hold: the influence of the maximum fictive gain on the next bet must be positive, and the relationship of the actual gain to the next bet must be negative. Indeed, we found that the influence of the reward (here the actual gain) was negative in nonsmokers; that is, there was a strong negative relationship between actual gain and next bet across all gain levels (**Fig. 2a** and **Supplementary Table 1**). This was not true for smokers (**Fig. 2a**). Considered together, these findings suggest that actual gains are not being treated as rewards in the smoker group, given the typical designation of reward as a positive reinforcer on which learning occurs and on which approach behavior is predicated.

The present data suggest an extension to existing temporal difference models of reinforcement learning that offer a theoretical mechanism by which drugs of addiction compromise dopaminergic systems¹. Specifically, fictive prediction errors (signaling the ongoing difference between what might have been obtained and what was actually obtained) may be incorporated into the reinforcement learning framework using a modification of an actor-critic architecture called supervised reinforcement learning²⁸. In supervised reinforcement learning, the usual TD learning algorithm generates a TD-error critic signal that is used to guide behavior through its influence on the behavioral policy of the actor. The behavioral policy is further guided by input from an external supervisor with a more global assessment of the quality of a behavioral response. Fictive errors in living organisms may then be considered to be the output of an 'endogenous supervisor' that modulates or complements the error signal coming to the actor from the TD critic. A natural question that arises in this theoretical framework, and in light of the current findings, is how the actor balances the inputs of these two critics.

Many studies have highlighted the presence of increased levels of dopamine in the ventral striatum (nucleus accumbens, putamen) following acute exposure to nicotine^{13,26}. Such data, and those showing that dopamine enhancing drugs (for example, levodopa) confer increased ventral striatal response to experiential learning signals^{5,29}, suggest a similar contribution of nicotine to the increased TD signal shown here in the ventral striatum of sated smokers. Identifying the possible physiological substrates of the absent effect of the neural fictive error signal in chronic smokers is perhaps a greater challenge. Nevertheless, some insight may come from data that identify behavioral risk factors for drug use and show a variety of changes in neurotransmitter function resulting from chronic drug abuse, providing a mechanistic account for the behavioral phenomena characteristic of addicted individuals^{30–35}. In chronic smokers, the discordance between the robust neural fictive error and absent behavioral influence of this signal may fall in the constellation of features associated with chronic drug use or vulnerability to addiction.

Although potential mechanisms underlying the dissociation between smokers' neural activation and behavior were not tested in the current task, we draw insight from recent work in the neural substrates of temporal discounting and cognitive control. These studies show that smokers, relative to nonsmokers, consistently choose smaller immediate outcomes over larger, but more delayed, outcomes across both hypothetical and actual rewards^{36–39}. We speculate that the steep delay discounting seen in smokers may be a specific manifestation of a more general deficit in using striatal learning signals derived from fictive outcomes to guide behavior. In controls, functional neuroimaging data show that decision making about immediate outcomes elicits activation in the striatum, whereas decisions about delayed, or fictive, outcomes elicit additional activation in lateral prefrontal cortical regions that have been implicated in cognitive control^{40–42}. In smokers, higher-order control signals modulating the influence of fictive

outcomes on behavioral choice may be impaired, leaving smokers guided only by immediate or experiential rewards and uninfluenced by fictive learning signals.

In summary, we show here that nicotine satiation increases the expression of the TD error signal in ventral striatum and that the behavioral response to fictive error signals is diminished in chronic smokers, despite the ongoing and robust presence of a neural response to this signal. These findings are consonant with the emerging understanding of addiction at the molecular level, and the computational model-based approach implemented here facilitates a mechanistic understanding of continued drug use in the presence of potential negative consequences and foregone positive outcomes.

METHODS

Participants. Smokers with no known psychiatric disorders were recruited from the Houston metropolitan area by advertisements seeking individuals who regularly smoke at least 15 cigarettes a day, had been doing so for at least the past year and were not currently attempting to quit. After an initial screening in which basic fMRI contraindications were assessed, qualifying individuals were invited to the laboratory for further assessment and fMRI scanning.

Following these procedures, 34 participants were initially enrolled in the study. On arrival at the laboratory, three of these individuals reported smoking an average of 12 daily cigarettes; data from these subjects were not included in subsequent analyses. The demographic and cigarette-use profile of the 31 included smokers are presented in **Table 1**.

In accordance with the Institutional Review Board of Baylor College of Medicine, written informed consent was obtained at the first laboratory visit, and it was emphasized that participants could withdraw from the study at any time with no adverse consequences.

Experimental procedure. At each visit to the laboratory, participants completed a standard battery of questionnaires assessing smoking use and mood (including the Fagerstrom Test for Nicotine Dependence⁴³, the SJWS²⁰ and the Positive and Negative Affect Scale⁴⁴).

Smokers were asked to complete the fMRI experimental sessions on 2 separate scanning days: a 'sated' day and an 'unsated' day. For the sated day, smokers were instructed to smoke as usual throughout the day, and, on arrival at the laboratory, to smoke until satiated. For the unsated session, smokers were instructed to refrain from smoking, beginning at midnight before the experimental day. The order of the two sessions was approximately balanced; 16 smokers completed the sated day first and 15 smokers completed the unsated day first. On arrival at the laboratory for both sessions, exhaled carbon monoxide was obtained as an objective assay of satiety and blood nicotine saturation (E50 Smokerlyzer)^{16,17}. Scores on measures of cigarette use at each scanning session are presented in **Table 2**.

We used 31 nonsmoking previously identified² individuals as the nonsmoker comparison group. These individuals answered "no" to both of the following questions: (i) "Have you smoked over 50 cigarettes in your lifetime?" and (ii) "Have you smoked in the past 12 months?"

Stimuli and task. On each day of fMRI scanning, subjects performed the following 'Market Task'. Subjects were endowed with \$100 to invest in stock market fluctuations. The market task was carried out under two conditions, Live and Not Live. At the beginning of each Live block, an initial segment of price history from a historic stock market was displayed. Subjects used a

Table 1 Subject characteristics (n = 31; mean ± s.d.)

Age (years)	35.1 ± 11.0
Sex (no. of females)	20
Education (years)	15.5 ± 2.1
Current cigarettes per day	19.1 ± 4.2
No. of quit attempts in past 5 years	2.2 ± 2.0
Age of daily smoking (years)	19.6 ± 4.8

Table 2 Smoking measures in Unsated and Sated states (mean ± s.e.m.)

	Unsated	Sated
Subjective craving*	31.6 ± 1.6	24.2 ± 1.0
Exhaled carbon monoxide (ppm)	7.9 ± 0.9	21.4 ± 2.0
Time since last cigarette (h)	11 ± 0.5	0 ± 0.0

Unsated and Sated smokers differed on all measures ($P < 0.001$ for each measure).
*Responses on the Craving subscale of the Shiffman-Jarvik Withdrawal Scale.

two-button box to move a slider bar to invest a percentage (0 to 100%) of their portfolio in each market segment. On submission of the investment decision (with a button box in the opposite hand), the next market segment was revealed after a delay of 750 ms, and subjects' current portfolio value and fractional market change were updated (see **Fig. 1b**). When subsequent market segments were revealed, the previous history remained on display, but the history was re-centered to prevent telegraphing unintended information about the market. After a delay of 750 ms, the slider bar changed from gray to red, indicating the free response time for the subsequent portfolio allocation decision. Each market comprised 20 investment decisions.

Not Live markets were interspersed as visual controls. In the Not Live condition, the market display remained the same while the portfolio amounts were replaced with 'N/A'. Subjects used the slider bar to make a simple visual-discrimination choice, indicating whether the current market price was higher or lower than the price two segments before the current segment. This was also repeated for 20 choices.

At each fMRI experimental session, participants played ten Live and ten Not Live markets, presented in pseudorandom order. The specific markets played in the Live and Not Live conditions at the first and second scanning sessions did not overlap.

Behavioral analyses. To assess the impact of fictive error, $f^+ = (1 \cdot r_t^+) - (b_t r_t^+)$, on subjects' behavior, we carried out linear mixed-effects multiple regression analyses on the series of market returns and investment decisions extracted from each subject. Subjects' investments were first z-scored within subject. As further detailed in the main text (also see **Fig. 1b**), each investment decision at time t is denoted as a bet, b_t (percentage of current portfolio invested in the market). On submitting an investment decision, the incremental market change relative to the prior market price is the market return (r_t). Therefore, positive market returns are expressed as $r_t^+ = (p_t - p_{t-1})/(p_{t-1}) > 0$, where p_t is the market price at time t ; negative market returns are similarly defined as $r_t^- = (p_t - p_{t-1})/(p_{t-1}) < 0$. Subject gains then become $b_t r_t^+$ for positive market returns, and losses are $b_t r_t^-$ for negative market returns. The final return of each market was excluded from the regression, as there is no investment decision following the final market segment. Using the variables defined above, the following multiple regression was carried out within each group in R (function *lm*, The R Foundation for Statistical Computing):

$$\tilde{b}_{t+1} = \beta_0 + \beta_1 \tilde{b}_t + \beta_2 r_t^+ + \beta_3 r_t^- + \beta_4 b_t r_t^+ + \beta_5 b_t r_t^-,$$

where \tilde{b}_t is the within-subject z-scored bet. The regression was performed simultaneously across all three groups by coding smoker status (unsated, sated and nonsmoker) as three indicator variables (SATED, UNSATED and NON-SMOKER) and by including a term in the regression of the form STATUS × regressor for each indicator and regressor (see **Supplementary Methods** for complete regressor list).

To assess the impact of nicotine satiation on smokers' behavioral responses to experiential learning signals, we carried out a multiple regression predicting subjects' behavior with state of smoking satiety and a model-free version of a TD error signal. Individual subjects' TD (experiential) errors were computed as the ongoing difference between z-scored rewards already gained (\tilde{g}_t) and z-scored rewards expected (\tilde{b}_t), where the subject's bet serves as a proxy for expected reward. At market return at time t , the z-scored subject gain, $g_t = (br)_t$, was defined as $\tilde{g}_t = (g_t - \text{mean}(g))/\text{s.d.}(g)$, where the mean and s.d. are taken over the subject gains already experienced. The subject's

z-scored bet, \tilde{b}_t , was defined similarly. The model-free TD error is $TD_t = \tilde{g}_t - \tilde{b}_t$. Smoker status was coded as described above using the indicators SATED, UNSATED and NONSMOKER. The following multiple regression was carried out in R (function *lm*, The R Foundation for Statistical Computing):

$$\begin{aligned} \tilde{b}_{t+1} = & \beta_0 \text{SATED} + \beta_1 \text{UNSATED} + \beta_2 \text{NONSMOKER} + \beta_3 \text{SATED} \cdot \tilde{b}_t \\ & + \beta_4 \text{UNSATED} \cdot \tilde{b}_t + \beta_5 \text{NONSMOKER} \cdot \tilde{b}_t \\ & + \beta_6 \text{SATED} \cdot TD_t + \beta_7 \text{UNSATED} \cdot TD_t + \beta_8 \text{NONSMOKER} \cdot TD_t \end{aligned}$$

fMRI data acquisition and reduction. All scans were carried out on a Siemens 3.0 Tesla Allegra scanner. Initial high-resolution T1-weighted scans were acquired using a magnetization-prepared rapid-acquisition gradient-echo sequence (MP-RAGE; Siemens). Continuous whole-brain imaging was performed as participants engaged in the interpersonal exchange task. Functional run details were as follows: echo-planar imaging, gradient recalled echo, repetition time (TR) = 2,000 ms, echo time (TE) = 40 ms, flip angle = 90°, 64 × 64 matrix, 26 4-mm axial slices acquired parallel to the anteroposterior commissural line for measurement of the blood oxygenation level-dependent effect^{45–47}. Scanning yielded functional 3.3-mm × 3.3-mm × 4.0-mm voxels.

Data reduction was carried out using SPM2 (<http://www.fil.ion.ucl.ac.uk/spm/>). Motion correction to the first functional scan was performed using a six-parameter rigid-body transformation within subjects⁴⁸. The average of the motion-corrected images was co-registered to each individual's structural MRI using a 12-parameter affine transformation. Slice timing artifact was corrected, and images were subsequently spatially normalized to the Montreal Neurological Institute (MNI) template⁴⁹ by applying a 12-parameter affine transformation, followed by nonlinear warping using standard basis functions⁵⁰. Finally, images were smoothed with an 8-mm isotropic Gaussian kernel and high-pass filtered in the temporal domain (filter width of 128 s).

General linear model analyses. To identify distinct blood oxygenation level-dependent responses associated with experiential and fictive error, two regressors, TD and fictive error, were added to the basic model (a complete list of regressors is in the **Supplementary Methods**). The TD and fictive error regressors were constructed as follows. The fictive error regressor was formed by convolving the market revelation (punctuate) regressor with the fictive error, $f^+ = (1 \cdot r_t^+) - (b_t r_t^+)$. The TD regressor was constructed by convolving the market revelation regressor with TD error. The initial market and first and final market revelation data were omitted from both the fictive and TD regressors. Orthogonalization of the fictive error regressor with respect to the TD error regressor was accomplished by subtracting the orthogonal projection of the fictive error onto the TD error from the fictive error regressor (fictive error orthogonalized with respect to the TD error regressor; **Fig. 3b**).

SPM2 beta maps were constructed for regressors of interest and then entered into a random-effects analysis by using the one-sample *t*-test function. All regions of activation positively associated with the TD and fictive error variables in the nonsmoker, unsated smoker and sated smoker groups are presented in **Supplementary Table 3** online (SPM2 Talairach coordinates in MNI space). Beta maps were thresholded at $P < 0.001$, uncorrected, cluster size ≥ 5 .

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank B. King-Casas (supported by the National Institute of Mental Health F32 MH078485) and R. Salas for scientific discussion, the Hyperscan Development Team at Baylor College of Medicine for Network Experiment Management Object (NEMO) software implementation (<http://www.hnl.bcm.tmc.edu/nemo>), X. Cui and J. Li for xjView image viewing and presentation software (<http://people.hnl.bcm.tmc.edu/cuixu/xjView>) and C. Bracero, J. McGee and S. Moore for technical assistance. This work was supported by the Kane Family Foundation (P.R.M.), the US National Institute on Drug Abuse (R01 DA11723 to P.R.M.), the US National Institute of Neurological Disorders and Stroke (R01 NS045790 to P.R.M.), the Angel Williamson Imaging Center and the American Psychological Association (T32 MH18882 to P.H.C.).

Published online at <http://www.nature.com/natureneuroscience>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions>

- Redish, A.D. Addiction as a computational process gone awry. *Science* **306**, 1944–1947 (2004).
- Lohrenz, T., McCabe, K., Camerer, C.F. & Montague, P.R. Neural signature of fictive learning signals in a sequential investment task. *Proc. Natl. Acad. Sci. USA* **104**, 9493–9498 (2007).
- Li, J., McClure, S.M., King-Casas, B. & Montague, P.R. Policy adjustment in a dynamic economic game. *PLoS ONE* **1**, e103 (2006).
- O'Doherty, J.P., Buchanan, T.W., Seymour, B. & Dolan, R.J. Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* **49**, 157–166 (2006).
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R.J. & Frith, C.D. Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* **442**, 1042–1045 (2006).
- O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* **304**, 452–454 (2004).
- McClure, S.M., Berns, G.S. & Montague, P.R. Temporal prediction errors in a passive learning task activate human striatum. *Neuron* **38**, 339–346 (2003).
- Schultz, W., Dayan, P. & Montague, P.R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
- Montague, P.R., Dayan, P. & Sejnowski, T.J. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**, 1936–1947 (1996).
- Montague, P.R., Hyman, S.E. & Cohen, J.D. Computational roles for dopamine in behavioural control. *Nature* **431**, 760–767 (2004).
- Daw, N.D. & Doya, K. The computational neurobiology of learning and reward. *Curr. Opin. Neurobiol.* **16**, 199–204 (2006).
- Koob, G.F., Sanna, P.P. & Bloom, F.E. Neuroscience of addiction. *Neuron* **21**, 467–476 (1998).
- Zhou, F.M., Liang, Y. & Dani, J.A. Endogenous nicotinic cholinergic activity regulates dopamine release in the striatum. *Nat. Neurosci.* **4**, 1224–1229 (2001).
- Volkow, N.D., Fowler, J.S. & Wang, G.J. Role of dopamine in drug reinforcement and addiction in humans: results from imaging studies. *Behav. Pharmacol.* **13**, 355–366 (2002).
- American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders: DSM-IV* (Washington, DC: American Psychiatric Association, 2000).
- Jarvik, M.E. *et al.* Nicotine blood levels and subjective craving for cigarettes. *Pharmacol. Biochem. Behav.* **66**, 553–558 (2000).
- Wald, N.J., Idle, M., Boreham, J. & Bailey, A. Carbon monoxide in breath in relation to smoking and carboxyhaemoglobin levels. *Thorax* **36**, 366–369 (1981).
- Kahneman, D. & Tversky, A. Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–291 (1979).
- Friston, K.J. *et al.* Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* **2**, 189–210 (1995).
- Shiffman, S.M. & Jarvik, M.E. Smoking withdrawal symptoms in two weeks of abstinence. *Psychopharmacology (Berl.)* **50**, 35–39 (1976).
- Bush, G., Luu, P. & Posner, M.I. Cognitive and emotional influences in anterior cingulate cortex. *Trends Cogn. Sci.* **4**, 215–222 (2000).
- Whalen, P.J. *et al.* The emotional counting Stroop paradigm: a functional magnetic resonance imaging probe of the anterior cingulate affective division. *Biol. Psychiatry* **44**, 1219–1228 (1998).
- Damasio, A.R. *et al.* Subcortical and cortical brain activity during the feeling of self-generated emotions. *Nat. Neurosci.* **3**, 1049–1056 (2000).
- Etkin, A., Egner, T., Peraza, D.M., Kandel, E.R. & Hirsch, J. Resolving emotional conflict: a role for the rostral anterior cingulate cortex in modulating activity in the amygdala. *Neuron* **51**, 871–882 (2006).
- Keightley, M.L. *et al.* An fMRI study investigating cognitive modulation of brain regions associated with emotional processing of visual stimuli. *Neuropsychologia* **41**, 585–596 (2003).
- Brody, A.L. Functional brain imaging of tobacco use and dependence. *J. Psychiatr. Res.* **40**, 404–418 (2006).
- Montague, P.R., King-Casas, B. & Cohen, J.D. Imaging valuation models in human choice. *Annu. Rev. Neurosci.* **29**, 417–448 (2006).
- Rosenstein, M.T. & Barto, A.G. in *Learning and Approximate Dynamic Programming: Scaling Up to the Real World* (eds Si, J., Barto, A., Powell, W. & Wunsch, D.) 359–380 (John Wiley & Sons, New York, 2004).
- Frank, M.J., Seeberger, L.C. & O'Reilly, R.C. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* **306**, 1940–1943 (2004).
- Hyman, S.E., Malenka, R.C. & Nestler, E.J. Neural mechanisms of addiction: the role of reward-related learning and memory. *Annu. Rev. Neurosci.* **29**, 565–598 (2006).
- Nestler, E.J. Molecular mechanisms of drug addiction. *Neuropharmacology* **47** Suppl 1, 24–32 (2004).
- McClung, C.A. *et al.* DeltaFosB: a molecular switch for long-term adaptation in the brain. *Brain Res. Mol. Brain Res.* **132**, 146–154 (2004).
- Marttila, K., Raattamaa, H. & Ahtee, L. Effects of chronic nicotine administration and its withdrawal on striatal FosB/DeltaFosB and c-Fos expression in rats and mice. *Neuropharmacology* **51**, 44–51 (2006).
- Hope, B.T., Nagarkar, D., Leonard, S. & Wise, R.A. Long-term upregulation of protein kinase A and adenylate cyclase levels in human smokers. *J. Neurosci.* **27**, 1964–1972 (2007).



35. Schepis, T.S., Adinoff, B. & Rao, U. Neurobiological processes in adolescent addictive disorders. *Am. J. Addict.* **17**, 6–23 (2008).
36. Bickel, W.K. *et al.* Behavioral and neuroeconomics of drug addiction: competing neural systems and temporal discounting processes. *Drug Alcohol Depend.* **90** Suppl 1, S85–S91 (2007).
37. Baker, F., Johnson, M.W. & Bickel, W.K. Delay discounting in current and never-before cigarette smokers: similarities and differences across commodity, sign and magnitude. *J. Abnorm. Psychol.* **112**, 382–392 (2003).
38. Odum, A.L., Madden, G.J. & Bickel, W.K. Discounting of delayed health gains and losses by current, never- and ex-smokers of cigarettes. *Nicotine Tob. Res.* **4**, 295–303 (2002).
39. Ohmura, Y., Takahashi, T. & Kitamura, N. Discounting delayed and probabilistic monetary gains and losses by smokers of cigarettes. *Psychopharmacology (Berl.)* **182**, 508–515 (2005).
40. McClure, S.M., Ericson, K.M., Laibson, D.I., Loewenstein, G. & Cohen, J.D. Time discounting for primary rewards. *J. Neurosci.* **27**, 5796–5804 (2007).
41. McClure, S.M., Laibson, D.I., Loewenstein, G. & Cohen, J.D. Separate neural systems value immediate and delayed monetary rewards. *Science* **306**, 503–507 (2004).
42. Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S. & Cohen, J.D. Conflict monitoring and cognitive control. *Psychol. Rev.* **108**, 624–652 (2001).
43. Heatherton, T.F., Kozlowski, L.T., Frecker, R.C. & Fagerstrom, K.O. The Fagerstrom Test for Nicotine Dependence: a revision of the Fagerstrom Tolerance Questionnaire. *Br. J. Addict.* **86**, 1119–1127 (1991).
44. Watson, D., Clark, L.A. & Tellegen, A. Development and validation of brief measures of positive and negative affect: the PANAS scales. *J. Pers. Soc. Psychol.* **54**, 1063–1070 (1988).
45. Kwong, K.K. *et al.* Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc. Natl. Acad. Sci. USA* **89**, 5675–5679 (1992).
46. Ogawa, S., Lee, T.M., Kay, A.R. & Tank, D.W. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc. Natl. Acad. Sci. USA* **87**, 9868–9872 (1990).
47. Ogawa, S., Lee, T.M., Nayak, A.S. & Glynn, P. Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. *Magn. Reson. Med.* **14**, 68–78 (1990).
48. Friston, K.J., Williams, S., Howard, R., Frackowiak, R.S. & Turner, R. Movement-related effects in fMRI time series. *Magn. Reson. Med.* **35**, 346–355 (1996).
49. Talairach, J. & Tournoux, P. *Co-Planar Stereotaxic Atlas of the Human Brain* (Thieme Medical Publishers, New York, 1988).
50. Ashburner, J. & Friston, K.J. Nonlinear spatial normalization using basis functions. *Hum. Brain Mapp.* **7**, 254–266 (1999).