OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

Dopamine Function Across Species and Behaviours: A Comparative Neuroscience Review

Olivia Zhang
oliviazhang770330@gmail.com

ABSTRACT

Reinforcement learning is the process by which animals change their behaviour from experience. When an action leads to a better-than-expected outcome, the animal becomes more likely to repeat that action. Here, I review studies that use self-stimulation, electrophysiological recordings, and fibre photometry to show that dopamine contributes to reward processing and reinforcement learning across rats, monkeys, mice, songbirds, and humans. First, stimulation of dopamine-related brain regions in rats reinforces actions in a way similar to natural rewards. Second, electrophysiological recordings from single dopamine neurons in thirsty monkeys show that dopamine spiking activity increases when juice is delivered, but only when that juice reward is unexpected, providing a learning signal that can reinforce the action associated with the reward. In mice, a fibre photometry study shows that dopamine signalling is also activated when a mother retrieves her crying pup. In songbirds, recordings show that dopamine responds to performance errors during birdsong learning and changes during courtship depending on social context. In humans, dopamine neuron recordings show that firing changes in response to unexpected financial gains and losses during decision-making. Across the recording studies, dopamine signals were shaped by whether an outcome was better or worse than expected. Together, these studies show a conserved core pattern that dopamine activity increases when an outcome is better than expected and decreases when an outcome is worse than expected, indicating that dopamine acts as a flexible learning signal across species and behaviours, although reward prediction error does not explain all of dopamine's functions.

INTRODUCTION

Reinforcement learning (RL) is a process where organisms learn by associating their actions with the outcomes that follow through experience (Dayan & Niv, 2008). According to Thorndike's Law of Effect, if an animal performs an action and the outcome is better than it expected, the animal is then more likely to repeat that action again in the future. If the outcome is worse than expected, then the behaviour becomes less likely to happen again. Reinforcement learning applies across a wide range of behaviours, including foraging, in which animals learn which actions lead to food; performance learning, like when a bird improves its song; parental behaviour, where a mother learns to retrieve a crying pup; and courtship, where animals change their behaviour in response to social feedback from a potential mate.

Dopamine is important in understanding reinforcement learning because it signals the difference between the outcome the animal expects and the outcome it actually receives. This difference is called a reward prediction error (RPE) (Schultz et al., 1997). A positive prediction error occurs when the outcome is better than expected, such as receiving an unexpected reward. A negative prediction error occurs when the outcome is worse than predicted, such as expecting a reward but receiving none. When the outcome matches the expectation, there is neither surprise nor disappointment.

REVIEW METHODOLOGY

This narrative review was developed during 2026. The primary research papers were identified in collaboration with the research advisor rather than through a formal systematic database search. Relevant review articles on dopamine and reinforcement learning, including Dayan and Niv (2008), were also read to provide background and help identify central themes. Studies were included when they directly measured or manipulated dopamine-related activity during a defined behavioural task and contributed to comparisons across species, behaviours, or experimental methods. Priority was given to landmark studies representing rats, monkeys, mice, songbirds, and humans and using self-stimulation, electrophysiology, or fibre photometry. Because this is a narrative rather than systematic review, the literature included is selective and not exhaustive.

Table 1: Comparative Summary of the Reviewed Studies

Study	Species	Behavioural context	Method	Main finding
Olds & Milner (1954); Fouriezos & Wise (1976)	Rats	Lever pressing and intracranial self-stimulation	Electrical stimulation and dopamine receptor blockade	Activation of dopamine-related circuits reinforced lever pressing, and pimozide reduced self-stimulation.
Schultz et al. (1997)	Monkeys	Juice-reward learning	Single-neuron electrophysiology	Dopamine firing reflected positive, expected, and negative reward prediction errors.
Xie et al. (2023)	Mice	Maternal pup retrieval	Fibre photometry	Dopamine activity reflected expected and unexpected social outcomes

				involving pup contact.
Gadagkar et al. (2016)	Songbirds	Birdsong learning	Single-neuron electrophysiology	Dopamine neurons encoded vocal performance errors.
Roeser et al. (2023)	Songbirds	Courtship, thirst, and song	Electrophysiology and fibre photometry	Dopamine signals retuned according to social context and current behavioural priorities.
Zaghloul et al. (2009)	Humans	Financial decision-making	Single-neuron electrophysiology	Putative dopamine neurons responded to unexpected financial gains and losses.

Self-Stimulation Experiments Implicate Dopamine-Related Circuits in Reinforcement

One of the earliest studies showing dopamine's role in reinforcement came from Olds and Milner's self-stimulation experiments (Olds & Milner, 1954). The main question was whether specific regions of the brain, when activated, would reinforce an associated behaviour. In these experiments, rats had electrodes implanted into different parts of their brains. The rats were placed in a box where they could press a lever to stimulate a specific brain region. The behaviour in this experiment was lever pressing, and the outcome was electrical stimulation in the brain. The experiment aimed to determine which brain regions would reinforce behaviour and which would not. At this time in the 1950s, scientists did not yet know whether animals had reinforcement systems in their brains or, if so, which brain regions were involved. In this task, the behaviour was the rat pressing the lever, and the reinforcing outcome was the activation of dopamine-related brain regions through electrical stimulation. When the electrode was placed in brain areas such as the visual cortex, rats pressed the lever only once or twice per hour. However, when the electrode was placed in a region near midbrain dopamine neurons, the rats pressed the lever thousands of times per hour (Olds & Milner, 1954).

This showed that stimulation of those dopamine-related brain regions was reinforcing. The rats repeatedly pressed the lever even though they received no natural reward, such as food or water, because the brain stimulation itself served as a reinforcing outcome.

This experiment helped identify dopamine-related brain circuits within the reinforcement system. It showed that when a behaviour activates the dopamine system, the animal is more likely to repeat it. Subsequent studies strengthened the link to dopamine by showing that the dopamine receptor antagonist pimozide reduced intracranial self-stimulation, with response patterns that could not be explained only by motor or performance deficits (Fouriezos & Wise, 1976).

Neural Recording Experiments Test What Dopamine Neurons Encode During Behaviour

After self-stimulation studies linked dopamine-related circuits to reinforcement, the next major question was how naturally occurring dopamine neuron activity changed when reward was experimentally controlled. A major study that tested what dopamine neurons encode was conducted by Schultz, Dayan, and Montague (1997), "A Neural Substrate of Prediction and Reward." In this experiment, researchers used electrophysiological recordings to monitor dopamine neurons in thirsty monkeys seeking a juice reward (Schultz et al., 1997).

Electrodes were implanted into the monkey's midbrain, where dopamine neurons are located. The researchers recorded one dopamine neuron at a time and measured its spike activity, also called action potentials. The monkeys were thirsty and completed a Pavlovian conditioning task in which a conditioned stimulus predicted a juice reward one second later. In this experiment, the cue (CS) was the visual stimulus presented before the reward; the behaviour was the monkey observing the cue and anticipating the outcome; and the reward was the delivery of juice. The main questions were whether dopamine neurons responded to reward, whether they responded to reward-predicting cues, and whether the animal's expectations influenced the dopamine responses.

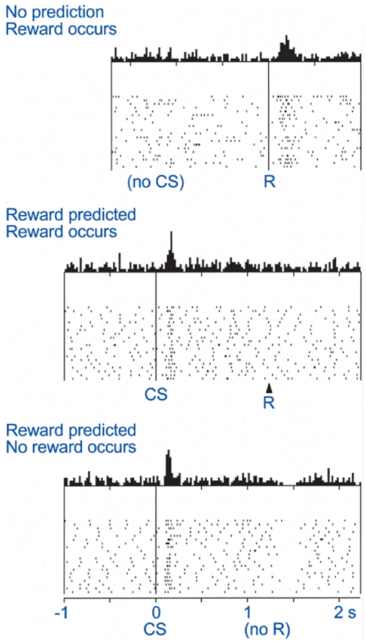

Figure 1. Dopamine neuron activity changes depending on whether the reward is expected or unexpected. *The x-axis represents time relative to the cue and reward delivery, while the y-axis represents dopamine neuron firing rate measured in spikes per second. Raster plots show individual spikes across trials. In each of the three panels, the top graph shows the average firing rate above the spike raster plot. In the first condition, there is no predictive cue, so the juice reward (R) is unexpected. Dopamine activity increases at the time of this unexpected reward delivery, shown as a peak in firing rate. This represents a positive prediction error. In the second condition (middle panel), a cue (CS) predicts the reward. The dopamine activation shifts to the time of the cue rather than the time of the reward. The peak at the cue indicates that the cue now carries predictive value, while the absence of a peak at reward delivery shows that the outcome is expected. In the third condition (lower panel), the cue is present, but the reward is unexpectedly omitted. Dopamine activity increases at the cue and then drops below baseline at the time when the reward was expected. This decrease in dopamine activity shows a negative prediction error following a disappointing reward outcome. Adapted from Schultz et al. (1997).*

Figure 1 shows the results of the recording experiments. Dopamine neurons responded strongly to rewards that were not predicted by the cue (Figure 1, top). When a conditioned stimulus (CS) predicted the reward, the dopamine neuron was activated by the CS but not by the delivery and consumption of the juice itself (Figure 1, middle). This showed that dopamine neurons are not simply activated by rewards when those rewards are predicted. Finally, if a CS predicted a reward but the reward was omitted, dopamine activity was suppressed at the moment of the disappointing outcome. Together, these results showed that dopamine encodes reward prediction error: activity increases when the monkey receives an unexpected juice reward, remains relatively unchanged when the juice reward is predicted, and decreases when the expected juice reward is not given (Schultz et al., 1997). This finding became highly influential in the neuroscience of learning and reward.

Measuring Dopamine Signals in the Brains of Awake, Behaving Animals: Fibre Photometry Versus Electrophysiology

Dopamine activity can be measured using different techniques, each providing distinct types of information. Electrophysiology dates back more than 70 years and involves inserting electrodes into the brain and recording electrical spikes from individual neurons. More recently, optical methods have been developed. Fibre photometry, for example, measures neural activity indirectly using light. In fibre photometry, a viral vector is used to make dopamine neurons express GCaMP, a fluorescent protein that becomes brighter when it binds calcium. When these neurons become active, calcium levels increase inside the cell, and the fluorescence becomes brighter. By placing an optical fibre over the dopamine neurons, scientists can observe dopamine-related activity through changes in fluorescence. Dopamine release can also be measured at release sites in the striatum, a region where dopamine axons release the neuromodulator. In this approach, genetically encoded dopamine sensors, such as dLight or GRAB-DA, become brighter when they bind dopamine. In photometry-based figures, the x-axis typically represents time aligned to behavioural events, while the y-axis represents normalized fluorescence signals. Changes in brightness correspond to increases or decreases in dopamine-related activity. However,

electrophysiology records millisecond-scale spikes from individual neurons, whereas fibre photometry measures a slower, combined signal from a population of neurons or from dopamine release within a brain region. Therefore, differences in the timing or size of signals across studies may reflect the methods used rather than biological differences between species.

Dopamine Encodes Reward Prediction Error During Maternal Pup Retrieval

Once dopamine signalling of reward prediction error was discovered in the 1990s in thirsty monkeys, a major question was whether the principle applied to other behavioural paradigms in which animals may also need to learn from reinforcing feedback. For example, most mammals exhibit parental behaviours. Even mouse mothers are highly motivated to find and comfort their crying pups if they wander from the nest. In Xie et al. (2023), dopamine activity during mouse maternal behaviour was measured using fibre photometry. Mouse mothers with several pups were observed in nests while their dopamine activity was recorded. Occasionally, experimentalists took a pup from the nest and moved it to a different part of the environment. When a pup is isolated from the nest, its temperature drops, and it begins to produce ultrasonic vocalizations that may be analogous to crying. The study first showed that when mothers sought out their pups, gently placed their mouths around them, and brought them back to the nest, there was a large dopamine activation when the mother made contact with the pup. This showed that dopamine activation can occur during parental behaviour. The next major question was whether this dopamine signal was influenced by what the mother predicted would happen; in other words, whether dopamine neurons encoded RPE in a parental context. To test this, mothers hearing the cries of an isolated pup also heard either a tone that predicted the pup would be retrievable (CS+) or a tone that predicted the pup would be non-retrievable because it was placed behind a door that could not be opened (CS-). After the tone, the mouse had the opportunity to seek out its pup (Figure 2; Xie et al., 2023).

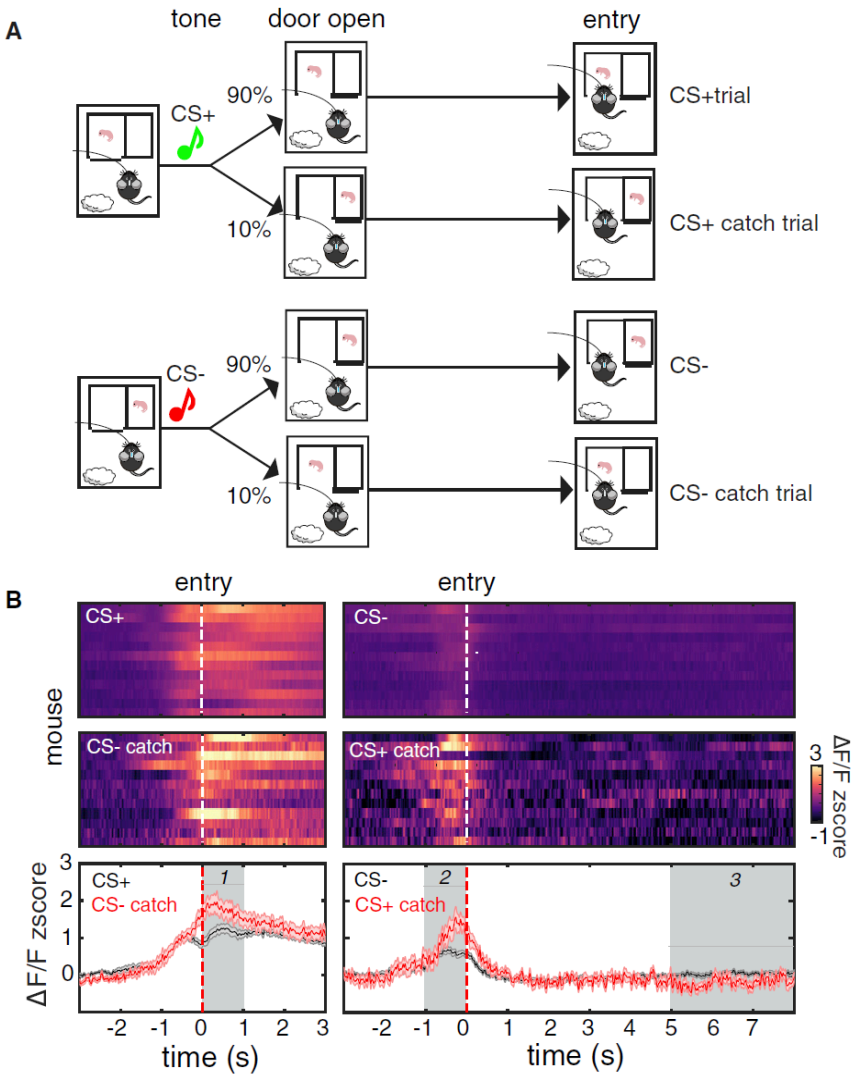

Figure 2. Dopamine encodes reward prediction error during parental behaviour in mice. (A) Experimental design showing four trial types. *In CS+ trials, a tone predicts that a pup will be retrievable. In CS- trials, a tone predicts that the pup will not be retrievable because it is placed behind a closed door. In CS+ catch trials, the tone predicts a pup, but no pup is present behind the open door, creating a worse-than-expected outcome. In CS- catch trials, the tone predicts no pup, but a pup is unexpectedly present, creating a better-than-expected outcome. (B) Dopamine activity aligned with the four trial types. The x-axis represents time-aligned events such as the tone, door opening, and chamber entry (time = 0). The y-axis represents dopamine activity, measured as a z-scored fluorescence signal. Heatmaps show activity across trials, with each row representing a single trial and colour intensity showing the level of fluorescence. Brighter colours indicate higher fluorescence, corresponding to increased dopamine-related activity. The line graph shows the average response across trials for each*

condition, allowing comparison of dopamine activity between trial types. The strongest signal occurs when the mouse unexpectedly encounters the pup on CS- catch trials, indicating a positive prediction error for a social reward. Reduced activity occurs when no pup is present on CS+ catch trials, reflecting a worse-than-expected outcome. Adapted from Xie et al. (2023).

The critical results were that dopamine activity increased following the CS+ tone and increased further when the mouse encountered the pup. The strongest signal occurred when the mouse unexpectedly encountered the pup on CS- catch trials, indicating a positive prediction error for a social reward. In contrast, dopamine signals at entry were weaker in the CS- condition. Importantly, activity was reduced when no pup was present on CS+ catch trials, reflecting a worse-than-expected outcome. Together, these data show that dopamine encodes both the prediction and the actual reward outcome in a social context.

Dopamine Encodes Reward Prediction Error During Birdsong Learning

These previous studies showed that when thirsty monkeys receive juice or when a caring parent makes contact with its pup, dopamine neurons are activated as if these events are rewarding. But what about behaviours, such as learning to talk or play an instrument, in which the quality of the outcome depends on how well performance matches internal goals? To address this question, Gadagkar et al. (2016) recorded dopamine activity in singing birds to test for reward prediction error-like signals related to vocal outcomes. First, the researchers used electrophysiology to measure spike activity while the birds sang. Next, the researchers distorted parts of the song so that the bird would treat them as errors. This allowed the researchers to control whether the outcome at a given part of the song was better or worse than expected (Figure 3).

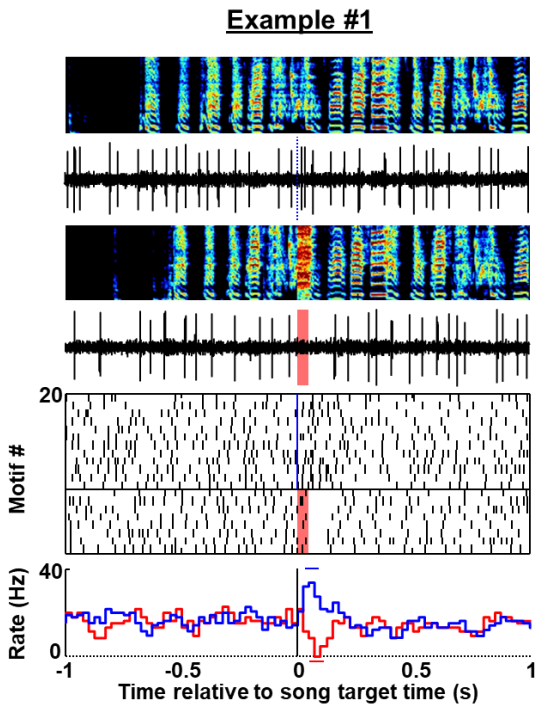

Figure 3. Dopamine neurons encode prediction error in singing birds. *Organized from top to bottom, the top panel shows a spectrogram of the undistorted rendition of the bird's song plotted above the spiking activity of a single dopamine neuron. Next is the spectrogram of a distorted rendition (red bar in the spectrogram). The middle panel shows raster plots, with each row representing a trial and each dot representing a spike (top, undistorted trials; bottom, distorted trials). Note the pause in spiking after the distortion and the burst in firing during the undistorted rendition. The bottom panel shows the averaged firing rate across undistorted (blue) and distorted (red) trials. Adapted from Gadagkar et al. (2016).*

As shown in Figure 3, dopamine activity increased when performance exceeded expectations and decreased when performance was erroneous. This shows that dopamine can encode performance error, not only external rewards (Gadagkar et al., 2016).

Dopamine Signals Dynamically Retune When an Animal Shifts Its Priorities

So far, the reviewed studies have examined conditions in which an animal has a relatively singular objective: obtaining water when thirsty, finding a pup when parenting, or producing the correct note when singing. However, it remained unknown how dopamine neurons would respond when an animal's objectives shifted. For example, would dopamine signals remain fixed, or would they change depending on behavioural priorities such as social interaction during courtship? In earlier experiments, dopamine encoded reward prediction error in tasks involving water, pups, or performance. During courtship, however, animals may prioritize social interactions over other rewards. This raises the question of

whether dopamine neurons continue to encode reward prediction error in the same way or whether these signals are retuned by social context (Roeser et al., 2023).

To test this, researchers measured dopamine activity in water-deprived male songbirds across four conditions: alone or with a female, and singing or not singing.

To measure dopamine activity, the researchers used fibre photometry, which involved injecting a viral vector into midbrain dopamine neurons so that they expressed GCaMP. This fluorescent protein becomes brighter when calcium levels rise. Because calcium levels increase when neurons are active, changes in fluorescence were used as a measure of dopamine-related activity. An optical fibre was implanted in the brain to record these fluorescence signals.

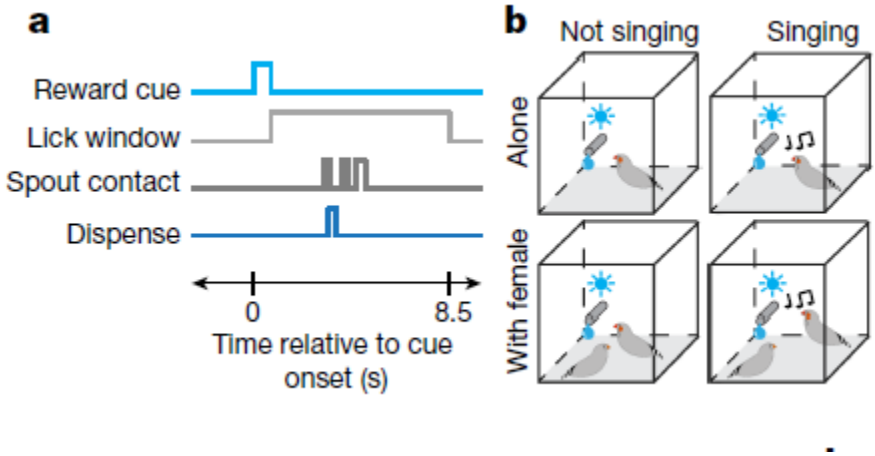

Figure 4. Experimental design used to test whether courtship changes the expression of thirst. In Figure A, the task design is shown. *The timeline includes a reward cue, a licking window, spout contact, and reward delivery. This part of the experiment is used as a standard reward task, similar to earlier dopamine studies, to compare how dopamine responds to predicted rewards.* ***In Figure B, the different behavioural conditions are shown.*** *The bird can be either alone or with a female, and in each case it can be either not singing or actively singing. This creates four conditions: alone and not singing, alone and singing, with a female and not singing, and with a female and singing. This allows the researchers to test how dopamine activity changes across social contexts and behaviours. Adapted from Roeser et al. (2023).*

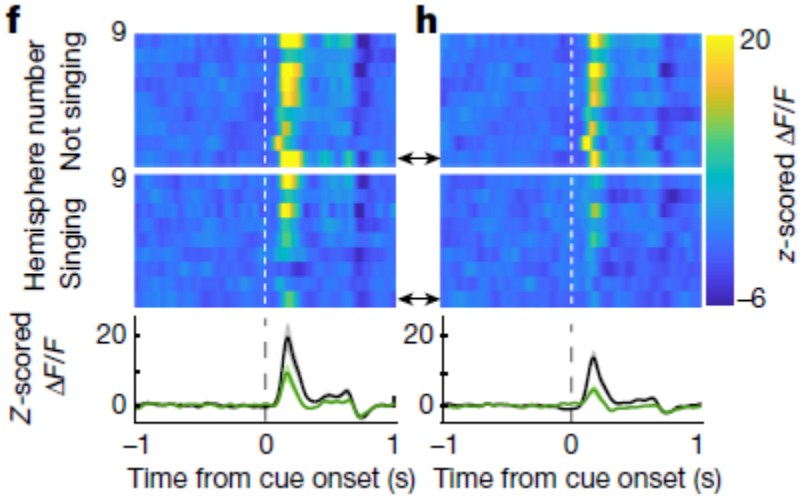

Figure 5. Dopamine activations retune according to social context. In Figure F, heatmaps of dopamine activity are shown for the different behavioural conditions. *The x-axis represents time aligned to the cue at time = 0, while the y-axis represents individual trials. Each row represents activity during one trial, and colour intensity indicates the level of fluorescence. Brighter colours correspond to higher fluorescence signals, reflecting increased dopamine-related activity. The heatmaps compare dopamine responses when the bird is not singing with responses when it is singing. After the cue is presented, the heatmaps become brighter, indicating an increase in fluorescence and dopamine-related activity.* ***In Figure H, the average dopamine signal over time is shown below the heatmaps.*** *When the bird is alone and not singing, dopamine shows a typical increase after the cue, consistent with a reward prediction signal. When the bird is with a female but not singing, the responses remain present but are reduced in intensity. The greatest change occurs when the bird is singing in the presence of a female. In this condition, the dopamine response has a lower peak amplitude than when the female is absent. This indicates that water-related dopamine responses are modulated by the social environment and the bird's current behavioural priorities (Roeser et al., 2023). Adapted from Roeser et al. (2023).*

Across conditions, dopamine responses were strongest when the bird was alone and were reduced when a female was present, especially during singing. The reduced peak height indicated that social context modulates the strength with which water-related cues and outcomes are represented. This suggests that dopamine signals are dynamically retuned according to the animal's current priorities rather than responding to reward in a fixed way (Roeser et al., 2023).

Human Dopamine Neurons Also Encode Reward Prediction Error

Animal studies show that dopamine prediction-error signals appear across many behaviours, including juice-reward learning in monkeys, pup retrieval in mice, birdsong performance, and courtship in zebra finches. However, to argue that dopamine functions as a general learning signal, it is important to ask whether the same pattern also appears in humans. To test this, researchers recorded from putative

dopamine neurons in awake, conscious humans engaged in a task in which reward was experimentally controlled (Zaghloul et al., 2009).

Zaghloul et al. (2009) examined how human substantia nigra neurons respond to reward prediction error. They used electrophysiological recordings to measure neural activity in patients with Parkinson's disease undergoing deep-brain stimulation surgery. During surgery, microelectrodes were inserted into the brain, allowing the researchers to measure the firing of putative dopamine neurons while the patients were awake (Zaghloul et al., 2009).

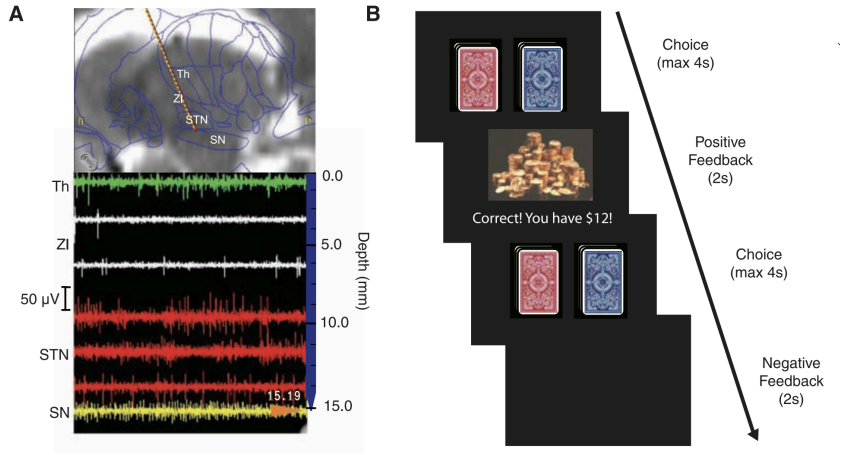

Figure 6. Recordings from putative dopamine neurons in awake, behaving humans. *The recording locations are shown across different brain regions. The coloured traces below represent neural activity recorded at different depths, and the researchers identified putative dopamine neurons based on their firing patterns.* ***In Figure B, the experimental task is shown. Participants were asked to choose between two card decks on each trial.*** *Each deck had a different probability of reward, so over time the participants had to learn which deck was better to maximize their earnings. After each choice, they received either positive or negative feedback, which could be expected or unexpected. Adapted from Zaghloul et al. (2009).*

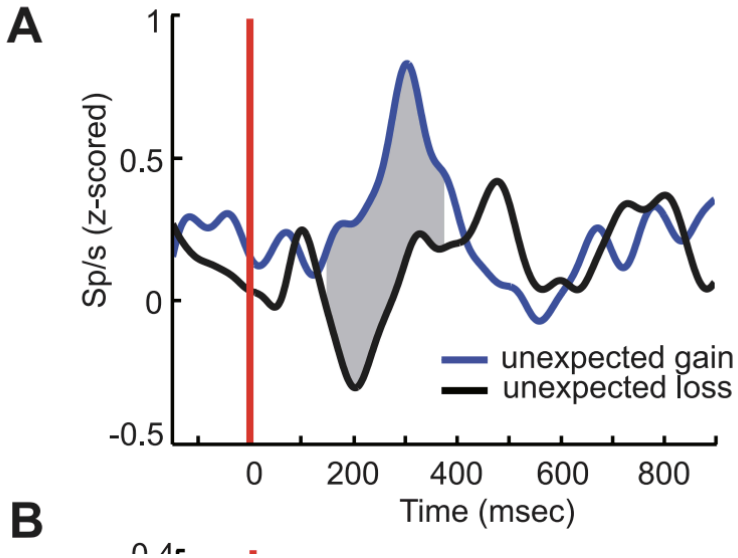

Figure 7. Human dopamine neurons encode reward prediction error. *Dopamine activity is plotted over time, with the x-axis showing time in milliseconds relative to feedback at time 0 and the y-axis showing the z-scored firing rate in spikes per second. The traces compare responses to unexpected gains and unexpected losses. After feedback is given, dopamine activity increases for unexpected gains and decreases for unexpected losses. Adapted from Zaghloul et al. (2009).*

This experiment shows that putative dopamine neurons in humans respond to unexpected outcomes. When an outcome is better than expected, firing increases; when it is worse than expected, firing decreases. When the outcome is expected, there is little or no change in activity (Zaghloul et al., 2009).

DISCUSSION

This review asked whether dopamine functions as a general prediction-error signal across species and behavioural contexts. Across the recording experiments, dopamine activity consistently reflected the relationship between expected and actual outcomes. In each recording study, a behavioural event was linked to a clearly defined outcome, and dopamine activity changed depending on whether that outcome was better or worse than expected.

In the self-stimulation experiment, lever pressing led to electrical activation of dopamine-related brain regions, which functioned as a reinforcing outcome (Olds & Milner, 1954). In the monkey experiment, a cue predicting juice shifted dopamine activity from the reward to the predictive cue (Schultz et al., 1997). In the maternal behaviour study, entering the chamber resulted in encountering a pup, which served as a social reward (Xie et al., 2023). In birdsong, singing led to an internally evaluated outcome based on how

closely the song matched a stored template (Gadagkar et al., 2016). In courtship, dopamine responses were influenced by social context and female feedback (Roeser et al., 2023). In humans, choosing a card deck resulted in monetary gain or loss (Zaghloul et al., 2009).

Despite these differences, the recording studies show the same core pattern: dopamine activity increases when a specific outcome is better than expected, decreases when it is worse than expected, and remains relatively stable when it is predicted. This consistency supports the idea that dopamine encodes reward prediction error across species and behavioural contexts.

The strongest similarity across studies is that dopamine signals changed according to the relationship between expected and actual outcomes despite major differences in species and reward type. However, direct comparisons should be made cautiously because self-stimulation manipulates broad neural circuits, electrophysiology records single-neuron spikes, and fibre photometry measures population-level fluorescence. Apparent differences among studies may therefore reflect temporal resolution, anatomical targeting, and neuronal sampling, as well as biological differences.

Dopamine signalling is not limited to primary rewards such as food or water. The comparison across experiments shows that dopamine also responds to social rewards, internal performance outcomes, and monetary rewards. In the maternal behaviour study, the relevant reward is contact with the pup, demonstrating that dopamine responds to socially meaningful outcomes (Xie et al., 2023). In the birdsong study, the relevant outcome is not an external reward but is based on internal performance accuracy, showing that dopamine can encode internally generated expectations (Gadagkar et al., 2016).

The courtship experiment further demonstrates that dopamine signals are not fixed but are modulated by context. When a female is present, water-related cue responses are reduced, showing that social context changes the strength and behavioural relevance of dopamine signals. This suggests that dopamine reflects not only reward prediction but also context, expectation, and current behavioural priorities (Roeser et al., 2023).

Another key point is that dopamine responses can vary with the probability and magnitude of a prediction error. In the birdsong experiment, the size of the error response depended on distortion probability, indicating stronger signals when outcomes were more surprising (Gadagkar et al., 2016). Schultz et al. (1997) similarly showed that dopamine firing distinguishes better-than-expected, expected, and worse-than-expected outcomes. These signals can help organisms update their behaviour more effectively when outcomes violate expectations.

Together, these findings challenge the common assumption that dopamine primarily represents pleasure or positive emotion. Instead, the evidence shows that dopamine functions as a teaching signal that supports learning by encoding differences between expected and actual outcomes.

However, dopamine does not exclusively encode reward prediction error. Different dopamine neuron populations can also encode movement, sensory and cognitive variables, aversive events, and contextual or motivational information (Engelhard et al., 2019; Menegas et al., 2018; Roeser et al., 2023). Therefore, the studies reviewed here support reward prediction error as an important and broadly conserved dopamine function, but not as a complete explanation of all dopamine signalling.

Finally, this broader view helps explain how dopamine may be involved in complex conditions such as addiction and Parkinson's disease. Drugs of abuse can produce unusually strong dopamine signals that reinforce drug-seeking behaviour, although addiction also involves additional neural and psychological mechanisms (Dayan & Niv, 2008). In Parkinson's disease, the loss of dopamine neurons contributes to movement difficulties and may also disrupt reinforcement learning.

LIMITATIONS

This review focuses on a relatively small number of landmark studies selected to represent different species, behaviours, and methods. Because the studies were identified collaboratively for a narrative review rather than through a systematic database search, selection bias is possible, and the review does not capture the full diversity of dopamine neurons or all findings in the field. In addition, the studies used different techniques, brain regions, behavioural tasks, and definitions of reward, which limits direct comparison. Finally, findings from animal models cannot be directly translated to humans, and the human study involved patients with Parkinson's disease undergoing surgery, which may limit generalization to healthy populations.

FUTURE DIRECTIONS

One possible future experiment could test how dopamine integrates internal performance errors and external reward prediction errors within the same task. Previous studies examined these prediction errors separately: birdsong studies focused on internal evaluation, while monkey and maternal behaviour studies focused on external rewards (Schultz et al., 1997; Gadagkar et al., 2016; Xie et al., 2023). However, it remains unclear how dopamine neurons respond when both types of prediction error occur simultaneously.

To test this, researchers could use singing birds and combine distorted auditory feedback with an external reward, such as food. In this experiment, the behaviour would be the bird singing, while the outcome would include both the quality of the song performance and the delivery of an external reward. Researchers could distort certain syllables during singing while unpredictably delivering or omitting food rewards. Dopamine neuron activity could then be recorded using electrophysiology or fibre photometry.

The hypothesis is that dopamine neurons integrate internal and external outcomes into a single prediction-error signal rather than encoding them independently. If this is true, dopamine activity should

change not only when song performance is distorted but also when external rewards are unexpectedly delivered or omitted. This experiment would help determine whether dopamine functions as a general learning signal across multiple forms of reinforcement learning.

CONCLUSION

Overall, the reviewed evidence supports reward prediction error as an important and flexible dopamine learning signal across reward learning, skill learning, social interaction, and decision-making. However, differences across recording methods, motivational states, and behavioural contexts, together with the heterogeneity of dopamine neuron populations, mean that reward prediction error should be understood as a major dopamine function rather than a complete account of all dopamine activity.

REFERENCES

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. Current Opinion in Neurobiology, 18(2), 185–196.

Engelhard, B., Finkelstein, J., Cox, J., Fleming, W., Jang, H. J., Ornelas, S., Koay, S. A., Thiberge, S. Y., Daw, N. D., Tank, D. W., & Witten, I. B. (2019). Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. Nature, 570(7762), 509–513.

Fouriezos, G., & Wise, R. A. (1976). Pimozide-induced extinction of intracranial self-stimulation: Response patterns rule out motor or performance deficits. Brain Research, 103(2), 377–380.

Gadagkar, V., Puzerey, P. A., Chen, R., Baird-Daniel, E., Farhang, A. R., & Goldberg, J. H. (2016). Dopamine neurons encode performance error in singing birds. Science, 354(6317), 1278–1282.

Menegas, W., Akiti, K., Amo, R., Uchida, N., & Watabe-Uchida, M. (2018). Dopamine neurons projecting to the posterior striatum reinforce avoidance of threatening stimuli. Nature Neuroscience, 21(10), 1421–1430.

Olds, J., & Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. Journal of Comparative and Physiological Psychology, 47(6), 419–427.

Roeser, A., Gadagkar, V., Das, A., Puzerey, P. A., Kardon, B., & Goldberg, J. H. (2023). Dopaminergic error signals retune to social feedback during courtship. Nature, 623(7986), 375–380.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. Science, 275(5306), 1593–1599.

Xie, Y., Huang, L., Corona, A., Pagliaro, A. H., & Shea, S. D. (2023). A dopaminergic reward prediction error signal shapes maternal behavior in mice. Neuron, 111(4), 557–570.e7.

Zaghloul, K. A., Blanco, J. A., Weidemann, C. T., McGill, K., Jaggi, J. L., Baltuch, G. H., & Kahana, M. J. (2009). Human substantia nigra neurons encode unexpected financial rewards. Science, 323(5920), 1496–1499.