

Research



Cite this article: Gardner MPH, Schoenbaum G, Gershman SJ. 2018 Rethinking dopamine as generalized prediction error. *Proc. R. Soc. B* **285**: 20181645.
<http://dx.doi.org/10.1098/rspb.2018.1645>

Received: 20 July 2018

Accepted: 2 November 2018

Subject Category:

Neuroscience and cognition

Subject Areas:

neuroscience, computational biology

Keywords:

reinforcement learning, successor representation, temporal difference learning

Author for correspondence:

Samuel J. Gershman

e-mail: gershman@fas.harvard.edu

Electronic supplementary material is available online at <https://dx.doi.org/10.6084/m9.figshare.c.4304231>.

Rethinking dopamine as generalized prediction error

Matthew P. H. Gardner¹, Geoffrey Schoenbaum^{1,2,3} and Samuel J. Gershman⁴

¹Intramural Research Program of the National Institute on Drug Abuse, NIH, Bethesda, MD, USA

²Department of Anatomy and Neurobiology, University of Maryland School of Medicine, Baltimore, MD, USA

³Department of Neuroscience, Johns Hopkins School of Medicine, Baltimore, MD, USA

⁴Department of Psychology and Center for Brain Science, Harvard University, Cambridge, MA, USA

SJG, 0000-0002-6546-3298

Midbrain dopamine neurons are commonly thought to report a reward prediction error (RPE), as hypothesized by reinforcement learning (RL) theory. While this theory has been highly successful, several lines of evidence suggest that dopamine activity also encodes sensory prediction errors unrelated to reward. Here, we develop a new theory of dopamine function that embraces a broader conceptualization of prediction errors. By signalling errors in both sensory and reward predictions, dopamine supports a form of RL that lies between model-based and model-free algorithms. This account remains consistent with current canon regarding the correspondence between dopamine transients and RPEs, while also accounting for new data suggesting a role for these signals in phenomena such as sensory preconditioning and identity unblocking, which ostensibly draw upon knowledge beyond reward predictions.

1. Introduction

The hypothesis that midbrain dopamine neurons report a reward prediction error (RPE)—the discrepancy between observed and expected reward—enjoys a seemingly unassailable accumulation of support from electrophysiology [1–5], calcium imaging [6,7], optogenetics [8–10], voltammetry [11,12] and human brain imaging [13,14]. The success of the RPE hypothesis is exciting because the RPE is precisely the signal a reinforcement learning (RL) system would need to update reward expectations [15,16]. Support for this RL interpretation of dopamine comes from findings that dopamine complies with basic postulates of RL theory [1], shapes the activity of downstream reward-predictive neurons in the striatum [11,17] and plays a causal role in the control of learning [8–10,13].

Despite these successes, however, there are a number of signs that this is not the whole story. First, it has long been known that dopamine neurons respond to novel or unexpected stimuli, even in the absence of changes in value [7,18–20]. While some theorists have tried to reconcile this observation with the RPE hypothesis by positing that value is affected by novelty [21] or uncertainty [22], others have argued that this response constitutes a distinct function of dopamine [23–25], possibly mediated by an anatomically segregated projection from midbrain to striatum [7]. A second challenge is that some dopamine neurons respond to aversive stimuli. If dopamine responses reflect RPEs, then one would expect aversive stimuli to *reduce* responses (as observed in some studies [26,27]). A third challenge is that dopamine activity is sensitive to movement-related variables such as action initiation and termination [28,29]. A fourth challenge is that dopamine activity [30] and its putative haemodynamic correlates [31] are influenced by information, such as changes in stimulus contingencies, that should in principle be invisible to a pure ‘model-free’ RL system that updates reward expectations using RPEs. This has led to elaborations of the RPE hypothesis according to which dopamine has access to some ‘model-based’ information, for examples in terms of probabilistic beliefs or samples from a model-based simulator [22,32–36].

While some of these puzzles can be resolved within the RPE framework by modifying assumptions about the inputs to and modulators of the RPE signal, recent findings have proven more unyielding. In this paper, we focus on three of these findings: (1) dopamine transients are necessary for learning induced by unexpected changes in the sensory features of expected rewards [37]; (2) dopamine neurons respond to unexpected changes in sensory features of expected rewards [38]; and (3) dopamine transients are both sufficient and necessary for learning stimulus–stimulus associations [39]. Taken together, these findings seem to contradict the RPE framework supported by so much other data.

Here we will suggest one possible way to reconcile the new and old findings, based on the idea that dopamine computes prediction errors over sensory features, much as was previously hypothesized for rewards. This sensory prediction error (SPE) hypothesis is motivated by normative considerations: SPEs can be used to estimate a predictive feature map known as the *successor representation* (SR) [40,41]. The key advantage of the SR is that it simplifies the computation of future rewards, combining the efficiency of model-free RL with some of the flexibility of model-based RL. Neural and behavioural evidence suggests that the SR is part of the brain’s computational repertoire [42,43], possibly subserved by the hippocampus [44,45]. Here, building on the pioneering work of Suri [46], we argue that dopamine transients previously understood to signal RPEs may instead constitute the SPE signal used to update the SR.

2. Theoretical framework

(a) The reinforcement learning problem

RL theories posit an environment in which an animal accumulates rewards as it traverses a sequence of ‘states’ governed by a transition function $T(s'|s)$, the probability of moving from state s to state s' , and a reward function $R(s)$, the expected reward in state s . The RL problem is to predict and optimize *value*, defined as the expected discounted future return (cumulative reward),

$$V(s_t) = \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k} \right], \quad (2.1)$$

where r_t is the reward received at time t in state s_t , and $\gamma \in [0, 1]$ is a discount factor that determines the weight of temporally distal rewards. Because the environment is assumed to obey the Markov property (transitions and rewards depend only on the current state), the value function can be written in a recursive form known as the *Bellman equation* [47]:

$$V(s_t) = \mathbb{E}[r_t + \gamma V(s_{t+1})]. \quad (2.2)$$

The Bellman equation allows us to define efficient RL algorithms for estimating values, as we explain next.

(b) Model-free and model-based learning

Model-free algorithms solve the RL problem by directly estimating V from interactions with the environment. The Bellman equation specifies a recursive consistency condition that the value estimate $\hat{V}(s_t)$ must satisfy in order to be accurate. By taking the difference between the two sides of the Bellman equation, $\mathbb{E}[r_t + \gamma \hat{V}(s_{t+1})] - \hat{V}(s_t)$, we can obtain a

measure of expected error; the direction and degree of the error is informative about how to correct $\hat{V}(s_t)$.

Because model-free algorithms do not have access to the underlying environment model (R and T) necessary to compute the expected error analytically, they typically rely on a stochastic sample of the error based on experienced transitions and rewards:

$$\delta_t = r_t + \gamma \hat{V}(s_{t+1}) - \hat{V}(s_t). \quad (2.3)$$

This quantity, commonly known as the temporal difference (TD) error, will on average be 0 when the value function has been perfectly estimated. The TD error is the basis of the classic TD learning algorithm [47], which in its simplest form updates the value estimate according to $\Delta \hat{V}(s_t) \propto \delta_t$. The RPE hypothesis states that dopamine reports the TD error [15,16].

Model-free algorithms like TD learning are efficient because they *cache* value estimates, which means that state evaluation (and by extension action selection) can be accomplished by simply inspecting the values cached in the relevant states. This efficiency comes at the cost of flexibility: if the reward function changes at a particular state, the entire value function must be re-estimated, since the Bellman equation implies a coupling of values between different states. For this reason, it has been proposed that the brain also makes use of model-based algorithms [48,49], which occupy the opposite end of the efficiency–flexibility spectrum. Model-based algorithms learn a model of the environment (R and T) and use this model to evaluate states, typically through some form of forward simulation or dynamic programming. This approach is flexible, because local changes in the reward or transition functions will instantly propagate across the entire value function, but at the cost of relying on comparatively inefficient simulation.

Some of the phenomena that we discuss in the Results section have been ascribed to model-based computations supported by dopamine [50], thus transgressing the clean boundary between the model-free function of dopamine and putatively non-dopaminergic model-based computations. The problem with this reformulation is that it is unclear what exactly dopamine is contributing to model-based learning. Although prediction errors are useful for updating estimates of the reward and transition functions used in model-based algorithms, these do not require a TD error. A distinctive feature of the TD error is that it bootstraps a future value estimate (the $\gamma \hat{V}(s_{t+1})$ term); this is necessary because of the Bellman recursion. But learning reward and transition functions in model-based algorithms can avoid bootstrapping estimates because the updates are local thanks to the Markov property.

To make this concrete, a simple learning algorithm (guaranteed to converge to the maximum-likelihood solution under some assumptions about the learning rate) is to update the model parameters according to

$$\Delta R(s) \propto r_t - R(s_t) \quad (2.4)$$

and

$$\Delta T(s' | s_t) \propto \mathbb{1}(s_{t+1} = s') - T(s' | s_t), \quad (2.5)$$

where $\mathbb{1}(\cdot) = 1$ if its argument is true, and 0 otherwise [51]. These updates can be understood in terms of prediction errors, but *not* TD errors (they do not bootstrap future value estimates). The TD interpretation is important for explaining phenomena like the shift in signalling to earlier reward-predicting cues [16], the temporal specificity of dopamine responses [52,53]

and the sensitivity to long-term values [54]. Thus, it remains mysterious how to retain the TD error interpretation of dopamine, which has been highly successful as an empirical hypothesis, while simultaneously accounting for the sensitivity of dopamine to SPEs.

(c) The successor representation

To reconcile these data, we will develop the argument that dopamine reflects sensory TD errors, encompassing both reward and non-reward features of a stimulus. In order to introduce some context to this idea, let us revisit the fundamental efficiency–flexibility trade-off. One way to find a middle-ground between the extremes occupied by model-free and model-based algorithms is to think about different ways to *compile* a model of the environment. By analogy with programming, a compiled program gains efficiency (in terms of runtime) at the expense of flexibility (the internal structure of the program is no longer directly accessible). Model-based algorithms are maximally uncompiled: they explicitly represent the parameters of the model, thus providing a representation that can be flexibly altered for new tasks. Model-free algorithms are maximally compiled: they only represent the summary statistics (state values) that are needed for reward prediction, bypassing a flexible representation of the environment in favour of computational efficiency.

A third possibility is a partially compiled model. Dayan [40] presented one such scheme, based on the following mathematical identity:

$$V(s_t) = \sum_{s'} M(s_t, s') R(s'), \quad (2.6)$$

where M denotes the SR, the expected discounted future state occupancy,

$$M(s_t, s') = \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k \mathbb{1}(s_{t+k} = s') \right]. \quad (2.7)$$

Intuitively, the SR represents states in terms of the frequency of their successor states. From a computational perspective, the SR is appealing for two reasons. First, it renders value computation a linear operation, yielding efficiency comparable to model-free evaluation. Second, it retains some of the flexibility of model-based evaluation. Specifically, changes in rewards will instantly affect values because the reward function is represented separately from the SR. On the other hand, the SR will be relatively insensitive to changes in transition structure, because it does not explicitly represent transitions—these have been compiled into a convenient but inflexible format. Behaviour reliant upon such a partially compiled model of the environment should be more sensitive to reward changes than transition changes, a prediction recently confirmed in humans [42].

The SR obeys a recursion analogous to the Bellman equation:

$$M(s_t, s') = \mathbb{E}[\mathbb{1}(s_t = s') + \gamma M(s_{t+1}, s')]. \quad (2.8)$$

Following the logic of the previous section, this implies that a TD learning algorithm can be used to estimate the SR,

$$\Delta \hat{M}(s_t, s') \propto \delta_t^M(s') = \mathbb{1}(s_t = s') + \gamma \hat{M}(s_{t+1}, s') - \hat{M}(s_t, s'), \quad (2.9)$$

where \hat{M} denotes the approximation of M .

One challenge facing this formulation is the *curse of dimensionality*: in large state spaces, it is impossible to accurately estimate the SR for all states. Generalization across states can be achieved by defining the SR over state features (indexed by j) and modelling this feature-based SR with linear function approximation,

$$\hat{M}(s_t, j) = \sum_i f_i(s_t) W_{ij}, \quad (2.10)$$

where $f_i(s)$ denotes the i th feature of state s and W is a weight matrix that parametrizes the approximation. In general, the features can be arbitrary, but for the purposes of this paper, we will assume that the features correspond to distinct stimulus identities; thus $f_i(s) = 1$ if stimulus i is present in state s , and 0 otherwise. Linear function approximation leads to the following learning rule for the weights:

$$\Delta W_{ij} \propto \delta_t^M(j) f_i(s_t), \quad (2.11)$$

where

$$\delta_t^M(j) = f_j(s_t) + \gamma \hat{M}(s_{t+1}, j) - \hat{M}(s_t, j) \quad (2.12)$$

is the TD error under linear function approximation. We will argue that dopamine encodes this TD error.

One issue with comparing this vector-valued TD error to experimental data is that we do not yet know how particular dopamine neurons map onto particular features. In order to make minimal assumptions, we will assume that each neuron has a uniform prior probability of encoding any given feature. Under ignorance about feature tuning, the expected TD error is then proportional to the superposition of feature-specific TD errors, $\sum_j \delta_t^M(j)$. In our simulations of dopamine, we take this superposition to be the ‘dopamine signal’ (see also [32]), but we wish to make clear that this is a provisional assumption that we ultimately hope to replace once the feature tuning of dopamine neurons is better understood.

There are several notable aspects of this new model of dopamine. First, it naturally captures SPEs, as we will illustrate shortly. Second, it also captures RPEs if reward is one of the features. Specifically, if $f_j(s_t) = r_t$, then the corresponding column of the SR is equivalent to the value function, $M(s, j) = V(s)$, and the corresponding TD error is the classical RPE, $\delta_t^M(j) = \delta_t$. Third, the TD error is now vector-valued, which means that dopamine neurons may be heterogeneously tuned to particular features (as hypothesized by some authors [55]), or they multiplex several features [56], or both. Notably, although the RPE correlate has famously been evident in single units, representation of these more complex or subtle prediction errors may be an ensemble property.

3. Simulations

Some of the most direct evidence for our hypothesis comes from a recent study by Chang *et al.* [37], who examined whether dopamine is necessary for learning about changes in reward identity (figure 1a). Animals first learned to associate two stimuli (X_B and X_{UB}) with different reward flavours. These stimuli were then reinforced in compound with other stimuli (A_B and A_{UB}). Critically, the $X_{UB}A_{UB}$ trials were accompanied by a change in reward flavour, a procedure known as ‘identity unblocking’ that attenuates the blocking effect [57–59]. This effect eludes explanation in terms of model-free mechanisms, but is naturally accommodated by

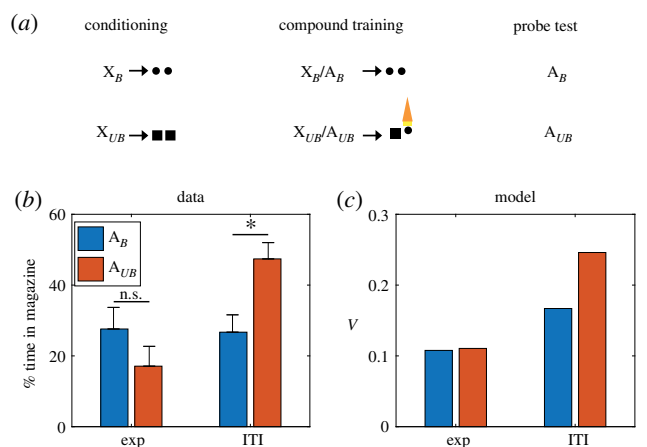


Figure 1. Inhibition of dopamine neurons prevents learning induced by changes in reward identity. (a) Identity unblocking paradigm. Circles and squares denote distinct reward flavours. Orange light symbol indicates when dopamine neurons were suppressed optogenetically to disrupt any positive SPE; this spanned a 5 s period beginning 500 ms prior to delivery of the second reward. (b) Conditioned responding on the probe test. Exp: experimental group, receiving inhibition during reward outcome. ITI: control group, receiving inhibition during the intertrial interval. Asterisk indicates significant difference ($p < 0.05$). Error bars show standard error of the mean. Data replotted from [37]. (c) Model simulation of the value function. (Online version in colour.)

the SR since changes in reward identity induce SPEs. Chang *et al.* [37] showed that optogenetic inhibition of dopamine at the time of the flavour change prevents this unblocking effect (figure 1b). Our model accounts for this finding (figure 1c), because inhibition suppresses SPEs that are necessary for driving learning.

One discrepant observation is a simulated increase in V in the ITI condition relative to the Exp condition for A_B , which does not appear in the experimental data. During the second stage of learning, X_{UB} , A_{UB} and the sensory features of both pellet types are presented together. Because of the co-occurrence of these features, associations develop between them such that the sensory features of the pellets now have slight associations with one another as well as the cues that predict them. This causes A_B and X_B to have a slight association with the sensory features of the pellet that it never predicted since both pellets now have mild associations with one another.

Electrophysiological experiments have confirmed that dopamine neurons respond to changes in identity, demonstrating a neural signal that is capable of explaining the data from Chang *et al.* [37]. We have already mentioned the sizable literature on novelty responses, but the significance of this activity is open to question, because the animal's prior value expectation is typically unclear. A study reported by Takahashi *et al.* [38] provides more direct evidence for an SPE signal, using a task (figure 2a) in which animals experience both shifts in value (amount of reward) and identity (reward flavour). Takahashi and colleagues found that individual dopamine neurons exhibited the expected changes in firing to shifts in value (figure 2b, reward addition and omission) and also showed a stronger response following a value-neutral change in reward identity (figure 2b, identity switch), changes in firing similar to those predicted by the model under these conditions (figure 2c).

A strong form of our proposal is that dopamine transients are both sufficient and necessary for learning stimulus–stimulus

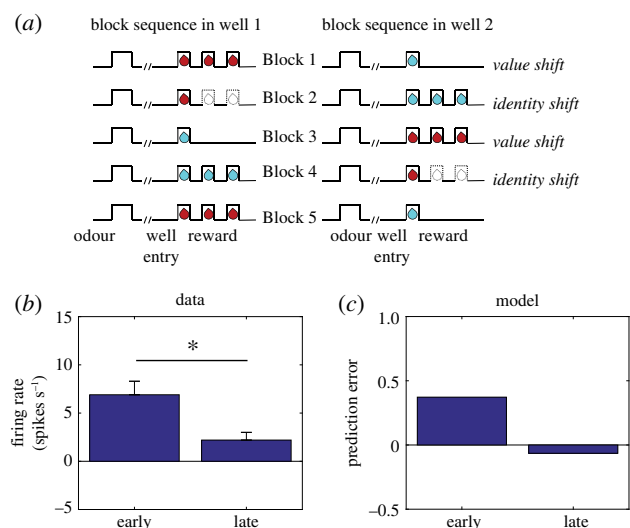


Figure 2. Dopamine neurons respond to changes in reward identity. (a) Time course of stimuli presented to the animal on each trial. Dotted line indicates reward omission, solid line indicates reward delivery. At the start of each session, one well was randomly designated as short (a 5 s delay before the reward) and the other long (a 1–7 s delay before the reward; see Block 1). In Block 2, these contingencies were switched. In Block 3, the delay was held constant, while the number of rewards was manipulated; one well was designated a big reward, in which a second bolus of reward was delivered (big reward), and a small (single bolus) reward was delivered in the other well. In Block 4, these contingencies were switched again. (b) Firing rate of dopamine neurons on trials that occurred early (first 5 trials) or late (last 5 trials) during an identity shift block. Error bars show standard error of the mean. Data replotted from [38]. (c) Model simulation of TD error. (Online version in colour.)

associations. Recent experiments using a sensory preconditioning paradigm [39] have tested this using sensory preconditioning. In this paradigm (figure 3a), various stimuli and stimulus compounds (denoted A, EF, AD, AC) are associated with another stimulus X through repeated pairing in an initial preconditioning phase. In a subsequent conditioning phase, X is associated with reward (sucrose pellets). In a final probe test, conditioned responding to a subset of the individual stimuli (F, D, C) is measured in terms of the number of food cup entries elicited by the presentation of the stimuli. During the preconditioning phase, one group of animals received optogenetic activation of dopamine neurons via channelrhodopsin (ChR2) expressed in the ventral tegmental area of the midbrain. In particular, optogenetic activation was applied either coincident with the onset of X on AC → X trials, or (as a temporal control) 120–180 s after X on AD → X trials. Another control group of animals received the same training and optogenetic activation, but expressed light-insensitive enhanced yellow fluorescent protein (eYFP).

A blocking effect was discernible in the control (eYFP) group, whereby A reduced acquisition of conditioned responding to C and D, compared to F, which was trained in compound with a novel stimulus (figure 3b). The blocking effect was eliminated by optogenetic activation in the experimental (ChR2) group, specifically for C, which received activation coincident with X. Thus, activation of dopamine neurons was sufficient to drive stimulus–stimulus learning in a temporally specific manner.

These findings raise a number of questions. First, how does one explain blocking of stimulus–stimulus associations?

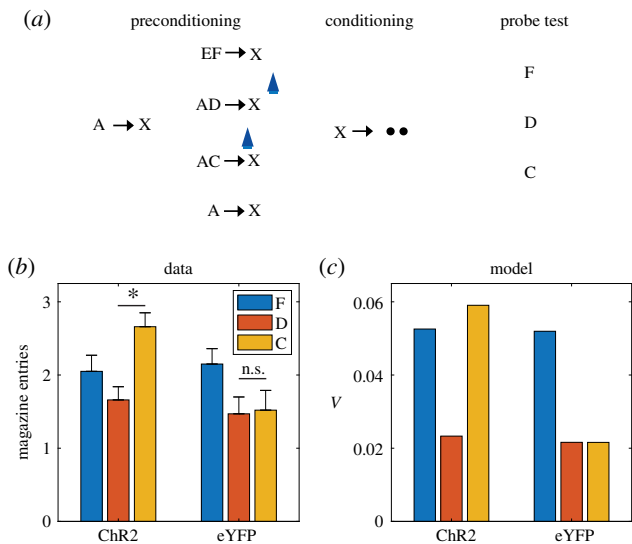


Figure 3. Dopamine transients are sufficient for learning stimulus–stimulus associations. (a) Sensory preconditioning paradigm. The initial preconditioning phase is broken down into two sub-phases. Letters denote stimuli, arrows denote temporal contingencies and circles denote rewards. Blue light symbol indicates when dopamine neurons were activated optogenetically to mimic a positive SPE; this spanned a 1 s period beginning at the start of X. (b) Number of food cup entries occurring during the probe test for experimental (Chr2) and control (eYFP) groups. Error bars show standard error of the mean. Data replotted from [39]. (c) Model simulation, using the value estimate as a proxy for conditioned responding. (Online version in colour.)

Second, how does one explain why dopamine affects this learning in the apparent absence of new reward information?

In answer to the first question, we can appeal to an analogy with blocking of stimulus–reward associations. The classic approach to modelling this phenomenon is to assume that each stimulus acquires an independent association and that these associations summate when the stimuli are presented in compound [60]. While there are boundary conditions on this assumption [61], it has proven remarkably successful at capturing a broad range of learning phenomenon, and is inherited by TD models with linear function approximation (e.g. [16,22,62]). Summation implies that if one stimulus (A) perfectly predicts reward, then a second stimulus (C) with no pre-existing association will fail to acquire an association when presented in compound with A, because the sum of the two associations will perfectly predict reward and hence generate an RPE of 0. The same logic can be applied to stimulus–stimulus learning by using linear function approximation of the SR, which implies that stimulus–stimulus associations will summate and hence produce blocking, as observed by Sharpe *et al.* [39].

In answer to the second question, we argue that dopamine is involved in stimulus–stimulus learning because it reflects a multifaceted SPE, as described in the previous section. By assuming that optogenetic activation adds a constant to the SPE (see Methods), we can capture the unblocking findings reported by Sharpe and colleagues (figure 3c). The mechanism by which optogenetic activation induces unblocking is essentially the same as the one suggested by the results of Steinberg *et al.* [9] for conventional stimulus–reward blocking: by elevating the prediction error, a learning signal is engendered where none would exist otherwise. However, while the results of Steinberg and colleagues are consistent with the original RPE hypothesis of dopamine, the results of Sharpe *et al.* [39] cannot be explained

by this model and instead require the analogous dopamine-mediated mechanism for driving learning with SPEs.

In addition to establishing the sufficiency of dopamine transients for learning, [39] also established their necessity, using optogenetic inactivation. In a variation of the sensory preconditioning paradigm (figure 4a), two pairs of stimulus–stimulus associations were learned (A → X and B → Y). Subsequently, X and Y were paired with different reward flavours, and finally conditioned responding to A and B was evaluated in a probe test. In one group of animals expressing halorhodopsin in dopamine neurons (NpHR), optogenetic inhibition was applied coincident with the transition between the stimuli on B → Y trials. A control group expressing light-insensitive eYFP was exposed to the same stimulation protocol. Sharpe and colleagues found that inhibition of dopamine selectively reduced responding to B (figure 4b), consistent with our model prediction that disrupting dopamine transients (a negative prediction error signal) should attenuate stimulus–stimulus learning (figure 4c).

4. Limitations and extensions

One way to drive a wedge between model-based and model-free algorithms is to devalue rewards (e.g. through pairing the reward with illness or selective satiation) and show effects on previously acquired conditioned responses to stimuli that predict those rewards. Because model-free algorithms like TD learning need to experience unbroken stimulus–reward sequences to update stimulus values, the behaviours they support are insensitive to such reward devaluation. Model-based algorithms, in contrast, are able to propagate the devaluation to the stimulus without direct experience, and hence allow behaviour to be devaluation-sensitive. Because of this, devaluation-sensitivity has frequently been viewed as an assay of model-based RL [48].

However, such sensitivity can also be a property of SR-based RL, since the SR represents the association between the stimulus and food, and is also able to update the reward function of the food as a result of devaluation. Thus, like model-based accounts, an SR model can account for changes in previously learned behaviour to reward-predicting stimuli after devaluation, both in normal situations [42,43] and when learning about those stimuli is unblocked by dopamine activation [63]. However, the SR model cannot spontaneously acquire transitions between states that are not directly experienced [42,43]. With this in mind, we consider the finding that reward devaluation alters the learning induced by activation of dopamine neurons in the sensory preconditioning paradigm of Sharpe *et al.* [39].

A key aspect of the reward devaluation procedure is that the food was paired with illness after the end of the entire preconditioning procedure and in the absence of any of the stimuli (and in fact not in the training chamber). In the SR model, only stimuli already predictive of the food can change their values after devaluation. In the paradigm of Sharpe and colleagues, X was associated with food but C was not. Moreover, C was associated with X before any association with food was established. Because of this, C is not updated in the SR model to incorporate an association with food. It follows that, unlike the animals in Sharpe *et al.* [39], the model will not be devaluation-sensitive when probed with C (figure 5c).

It is possible to address this failure within our theoretical framework in a number of different ways. One way we

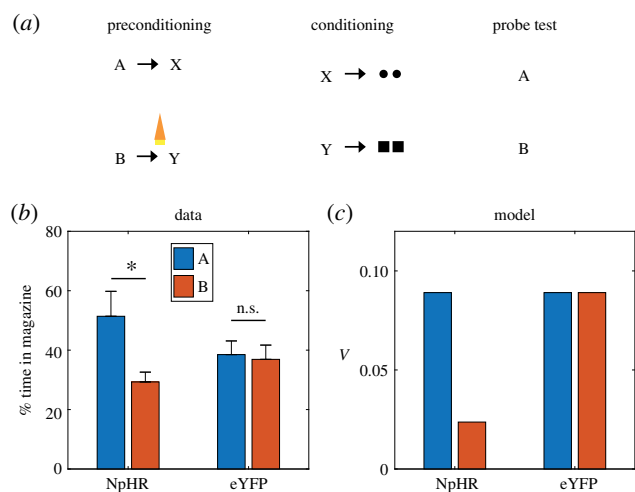


Figure 4. Dopamine transients are necessary for learning stimulus–stimulus associations. (a) Sensory preconditioning paradigm. Circles and squares denote distinct reward flavours. Orange light symbol indicates when dopamine neurons were suppressed optogenetically to disrupt any positive SPE; this spanned a 2.5 s period beginning 500 ms prior to the end of B. (b) Number of food cup entries occurring during the probe test for experimental (NpHR) and control (eYFP) groups. Error bars show standard error of the mean. Data replotted from [39]. (c) Model simulation. (Online version in colour.)

considered was to allow optogenetic activation to increment predictions for *all* possible features, instead of being restricted to recently active features by a feature eligibility trace (see Methods), as in the simulations thus far. With such a promiscuous artificial error signal, the model can recapitulate the devaluation effect, because C would then become associated with food (along with everything else) in the preconditioning phase itself. The problem with this work-around is that it also predicts that animals should develop a conditioned response to the food cup for all the cues during preconditioning, since food cup shaping prior to preconditioning seeds the food state with reward value. As a result, any cue paired with the food state immediately begins to induce responding at the food cup. Such behaviour is not observed, suggesting that the artificial update caused by optogenetic activation of the dopamine neurons is locally restricted.

A second, more conventional way to address this failure within our theoretical framework is to assume that there is some form of offline rehearsal or simulation that is used to update cached predictions [33,64,65]. Russek *et al.* [43] have shown that such a mechanism is able to endow SR-based learning with the ability to retrospectively update predictions even in the absence of direct experience. A minimal implementation of such a mechanism in our model, simply by ‘confabulating’ the presence of X during reward devaluation, is sufficient to capture the effects of devaluation following optogenetic activation of dopamine neurons (figure 5d). This solution makes the experimental prediction that the devaluation-sensitivity of this artificially unblocked cue should be time-dependent, under the assumption that the amount of offline rehearsal is proportional to the retention interval.

5. Discussion

The RPE hypothesis of dopamine has been one of theoretical neuroscience’s signature success stories. This paper has set forth a significant generalization of the RPE hypothesis

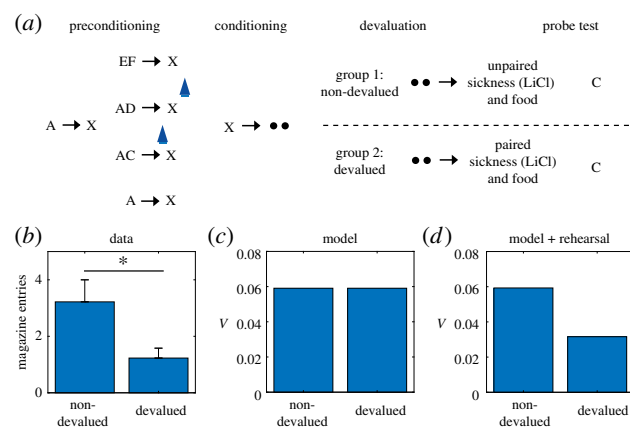


Figure 5. Behaviour to preconditioned cue that is unblocked by activation of dopamine neurons is sensitive to devaluation of the predicted reward. (a) Data replotted from [39] and (b) model simulation for conditioned responding to stimulus C in the probe test. Animals in the devalued group were injected with lithium chloride in conjunction with ingestion of the reward (sucrose pellets), causing a strong aversion to the reward. Animals in the non-devalued group were injected with lithium chloride approximately 6 h after ingestion of the reward. Error bars show standard error of the mean. (c) A version of the model with rehearsal of stimulus X during reward devaluation was able to capture the devaluation-sensitivity of animals. (Online version in colour.)

that enables it to account for a number of anomalous phenomena, without discarding the core ideas that motivated the original hypothesis. The proposal that dopamine reports a SPE is grounded in a normative theory of RL [40], motivated independently by a number of computational [43,66], behavioural [42,67,68] and neural [44,45,69,70] considerations.

An important strength of the proposal is that it extends the functional role of dopamine beyond RPEs, while still accounting for the data that motivated the original RPE hypothesis. This is because, if reward is treated as a sensory feature, then one dimension of the vector-valued SPE will be the RPE. Indeed, dopamine SPEs should behave systematically like RPEs, except that they respond to features: they should pause when expected features are unexpectedly omitted, they should shift back to the earliest feature-predicting cue, and they should exhibit signatures of cue competition, such as over-expectation. SPEs are used to update cached predictions, analogous to the RPE in model-free algorithms. However, these cached predictions extend beyond value to include information about the occupancy of future states (the SR). The SR can be used in a semi-flexible manner that allows behaviour to be sensitive to changes in the reward structure, such as devaluation by pairing a reward with illness. As a result, even if dopamine is constrained by the model proposed here, it would support significantly more flexible behaviour than supposed by classical model-free accounts [15,16], even without moving completely to an account of model-based computation in the dopamine system [50].

Nevertheless, the theory proposed here—particularly if it incorporates offline rehearsal in order to fully explain the results of Sharpe *et al.* [39]—does strain the dichotomy between model-based and model-free algorithms that has been at the heart of modern RL theories [48]. However, as noted earlier, SR requires offline rehearsal to incorporate the effects of devaluation after preconditioning in Sharpe *et al.* or manipulations of the transition structures of tasks [42]. If these effects, and particularly dopamine’s involvement in them, are mediated by an

SR mechanism, then we should be able to interfere with it by manipulating retention intervals or attention [33]. For example, the strength of devaluation sensitivity in Sharpe *et al.* should be diminished by a very short retention interval prior to the probe test, since this would reduce time available for rehearsal. If these effects do not show any dependence on the length of the retention interval, then this would be more consistent with model-based algorithms, which do not require any rehearsal.

Another testable prediction of the theory is that we should see heterogeneity in the dopamine response, reflecting the vector-valued nature of the SPE. Importantly, such tuning need not be statistically evident in the spiking of an individual neuron. It might show up in the pattern of response across the entire population or even in subpopulations determined by target or other criteria. Indeed, target-based heterogeneity is already evident in some studies of dopamine release or function in downstream regions [63,71,72]. Related to this, the theory also predicts the existence of a negative SPE to allow reductions in the strength of weights in the SR. In its simplest form, the omission of an expected stimulus could result in suppression of firing, analogous to reward omission responses [16]. However, this effect might be subtle if SPEs are population-coded by the dopamine signal, as suggested above; the negative SPE may simply reflect a particular pattern across the population rather than overt suppression at the level of single neurons. Distinctive patterns of activity identifying the source of the error and differentiating the addition of information versus its omission sets our proposal apart from explanations based on salience signals, which are typically thought to be both non-specific and unsigned [73]. These predictions set an exciting new agenda for dopamine research by embracing a broader conception of dopamine function.

While we have focused on dopamine in this paper, a complete account will obviously need to integrate the computational functions of other brain regions. Where does information relevant to computing SPE's come from? One obvious possibility is from sensory regions. Sensory areas respond both to and in expectation of external events [74,75], and these areas send input to brainstem, thus they are positioned to feed information to the dopaminergic system. Beyond this, the hippocampus and orbitofrontal cortex seem likely to be particularly important. Many lines of evidence are consistent with the idea that the hippocampus encodes a 'predictive map' resembling the SR [44]. For example, hippocampal place cells alter their tuning with repeated experience to fire in anticipation of future locations [76], and fMRI studies have found predictive coding of non-spatial states [45,77]. The orbitofrontal cortex has also been repeatedly implicated in predictive coding, particularly of reward outcomes [78,79], but also of sensory events [80], and the orbitofrontal cortex is critical for sensory-specific outcome expectations in Pavlovian conditioning [81]. Wilson *et al.* [82] have proposed that the orbitofrontal cortex encodes a 'cognitive map' of state space, which presumably underpins this diversity of stimulus expectations. Thus, evidence suggests that both hippocampus and orbitofrontal cortex encode some form of predictive representation [83]. Further, dopaminergic modulation of these regions is well-established [84,85]. It is tempting to speculate that afferent input from and dopaminergic modulation of the hippocampus and orbitofrontal cortex may be especially critical to the SPE function proposed here.

The influence of these representations may be filtered through interactions with more basic value representations in striatum. This proposal fits with the observation that the hippocampus and orbitofrontal cortex appear to confer stimulus specificity on value-sensitive neurons in the striatum [86,87]. Striatal value representations are already proposed to influence activity in VTA [88]. By this model, dopamine would still provide the RPE signal that drives striatal plasticity related to actions or 'value', as in most contemporary accounts, but in addition it would provide an SPE signal to update associative representations, perhaps in striatum but also in upstream orbitofrontal and hippocampal areas, which feed into the striatum. While speculative, this idea is consistent with findings showing heterogeneity of dopamine function based on projection target, at least within striatum [63,71]. It is also consistent with recent human imaging work, confirming the presence of an SPE-like signal in human VTA, and reporting that the strength of this signal during learning is correlated with the strength of new sensory-sensory correlates developed in the orbitofrontal cortex [70].

The view that dopamine reports the SR prediction error provides a bridge between sensory and RPE accounts of dopamine function. The tension between these views has long vexed computational theories, and has posed particular problems for pure RPE accounts of dopamine. We see our model as the first step towards resolving this tension. While we have shown that the notion of a generalized prediction error is consistent with a wealth of empirical data, this is just the beginning of the empirical enterprise. Armed with a quantitative framework, we can now pursue evidence for such prediction errors with greater precision and clarity.

6. Methods

(a) Linear value function approximation

Under the linear function approximation scheme described in the Results, the value function estimate is given by

$$\hat{V}(s_t) = \sum_i f_i(s_t) \sum_j U(j) W_{ij}, \quad (6.1)$$

where $U(j)$ is the reward expectation for feature j , updated according to a delta rule,

$$\Delta U(j) = \alpha_U f_j(s_t) [r_t - \hat{V}(s_t)], \quad (6.2)$$

with learning rate α_U .

In the electronic supplementary material's figures, we report simulations of the value-based TD learning algorithm, TD(0), which approximates the value function using linear function approximation

$$\hat{V}(s_t) \approx \sum_j w_{jf}(s_t), \quad (6.3)$$

and updates the weights according to

$$\Delta w_j = \alpha f_j(s_t) [r_t + \gamma \hat{V}(s_{t+1}) - \hat{V}(s_t)]. \quad (6.4)$$

(b) Excitatory and inhibitory asymmetry in the TD error term

There is a large body of evidence in associative learning suggesting an imbalance between excitatory and inhibitory learning [89,90]. Mirroring this imbalance is an asymmetry

in the dynamic range of the firing rate of single dopaminergic neurons in the midbrain [2]. In accordance with these observations, we assume that the error terms (ΔW_{ij} and ΔU_i) are rescaled by a factor of $\frac{1}{4}$ for negative prediction errors. This is equivalent to assuming separate learning rates for positive and negative prediction errors [91]. Note that, following prior theoretical work (e.g. [16]), we consider negative prediction errors to be coded by real neurons relative to a baseline firing rate, acknowledging the fact that neurons cannot produce negative firing rates.

(c) Simulation parameters

We used the following parameters in the simulations of SR: $\gamma = 0.95$, $\alpha_W = 0.06$, $\alpha_U = 0.03$, where α_W is the learning rate for the weight matrix W , α_U is the learning rate for the reward function, and γ is the discount rate. For the model-free TD learning algorithm simulations, we used the following parameters: $\gamma = 0.95$, $\alpha = 0.05$. We used the same set of parameters across all simulations. However, our results are largely robust to variations in these parameters.

(d) Modelling optogenetic activation and inhibition

Optogenetic intervention was modelled by modifying the TD error as follows:

$$\delta_t^M(j) = \begin{cases} (1 + \eta)\delta_t^M(j) & \eta < 0 \\ [f_j(s_t)\eta + \delta_t^M(j)] & \eta > 0, \end{cases} \quad (6.5)$$

References

1. Waelti P, Dickinson A, Schultz W. 2001 Dopamine responses comply with basic assumptions of formal learning theory. *Nature* **412**, 43–48. (doi:10.1038/35083500)
2. Bayer HM, Glimcher PW. 2005 Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* **47**, 129–141. (doi:10.1016/j.neuron.2005.05.020)
3. Roesch MR, Calu DJ, Schoenbaum G. 2007 Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nat. Neurosci.* **10**, 1615–1624. (doi:10.1038/nn2013)
4. Eshel N, Bukwich M, Rao V, Hemmelder V, Tian J, Uchida N. 2015 Arithmetic and local circuitry underlying dopamine prediction errors. *Nature* **525**, 243–246. (doi:10.1038/nature14855)
5. Eshel N, Tian J, Bukwich M, Uchida N. 2016 Dopamine neurons share common response function for reward prediction error. *Nat. Neurosci.* **19**, 479–486. (doi:10.1038/nn.4239)
6. Parker NF, Cameron CM, Taliaferro JP, Lee J, Choi JY, Davidson TJ, Daw ND, Witten IB. 2016 Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nat. Neurosci.* **19**, 845–854. (doi:10.1038/nn.4287)
7. Menegas W, Babayan BM, Uchida N, Watabe-Uchida M. 2017 Opposite initialization to novel cues in dopamine signaling in ventral and posterior striatum in mice. *elife* **6**, e21886. (doi:10.7554/eLife.21886)
8. Tsai H-C, Zhang F, Adamantidis A, Stuber GD, Bonci A, De Lecea L, Deisseroth K. 2009 Phasic firing in dopaminergic neurons is sufficient for behavioral conditioning. *Science* **324**, 1080–1084. (doi:10.1126/science.1168878)
9. Steinberg EE, Keiflin R, Boivin JR, Witten IB, Deisseroth K, Janak PH. 2013 A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* **16**, 966–973. (doi:10.1038/nn.3413)
10. Chang CY, Esber GR, Marrero-Garcia Y, Yau H-J, Bonci A, Schoenbaum G. 2016 Brief optogenetic inhibition of dopamine neurons mimics endogenous negative reward prediction errors. *Nat. Neurosci.* **19**, 111–116. (doi:10.1038/nn.4191)
11. Day JJ, Roitman MF, MarkWightman R, Carelli RM. 2007 Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. *Nat. Neurosci.* **10**, 1020–1028. (doi:10.1038/nn1923)
12. Hart AS, Rutledge RB, Glimcher PW, Phillips PEM. 2014 Phasic dopamine release in the rat nucleus accumbens symmetrically encodes a reward prediction error term. *J. Neurosci.* **34**, 698–704. (doi:10.1523/JNEUROSCI.2489-13.2014)
13. Pessiglione M, Seymour B, Flandin G, Dolan RJ, Frith CD. 2006 Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* **442**, 1042–1045. (doi:10.1038/nature05051)
14. D'ardenne K, McClure SM, Nystrom LE, Cohen JD. 2008 BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* **319**, 1264–1267. (doi:10.1126/science.1150605)
15. Read Montague P, Dayan P, Sejnowski TJ. 1996 A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *J. Neurosci.* **16**, 1936–1947. (doi:10.1523/JNEUROSCI.16-05-01936.1996)
16. Schultz W, Dayan P, Read Montague P. 1997 A neural substrate of prediction and reward. *Science* **275**, 1593–1599. (doi:10.1126/science.275.5306.1593)
17. Cheer JF, Aragona BJ, Heien MLAV, Seipel AT, Carelli RM, MarkWightman R. 2007 Coordinated accumbal dopamine release and neural activity drive goal-directed behavior. *Neuron* **54**, 237–244. (doi:10.1016/j.neuron.2007.03.021)
18. Strecker RE, Jacobs BL. 1985 Substantia nigra dopaminergic unit activity in behaving cats: effect of arousal on spontaneous discharge and sensory evoked activity. *Brain Res.* **361**, 339–350. (doi:10.1016/0006-8993(85)91304-6)
19. Ljungberg T, Apicella P, Schultz W. 1992 Responses of monkey dopamine neurons during learning of behavioral reactions. *J. Neurophysiol.* **67**, 145–163. (doi:10.1152/jn.1992.67.1.145)
20. Horvitz JC. 2000 Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* **96**, 651–656. (doi:10.1016/S0306-4522(00)00019-1)
21. Kakade S, Dayan P. 2002 Dopamine: generalization and bonuses. *Neural Netw.* **15**, 549–559. (doi:10.1016/S0893-6080(02)00048-5)
22. Gershman SJ. 2017 Dopamine, inference, and uncertainty. *Neural Comput.* **29**, 3311–3326. (doi:10.1162/neco_a_01023)

where $\eta = 1.0$ for optogenetic activation and -0.8 for inhibition. The asymmetry between the functions for activation and inactivation was chosen to better match the hypothesized function of optogenetic stimulation based on empirical findings. For positive stimulation of dopamine, it is thought that the increased dopamine activity should enhance learning with the currently active features, which in the SR model is the $f_j(s_t)$ term. For optogenetic inhibition of dopamine, we have found that punctate versus prolonged inhibition causes differential effects, with punctate inhibition resulting in negative prediction errors and prolonged inhibition resulting in shunting of the error signal [10]. Our inhibition in the experiments included in this paper were prolonged, necessitating a different model of the inhibitory optogenetic manipulation.

Data accessibility. All simulation code is available at <https://github.com/mpgargner/TDSR>.

Authors' contributions. M.P.H.G., S.J.G. and G.S. conceived the ideas and wrote the manuscript. M.P.H.G. and S.J.G. carried out the model simulations and created the figures.

Competing interests. We have no competing interests.

Funding. This work was supported by the National Institutes of Health (CRCNS 1R01MH109177 to S.J.G.) and the Intramural Research Programat NIDA ZIA-DA000587 (to G.S.).

Acknowledgements. We are grateful to Yael Niv, Mingyu Song, Brian Sadacca and Andrew Wikenheiser for helpful discussions.

Disclaimer. The opinions expressed in this article are the authors' own and do not reflect the view of the NIH/DHHS.

23. Redgrave P, Gurney K, Reynolds J. 2008 What is reinforced by phasic dopamine signals? *Brain Res. Rev.* **58**, 322–339. (doi:10.1016/j.brainresrev.2007.10.007)
24. Bromberg-Martin ES, Matsumoto M, Hikosaka O. 2010 Dopamine in motivational control: rewarding, aversive, and alerting. *Neuron* **68**, 815–834. (doi:10.1016/j.neuron.2010.11.022)
25. Schultz W. 2016 Dopamine reward prediction-error signalling: a two-component response. *Nat. Rev. Neurosci.* **17**, 183–195. (doi:10.1038/nrn.2015.26)
26. Mirenowicz J, Schultz W. 1996 Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature* **379**, 449–451. (doi:10.1038/379449a0)
27. Ungless MA, Magill PJ, Paul Bolam J. 2004 Uniform inhibition of dopamine neurons in the ventral tegmental area by aversive stimuli. *Science* **303**, 2040–2042. (doi:10.1126/science.1093360)
28. Syed ECJ, Grima LL, Magill PJ, Bogacz R, Brown P, Walton ME. 2016 Action initiation shapes mesolimbic dopamine encoding of future rewards. *Nat. Neurosci.* **19**, 34–36. (doi:10.1038/nn.4187)
29. Jin X, Costa RM. 2010 Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature* **466**, 457–462. (doi:10.1038/nature09263)
30. Bromberg-Martin ES, Matsumoto M, Hong S, Hikosaka O. 2010 A pallidum–habenula–dopamine pathway signals inferred stimulus values. *J. Neurophysiol.* **104**, 1068–1076. (doi:10.1152/jn.00158.2010)
31. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. 2011 Model-based influences on humans' choices and striatal prediction errors. *Neuron* **69**, 1204–1215. (doi:10.1016/j.neuron.2011.02.027)
32. Daw ND, Courville AC, Touretzky DS. 2006 Representation and timing in theories of the dopamine system. *Neural Comput.* **18**, 1637–1677. (doi:10.1162/neco.2006.18.7.1637)
33. Gershman SJ, Markman AB, Ross Otto A. 2014 Retrospective revaluation in sequential decision making: a tale of two systems. *J. Exp. Psychol. Gen.* **143**, 182–194. (doi:10.1037/a0030844)
34. Kwon Starkweather C, Babayan BM, Uchida N, Gershman SJ. 2017 Dopamine reward prediction errors reflect hidden-state inference across time. *Nat. Neurosci.* **20**, 581–589. (doi:10.1038/nn.4520)
35. Nakahara H, Itoh H, Kawagoe R, Takikawa Y, Hikosaka O. 2004 Dopamine neurons can represent context-dependent prediction error. *Neuron* **41**, 269–280. (doi:10.1016/S0896-6273(03)00869-9)
36. Nakahara H, Hikosaka O. 2012 Learning to represent reward structure: a key to adapting to complex environments. *Neurosci. Res.* **74**, 177–183. (doi:10.1016/j.neures.2012.09.007)
37. Chang CY, Gardner M, Di Tillio MG, Schoenbaum G. 2017 Optogenetic blockade of dopamine transients prevents learning induced by changes in reward features. *Curr. Biol.* **27**, 3480–3486. (doi:10.1016/j.cub.2017.09.049)
38. Takahashi YK, Batchelor HM, Liu B, Khanna A, Morales M, Schoenbaum G. 2017 Dopamine neurons respond to errors in the prediction of sensory features of expected rewards. *Neuron* **95**, 1395–1405. (doi:10.1016/j.neuron.2017.08.025)
39. Sharpe MJ, Chang CY, Liu MA, Batchelor HM, Mueller LE, Jones JL, Niv Y, Schoenbaum G. 2017 Dopamine transients are sufficient and necessary for acquisition of model-based associations. *Nat. Neurosci.* **20**, 735–742. (doi:10.1038/nn.4538)
40. Dayan P. 1993 Improving generalization for temporal difference learning: the successor representation. *Neural Comput.* **5**, 613–624. (doi:10.1162/neco.1993.5.4.613)
41. Gershman SJ. 2018 The successor representation: its computational logic and neural substrates. *J. Neurosci.* **38**, 7193–7200. (doi:10.1523/JNEUROSCI.0151-18.2018)
42. Momennejad I, Russek EM, Cheong JH, Botvinick MM, Daw N, Gershman SJ. 2017 The successor representation in human reinforcement learning. *Nat. Hum. Behav.* **1**, 680–692. (doi:10.1038/s41562-017-0180-8)
43. Russek EM, Momennejad I, Botvinick MM, Gershman SJ, Daw ND. 2017 Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput. Biol.* **13**, e1005768. (doi:10.1371/journal.pcbi.1005768)
44. Stachenfeld KL, Botvinick MM, Gershman SJ. 2017 The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653. (doi:10.1038/nn.4650)
45. Garvert MM, Dolan RJ, Behrens TEJ. 2017 A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife* **6**, e17086. (doi:10.7554/eLife.17086)
46. Suri RE. 2001 Anticipatory responses of dopamine neurons and cortical neurons reproduced by internal model. *Exp. Brain Res.* **140**, 234–240. (doi:10.1007/s002210100814)
47. Sutton RS, Barto AG. 1998 *Reinforcement learning: an introduction*. New York, NY: MIT Press.
48. Daw ND, Niv Y, Dayan P. 2005 Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711. (doi:10.1038/nn1560)
49. Daw ND, Dayan P. 2014 The algorithmic anatomy of model-based evaluation. *Phil. Trans. R. Soc. B* **369**, 20130478. (doi:10.1098/rstb.2013.0478)
50. Langdon AJ, Sharpe MJ, Schoenbaum G, Niv Y. 2018 Model-based predictions for dopamine. *Curr. Opin. Neurobiol.* **49**, 1–7. (doi:10.1016/j.conb.2017.10.006)
51. Gläscher J, Daw N, Dayan P, O'Doherty JP. 2010 States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* **66**, 585–595. (doi:10.1016/j.neuron.2010.04.016)
52. Takahashi YK, Langdon AJ, Niv Y, Schoenbaum G. 2016 Temporal specificity of reward prediction errors signaled by putative dopamine neurons in rat vta depends on ventral striatum. *Neuron* **91**, 182–193. (doi:10.1016/j.neuron.2016.05.015)
53. Hollerman JR, Schultz W. 1998 Dopamine neurons report an error in the temporal prediction of reward during learning. *Nat. Neurosci.* **1**, 304–309. (doi:10.1038/1124)
54. Enomoto K, Matsumoto N, Nakai S, Satoh T, Sato TK, Ueda Y, Inokawa H, Haruno M, Kimura M. 2011 Dopamine neurons learn to encode the long-term value of multiple future rewards. *Proc. Natl Acad. Sci. USA* **108**, 15 462–15 467. (doi:10.1073/pnas.1014457108)
55. Lau B, Monteiro T, Paton JJ. 2017 The many worlds hypothesis of dopamine prediction error: implications of a parallel circuit architecture in the basal ganglia. *Curr. Opin. Neurobiol.* **46**, 241–247. (doi:10.1016/j.conb.2017.08.015)
56. Tian J, Huang R, Cohen JY, Osakada F, Kobak D, Machens CK, Callaway EM, Uchida N, Watabe-Uchida M. 2016 Distributed and mixed information in monosynaptic inputs to dopamine neurons. *Neuron* **91**, 1374–1389. (doi:10.1016/j.neuron.2016.08.018)
57. McDannald MA, Lucantonio F, Burke KA, Niv Y, Schoenbaum G. 2011 Ventral striatum and orbitofrontal cortex are both required for model-based, but not model-free, reinforcement learning. *J. Neurosci.* **31**, 2700–2705. (doi:10.1523/JNEUROSCI.5499-10.2011)
58. Rescorla RA. 1999 Learning about qualitatively different outcomes during a blocking procedure. *Learn. Behav.* **27**, 140–151. (doi:10.3758/BF03199671)
59. Blaisdell AP, Denniston JC, Miller RR. 1997 Unblocking with qualitative change of unconditioned stimulus. *Learn. Motiv.* **28**, 268–279. (doi:10.1006/lmot.1996.0961)
60. Rescorla RA, Wagner AR. 1972 A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: current research and theory* (eds AH Black, WF Prokasy), pp. 64–99. New York, NY: Appleton-Century-Crofts.
61. Soto FA, Gershman SJ, Niv Y. 2014 Explaining compound generalization in associative and causal learning through rational principles of dimensional generalization. *Psychol. Rev.* **121**, 526–558. (doi:10.1037/a0037018)
62. Ludvig EA, Sutton RS, James Kehoe E. 2008 Stimulus representation and the timing of reward-prediction errors in models of the dopamine system. *Neural Comput.* **20**, 3034–3054. (doi:10.1162/neco.2008.11-07-654)
63. Keiflin R, Pribut HJ, Shah NB, Janak PH. 2017 Phasic activation of ventral tegmental, but not substantia nigra, dopamine neurons promotes model-based pavlovian reward learning. *bioRxiv*.
64. Johnson A, David Redish A. 2005 Hippocampal replay contributes to within session learning in a temporal difference reinforcement learning model. *Neural Netw.* **18**, 1163–1171. (doi:10.1016/j.neunet.2005.08.009)
65. Pezzulo G, van der Meer MAA, Lansink CS, Pennartz CMA. 2014 Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* **18**, 647–657. (doi:10.1016/j.tics.2014.06.011)
66. Barreto A, Dabney W, Munos R, Hunt JJ, Schaul T, Silver D, van Hasselt HP. 2017 Successor features for

- transfer in reinforcement learning. 31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA. See <https://papers.nips.cc/paper/6994-successor-features-for-transfer-in-reinforcement-learning.pdf>.
67. Gershman SJ, Moore CD, Todd MT, Norman KA, Sederberg PB. 2012 The successor representation and temporal context. *Neural Comput.* **24**, 1553–1568. (doi:10.1162/NECO_a_00282)
 68. Smith TA, Hasinski AE, Sederberg PB. 2013 The context repetition effect: predicted events are remembered better, even when they don't happen. *J. Exp. Psychol. Gen.* **142**, 1298–1308. (doi:10.1037/a0034067)
 69. Brea J, Taml A, Urbanczik R, Senn W. 2016 Prospective coding by spiking neurons. *PLoS Comput. Biol.* **12**, e1005003. (doi:10.1371/journal.pcbi.1005003)
 70. Howard JD, Kahnt T. 2018 Identity prediction errors in the human midbrain update reward-identity expectations in the orbitofrontal cortex. *Nat. Commun.* **9**, 1611. (doi:10.1038/s41467-018-04055-5)
 71. Saunders BT, Richard JM, Margolis EB, Janak PH. 2018 Dopamine neurons create pavlovian conditioned stimuli with circuit-defined motivational properties. *Nat. Neurosci.* **21**, 1072–1083. (doi:10.1038/s41593-018-0191-4)
 72. Engelhard B *et al.* 2018 Specialized and spatially organized coding of sensory, motor, and cognitive variables in midbrain dopamine neurons. *bioRxiv*.
 73. Roesch MR, Esber GR, Li J, Daw ND, Schoenbaum G. 2012 Surprise! neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. *Eur. J. Neurosci.* **35**, 1190–1200. (doi:10.1111/j.1460-9568.2011.07986.x)
 74. Zelano C, Mohanty A, Gottfried JA. 2011 Olfactory predictive codes and stimulus templates in piriform cortex. *Neuron* **72**, 178–187.
 75. Gardner MPH, Fontanini A. 2014 Encoding and tracking of outcome-specific expectancy in the gustatory cortex of alert rats. *J. Neurosci.* **34**, 13 000–13 017.
 76. Mehta MR, Quirk MC, Wilson MA. 2000 Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* **25**, 707–715. (doi:10.1016/S0896-6273(00)81072-7)
 77. Schapiro AC, Turk-Browne NB, Norman KA, Botvinick MM. 2016 Statistical learning of temporal community structure in the hippocampus. *Hippocampus* **26**, 3–8. (doi:10.1002/hipo.22523)
 78. Schoenbaum G, Chiba AA, Gallagher M. 1998 Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat. Neurosci.* **1**, 155–159. (doi:10.1038/407)
 79. Gottfried JA, O'Doherty J, Dolan RJ. 2003 Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* **301**, 1104–1107. (doi:10.1126/science.1087919)
 80. Chaumon M, Kveraga K, Feldman Barrett L, Bar M. 2013 Visual predictions in the orbitofrontal cortex rely on associative content. *Cereb. Cortex* **24**, 2899–2907. (doi:10.1093/cercor/bht146)
 81. Ostlund SB, Balleine BW. 2007 Orbitofrontal cortex mediates outcome encoding in Pavlovian but not instrumental conditioning. *J. Neurosci.* **27**, 4819–4825. (doi:10.1523/JNEUROSCI.5443-06.2007)
 82. Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. 2014 Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267–279. (doi:10.1016/j.neuron.2013.11.005)
 83. Wikenheiser AM, Schoenbaum G. 2016 Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nat. Rev. Neurosci.* **17**, 513–523. (doi:10.1038/nrn.2016.56)
 84. Aou S, Oomura Y, Nishino H, Inokuchi A, Mizuno Y. 1983 Influence of catecholamines on reward-related neuronal activity in monkey orbitofrontal cortex. *Brain Res.* **267**, 165–170. (doi:10.1016/0006-8993(83)91052-1)
 85. Lisman JE, Grace AA. 2005 The hippocampal-VTA loop: controlling the entry of information into long-term memory. *Neuron* **46**, 703–713. (doi:10.1016/j.neuron.2005.05.002)
 86. Johnson A, van der Meer MAA, David Redish A. 2007 Integrating hippocampus and striatum in decision-making. *Curr. Opin Neurobiol.* **17**, 692–697. (doi:10.1016/j.conb.2008.01.003)
 87. Cooch NK, Stalnaker TA, Wied HM, Bali-Chaudhary S, McDannald MA, Liu T-L, Schoenbaum G. 2015 Orbitofrontal lesions eliminate signalling of biological significance in cue-responsive ventral striatal neurons. *Nat. Commun.* **6**, 7195. (doi:10.1038/ncomms8195)
 88. Joel D, Niv Y, Ruppin E. 2002 Actor–critic models of the basal ganglia: new anatomical and computational perspectives. *Neural Netw.* **15**, 535–547. (doi:10.1016/S0893-6080(02)00047-3)
 89. Konorski J 1948 *Conditioned reflexes and neuron organization*. Cambridge, UK: Cambridge University Press.
 90. Bathellier B, Poh Tee S, Hrovat C, Rumpel S. 2013 A multiplicative reinforcement learning model capturing learning dynamics and interindividual variability in mice. *Proc. Natl Acad. Sci. USA* **110**, 19 950–19 955. (doi:10.1073/pnas.1312125110)
 91. Collins AGE, Frank MJ. 2014 Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychol. Rev.* **121**, 337–366. (doi:10.1037/a0037015)