Cognition 244 (2024) 105714

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

The spread of affective and semantic valence representations across states

Orit Heimer [a,*], Uri Hertz [b]

[a] Department of Psychology, University of Haifa, Haifa, Israel
[b] Department of Cognitive Sciences, University of Haifa, Haifa, Israel

A R T I C L E I N F O

Keywords:
Affective valence
Semantic valence
States
Associative learning
Computational models

A B S T R A C T

In many decision problems, outcomes are not reached after a single action but rather after a series of events or states. To optimize decisions over multiple states, representations of how good or bad the outcomes are, that is, the outcomes' valence, should spread across states. One mechanism for valence spreading is a temporal, state-independent process in which a single valence representation is updated when an outcome is experienced and fades away afterwards. Each state's valence is based on its temporal proximity to the experienced outcome. An alternative, state-dependent mechanism relies on the structure of transitions between states, updating a separate valence representation for each state according to its spatial distance from the outcomes. We examined how these mechanistic accounts shape the spread of two formats of valence representation, feelings (affective valence) and knowledge (semantic valence), between states.

In two pre-registered experiments ($N = 585$), we used a novel task in which participants move in a four-state maze, one of which contains an outcome. The participants provide self-reports of affective and semantic valence throughout the maze and after finishing it. Results show that the affective representation of negative valence is more localized in state-space than the semantic representation. We also found evidence for the relative reliance of the affective valence on a temporal, state-independent mechanism and of the semantic valence on a structured, state-dependent mechanism.

Our findings provide mechanistic accounts for the differences between affective and semantic valence representations and indicate how such representations may play a role in associative learning and decision-making.

1. Introduction

Imagine you commute to work by train. The train passes through four stations; one of them is under construction. The train's delay at this station is longer than in other stations, and it is very noisy (see Fig. 1a). Therefore, if you were asked about the *valence* of this station, that is, how good or bad, pleasant or unpleasant this station is (e.g., Barrett, 2006), you would probably say it has a negative valence. In this study, we are interested in the way the negativity of the station under construction may influence your judgment of the valence of the other stations, that is, the way the negativity spreads to other stations, the mechanisms governing such a spread, and its dependence on the *format* in which the negativity is represented, that is, feelings vs. knowledge. In the following, we will elaborate on the above research questions, starting with whether the valence spreads from station to station.

1.1. Does valence spread?

Our primary question is how and whether the negativity of the station under construction influences the valence of the other stations, that is, whether the negative valence of the station under construction spreads to the other stations' valence or stays localized. Notably, reinforcement learning (RL) models (e.g., Sutton & Barto, 2018) provide an initial positive answer to this question. RL models are computational models that describe how an agent learns a behavioral action policy based on trial and error. The agent learns action–outcome association, that is, which action will most likely lead to the best reward. Using RL modeling allows for generating computationally explicit and directly testable hypotheses about the characteristics of the behavior under investigation (Niv, 2009). In the RL literature, the train stations are termed "*States*." States are defined as discrete episodes in which an agent (e.g., a computer algorithm, an animal, or a human) can take some action. In some states, the agent also experiences an outcome (Gläscher, Daw, Dayan, & O'Doherty, 2010). In our train example, the agent is a

* Corresponding author.
E-mail address: heimer.orit@gmail.com (O. Heimer).

https://doi.org/10.1016/j.cognition.2023.105714
Received 10 June 2023; Received in revised form 22 December 2023; Accepted 24 December 2023
Available online 3 January 2024

train passenger. In each state (i.e., station), the passenger can decide to leave the train. In the under-construction state, the passenger experiences an adverse outcome resulting from the longer delay and the noise.

Significantly, RL models deal with cases where current decisions impact future outcomes, that is, where sequential-state planning is needed (Dolan & Dayan, 2013). In these cases, one needs to make a series of actions that lead him from state to state, and the outcome is revealed only in the final state. Two essential points are implied. First, implementing sequential-state planning computations requires knowing the *transition probabilities* between states, that is, the probability of moving from the current state to each other state. For example, if you are currently at Station 1 and stay on the train, there is a 100% chance you will move to Station 2, a 0% chance you will move to Station 3, and a 0% chance you will move to Station 4. In our example, the transition probabilities remain *fixed,* that is, the train follows the same sequence of stations every day. However, in other cases, the transition probabilities between states may vary. Second, the agent needs to evaluate the long-term consequences of his decision, that is, take into account both transition probabilities and final outcome. Unlike our simple train example, in most cases, evaluating all the possible long-run sequences of actions is very computationally demanding or even impossible.

One way to solve this taxing computational planning problem is to represent each state's valance not only in terms of its direct outcome but also in terms of its likelihood to lead to future outcomes, using a process called dynamic programming (e.g., O'Donoghue, Osband, Munos, & Mnih, 2018). In dynamic programming, once an outcome is experienced in one state, the valance of all possible preceding states is also updated according to their likelihood to lead to this final state. Moreover, the valence updating process does not stop at the immediate preceding states. Instead, it continues recursively to states further away from the final state, as the preceding states of each of the preceding states are also updated, and so on. Notably, this valence update mechanism frees the agent from the need to consider all future states in advance. As the valance of the current state already includes the valance of all future states, it is sufficient to decide based on the valance of the immediate following states. Thus, a core assumption of RL models is that valance spreads across states to guide behavior.

Numerous studies show that during associative learning tasks, humans and other animals form representations of the task environment, that is, the different states within the task and the expected transition probabilities between them, and update the states' valance following dynamic programming principals (e.g., Gershman, Markman, & Otto, 2014; Keramati, Smittenaar, Dolan, & Dayan, 2016; Prével & Krebs, 2021; Rmus, Ritz, Hunter, Bornstein, & Shenhav, 2022). Specifically, an influential multi-stage decision task showed direct evidence for the spread of valence from a rewarded state to the preceding state (Daw, Gershman, Seymour, Dayan, & Dolan, 2011). In the first stage, participants had to decide between two options, leading them to one out of two rooms, where another decision was made, and a reward was experienced. The magnitude of rewards experienced in the second stage

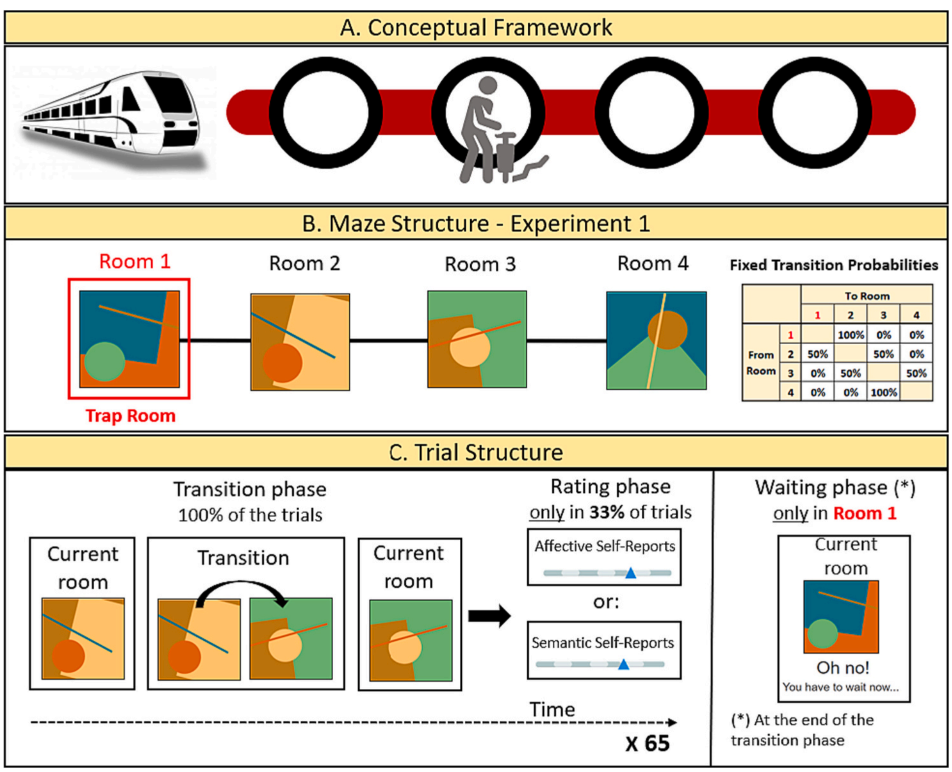

Fig. 1. (a) Conceptual Framework. A train passes through four stations. One of the stations is under construction. The train's delay at this station is longer than at other stations, and it is very noisy. We are interested in how the negativity of the station under construction may spread to other stations, the mechanisms governing such a spread, and its dependence on the format in which the negativity is represented, that is, feelings vs. knowledge. (b) Experiment 1's design – Maze Structure. The participants were led through a series of four rooms representing four states, each represented by a unique abstract image. The task included three identical mazes. In each, the participants visited 65 rooms. A unique abstract image represented each room in each maze. The series had fixed transition probabilities between the rooms. Specifically, the extremist rooms, Rooms 1 and 4, lead to the interior rooms, that is, Rooms 2 and 3, respectively, with a 100% probability. Room 2 leads to Room 1 with a 50% probability and Room 3 with a 50% probability. Room 3 leads to Room 2 with a 50% probability and Room 4 with a 50% probability. The leftmost room, Room 1, was a "trap" room. Usually, the participants self-timed their stay in a specific room. However, when visiting Room 1, the "trap" room, the participants had to wait at least 3 s before moving on. (c) Trial Structure. After visiting a specific room, the participants saw the room they left and the next room on the screen together. Notably, the participants were not explicitly informed of the series' structure, but they could use the transitions between the rooms to infer it. The leftmost room, Room 1, was a "trap" room. At the end of 33% of the trials, the participants were asked to rate either the affective or the semantic valence of the room they just visited.

changed over time, and the participants had to adapt their decisions accordingly. Crucially, the participants not only updated their decision behavior in the second stage, where the reward was actually experienced, but also in the first stage. We can conclude that the participants ensured being led to the more rewarded room using the spread of valence from the rewarded state to the preceding state.

The behavioral evidence received additional support in a series of brain studies that demonstrated the existence of specific regions involved in representations of the task environment and goal-directed control over behavior (e.g., Garvert, Dolan, & Behrens, 2017; Pauli et al., 2015; Pauli, Gentile, Collette, Tyszka, & O'Doherty, 2019; Seymour et al., 2004; for review, see Behrens et al., 2018; Bellmund, Gärdenfors, Moser, & Doeller, 2018; Epstein, Patai, Julian, & Spiers, 2017). For example, Klein-Flügge, Wittmann, Shpektor, Jensen, and Rushworth (2019) showed that representations of task knowledge are derived via multiple learning processes of state structure, operating at different time scales, associated with partially overlapping and partially specialized anatomical regions. Crucially, they also showed evidence for the spatial spread of reward between states by demonstrating a change in the amygdala activity to stimuli close in space to the location where the reward had occurred.

To sum up, the behavioral and brain evidence suggest that goal-directed RL algorithms, assuming valence spread, are not only efficient solutions to sequential-state planning but also used by humans. Still, the exact mechanism for valence spread is unclear. It is also unclear whether valence spread depends on the representation's format, that is, feelings versus knowledge. We will now introduce the two possible accounts for valence spread, that is, time versus structure, and then turn to the two formats of valence representation.

1.2. Time versus structure: two accounts for valence spread

What accounts may explain the spread of valence from one station to the other, for example, from Station 2 to Station 3? Will the valence spread from Station 2 to Station 3 because these two stations are close in **space,** that is, the train travels from Station 2 to Station 3, or because they are close in **time**, that is, in the train's schedule, the time it reaches Station 3 is close to the time it reaches Station 2?

1.2.1. Spread-through-time

According to behaviorists such as Pavlov (2010), animals learn to associate specific stimuli and outcomes, that is, Pavlovian conditioning, or learn to repeat specific stimuli–response patterns that lead to positive reinforcement, that is, instrumental conditioning (Skinner, 1938). This account is based on the history of experiences, that is, the spread of the negative valence of the station under construction to the other stations' valence will depend on the time that passed from the negative (/positive) experience of the under-construction station until visiting the other stations. Shorter times indicate a shorter temporal distance to the outcome and a greater spread of the outcome's valence. In our example, this account predicts the spread of negative valence from Station 2 (the under-construction station) to Station 3, which is in temporal proximity to Station 2, and less to Station 4, which is further in time from Station 2. It also predicts that the negative valance will be reduced when moving back from Station 4 to Station 3 (on the way back home), as these transitions are even further in time from the negative experience. More recent research used the terms "habitual," "model-free," and "retrospective" for this account.

1.2.2. Spread-through-structure

An alternative, reflective account, suggested by Tolman (1948), is to plan the behavior to achieve the organism's goals using an internal representation of the environment, that is, the different states (or contexts) in it and their connectivity, in the form of a cognitive map (for reviews, see Behrens et al., 2018; Bellmund et al., 2018; Epstein et al., 2017). According to this state-dependent account, the spread of the under-construction station's valence to other stations' valence will depend on the structure of the states' space, that is, the transition probabilities between them or their structural distance. Shorter paths between the outcome and the current state, that is, fewer steps needed to move from Station 2 to the current station, indicate a shorter structural distance to the outcome and, therefore, a greater spread of Station 2's negativity. Importantly, this does not depend on the direction of movement, to and from Station 2. More recent research has used the terms "goal-directed," "model-based," and "perspective" for this account.

1.3. Two formats for valence spread: semantic and affective representations

The last extension of our primary valence spread question is whether the spread of valence depends on the format in which the negativity is represented, that is, feelings vs. knowledge. Previous research supports two types of representations of a stimulus' valence: semantic and affective (e.g., Givon, Itzhak-Raz, Karmon-Presser, Danieli, & Meiran, 2019; Heimer, Kron, & Hertz, 2023, Itkes, Kimchi, Haj-Ali, Shapiro, & Kron, 2017; Robinson and Clore, 2002; Wang et al., 2021. See Itkes & Kron, 2019 for review). We assume that affective and semantic valence are different mental representations characterized by specific **formats**, that is, feelings and knowledge (Cardinal, Parkinson, Hall, & Everitt, 2002; Gazzaniga, Ivry, & Mangun, 2014). Affective valence is the conscious experience of pleasure or displeasure that follows exposure to an event, like the frustration felt when the train is delayed in the under-construction station, which indicates that passing through this station is an adverse event.[1] Semantic valence is a general, conceptual, non-experiential stored knowledge about the valence of events, for example, "delays are a bad thing" (e.g., Osgood, 1952; Russell, 1983).[2]

Do the different **formats** of valence representation differ in their dependency on the two accounts, that is, time and structure, and in their spread? Previous research suggests that the affective valence representation is more temporally local than the semantic valence representation. First, studies that utilized a habituation protocol, that is, repeated exposure to a stimulus (e.g., Itkes et al., 2017; Wang et al., 2021) showed that measures related to affective valence are attenuated with repeated exposure, whereas measures related to semantic valence do not attenuate or attenuate to a lesser extent. The logic behind this finding is that repeated exposure to the same stimulus causes a reduction in this stimulus' novelty and unpredictability. Therefore, it reduces the intensity of the emotional response to the stimulus, that is, its affective valence (Bradley, Lang, & Cuthbert, 1993; Öhman, Hamm, & Hugdahl, 2000). In contrast, the semantic valence is expected to remain relatively stable because semantic knowledge represents meaning retrieved from long-term memory (see Itkes et al., 2017). Heimer et al. (2023) used a neutral conditioned stimulus (CS) in the context of reversal learning, a type of associative learning (Izquierdo, Brigman, Radke, Rudebeck, & Holmes, 2017), to check whether affective and semantic valence representations follow different learning dynamics. They showed that when the outcomes are variable, affective valence representations are updated more rapidly, indicating a more temporally local effect of the outcome

[1] The negative feeling is accompanied by a profile of changes such as autonomic (e.g., Bradley et al., 1993; Hodes, Cook III, & Lang, 1985), motoric (e.g., Dael, Mortillaro, & Scherer, 2012; Du, Tao, & Martinez, 2014), and action tendencies (e.g., Carver, 2006).

[2] Semantic valence includes general factual knowledge (e.g., Delayes are bad), episodic knowledge about a specific event at a particular time and place (e.g., I was delayed yesterday and felt bad), and self-related knowledge that combines episodic and semantic knowledge to form information about the self (e.g., I hate delays). Notably, the critical unity between these forms of knowledge about the valence of the object is that they are non-experiential (they do not represent an experience).

on representation.

2. Predictions and simulations of the time and structure accounts

We used a novel associative learning task to study how structure and temporal dynamics shape the way valence representations spread across states. During the task, the participants were led through a series of four rooms representing four states, each represented by a unique abstract image (Fig. 1b). The task included three identical mazes. In each, the participants visited 65 rooms. A unique abstract image represented each room in each maze. The series had fixed transition probabilities between the rooms. Specifically, the extremist rooms, Rooms 1 and 4, lead to the interior rooms, that is, Rooms 2 and 3, respectively, with a 100% probability. Room 2 leads to Room 1 with a 50% probability and Room 3 with a 50% probability. Room 3 leads to Room 2 with a 50% probability and Room 4 with a 50% probability.

After visiting a specific room, the participants saw the room they left and the next room on the screen together. Notably, the participants were not explicitly informed of the series' structure, but they could use the transitions between the rooms to infer it (Fig. 1c). The leftmost room, Room 1, was a "trap" room. Usually, the participants self-timed their stay in a specific room. However, when visiting Room 1, the "trap" room, the participants had to wait at least 3 s before moving on. At the end of 33% of the trials, the participants were asked to rate either the affective **or** the semantic valence of the room they just visited.[3] At the end of the maze, the participants were asked to rate the affective and semantic valence of all four rooms. The mid and post-maze valence ratings were performed using affective and semantic self-reports (see Section 3.1.2).

2.1. Computational models of the two accounts

2.1.1. The state-independent model

Our state-independent model of the spread-through-time account is inspired by the computational model of momentary subjective well-being suggested by Rutledge, Skandali, Dayan, and Dolan (2014; see also Eldar, Rutledge, Dolan, & Niv, 2016; Keren et al., 2021). They used a lottery task, where the participants could earn or lose varying amounts of money in each trial. Their model showed that the participant's momentary happiness was influenced by the gap between their expectations and the actual outcome, that is, the prediction errors arising from those expectations. Similarly, our state-independent model assumes the participants form a single continuous valence representation, which relies on their expectations of what will happen next and the difference between these expectations and their actual experiences. When asked to report a specific room (/state) valence, the participants report on this continuous representation, which does not directly consider which room is currently being visited, that is, is state-independent. Specifically, we expect that a visit to the trap room will cause disappointment, that is, a negative prediction error, and therefore negatively influence their valence representation. On the other hand, visiting other rooms will cause a positive prediction error that will positively influence the valence representation (Fig. 2b, orange line).

According to our state-independent model, a model-free reinforcement learning model, the participants update a single, continuous state-valence representation while moving from one room to the other. This single representation reflects their emotional state, and is not dependent on the room they are visiting. The update is based on the difference between the participant's expectations and actual experiences, that is, prediction errors (Eq. (1)):

$$\delta_t = r_t - Q(S)_t \tag{1}$$

δ_t – Prediction error in trial t.
r_t – Outcome in trial t.
$Q(S)_t$ – Predicted valence of the states in trial t.

Furthermore, the update is done according to a learning rule (Eq. (2)):

$$Q(S)_{t+1} = Q(S)_t + \alpha \cdot \delta_t \tag{2}$$

α – Learning rate; $0 < \alpha < 1$.

The learning rate controls the weight given to the history of the long-term outcomes versus the current outcome and the resulting prediction error. The higher the learning rate, the more weight is given to the current outcome and the resulting prediction error, that is, the update of the continuous valence, occurs faster (Niv, 2009). Therefore, the higher the learning rate, the faster the negative experience of visiting the trap room fades away after experiencing positive outcomes. On the other hand, the lower the learning rate, the more time it takes for the negative influence of the trap room on the continuous valence to fade away. Therefore, a low learning rate leads to a higher spread rate of the negativity of the trap room between the states. In contrast, a high learning rate leads to a low spread rate, that is, localization, of this negativity (Fig. 2a).

2.1.2. The state-dependent model

Our state-dependent model of the spread-through-structure account is based on the successor representation (SR) model (Dayan, 1993; for reviews, see Momennejad, 2020; Momennejad et al., 2017; Russek, Momennejad, Botvinick, Gershman, & Daw, 2017). The SR model, a variant of model-based reinforcement learning models, uses the state space structure; when an outcome is experienced in one state, all state representations are updated according to this structure. In this model states represent the different rooms. While states proximate to the current state change a lot, those further away do not change that much. Notably, in this model, the valence representation of a specific room, for example, Room 2, is based not only on the outcomes received in Room 2 itself but also on the fact that Room 2 leads to Room 1 with 50% probability and to Room 3 with 50% probability, and each of these possible successive rooms is connected to a unique outcome. According to this account, when asked to report a specific room (/state) valence, the participants report on the unique representation of this specific state and not on a single, continuous representation as assumed by the spread-through-time account (Fig. 2b, green line).

To enable this, the model assumes that the participants learn the structure of the state space, that is, the transition relations between the states, using a connectivity rate parameter. The connectivity rate controls the furthest successor state that is "visible" from every starting state. As it gets larger, further states are taken into account and the spread of valence between the states increases. Lower connectivity will lead to a lower spread rate of the outcome's valence. Full details of the SR model can be found in the Supplementary material.

2.2. Model simulations

We used the state-dependent and state-independent models detailed above and the actual history of room transitions experienced by the 262 participants of Experiment 1 to stimulate the mid-maze valence self-reports. Full details of the model simulations are available in the Supplementary material.

2.2.1. Rate of valence spread

In our task, the structural distance from the outcome room is positively correlated with the time that passed since the outcome room was visited, that is, the temporal distance (Fig. 2b). Therefore, both the time and structure accounts should predict that mid-maze valence self-

[3] These ratings were the same for the "trap" room, i.e., Room 1 and all other three rooms.

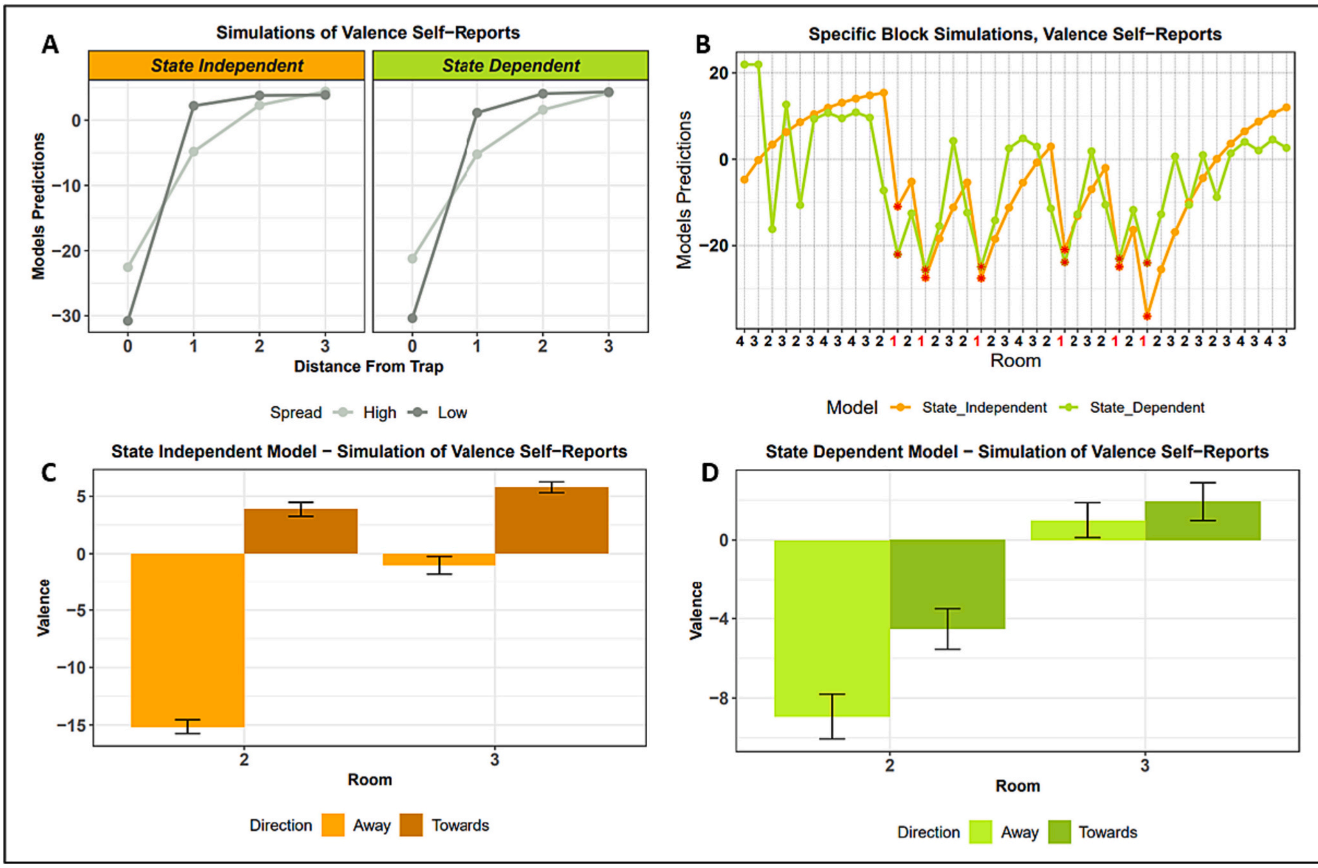

Fig. 2. Models Simulations for the Valence Self-Reports. (a) Whole Sample Simulations of the Two Models. Both models predict that average self-reports for each room will be related to the distance from the trap room, that is, more positive self-reports for Room 4 compared to rooms 3 and 2. In addition, both models can support high and low spread rates. In the state-independent model, the spread rate depends on the learning rate, whereas in the state-dependent model, it depends on the connectivity rate. (b) Simulation for a Specific Participant and Maze. In the state-independent model simulation, the negativity of the trap room spreads to the other rooms over time because it takes time until the single valence representation under this model recovers after a visit to the trap room. On the other hand, the state-dependent model produces a spread of valence through the structure of the state's space. The further away the room is from the trap, the more positive its valence is. Notably, the differences between the state's valence become smaller, that is, the spread of valence increases throughout the maze. (c) Direction and Structural effects on the State-Independent Model Simulation. The spread-through-time account predicts direction-dependent self-reports, that is, less positive self-reports when the direction is away from the trap room than toward the trap room. In addition, because in our design, the structural distance is positively correlated with the temporal distance, we can expect to see some structural effect. (d) Direction and Structural effects on the State-Dependent Model Simulation. According to the spread-through-structure account, once the participants learned the structure of the states, the self-reports should be room-dependent, that is, all the self-reports of the same room should be the same, regardless of the direction (away/toward the trap room).

reports will become more positive when the distance between the current room and the outcome room increases. Indeed, both model simulations predicted that average self-reports for each room would be related to the distance from the trap room, that is, more positive self-reports for Room 4 compared to Rooms 3 and 2 (Fig. 2a).

In addition, both models can support high and low valence spread rates. In the temporal, state-independent model, the spread rate depends on the learning rate, whereas in the structural, state-dependent model, it depends on the connectivity rate between states. The trend of more positive self-reports indicating a greater distance of the current room from the outcome room becomes more linear the higher the spread rate of the outcome's valence to the other three rooms. The trend becomes more quadratic when the spread rate is lower, that is, the more each room is isolated, and its valence depends gradually more on its own individual valence (Fig. 2a). Therefore, in our experimental design, it is possible to infer the rate of valence spread by examining the curvature and pairwise comparisons of the self-reports between rooms.

2.2.2. Dissociative predictions for direction and structural effects

The differences between the two models' predictions can be qualitatively appreciated when examining the trial-by-trial simulated valence representations (Fig. 2b). In the state-independent model, the negativity of the trap room spreads to the other rooms through temporal distance from experience. Therefore, the further in time a participant is from the last trap experience, the more positive the self-reports, even if the current room is Room 2, which is spatially close to the trap. On the other hand, the state-dependent model produces a spread of valence through structural distance. The further away the room is from the trap, the more positive its valence is. The distinction between the models can be observed when the participant moves from Room 4, the farthest from the trap room, toward the trap, that is, to Rooms 3 and 2. According to the state-independent model, self-reports become more positive because more time passes without visiting the trap room. According to the state-dependent model, self-reports become less positive because the participant moves to states closer to the trap room.

Notably, the structural distance between the rooms in our task is fixed, for example, the structural distance between Rooms 2 and 1 is always one, regardless of the direction of movement (from Room 1 to Room 2 or vice versa). However, the temporal distance between Rooms 2 and 1 is variable as sometimes participants visit Room 2 after Room 1, immediately after experiencing the negative outcome, and sometimes after visiting Room 3, with the negative outcome experience a long way in the past.

We can use the fact that the structural distance between the rooms is fixed and the temporal distance is variable to dissociate the two accounts' predictions for the valence self-reports of Rooms 2 and 3. According to the spread-through-structure account, valence self-reports are shaped by the *structural effect*. Representations for Room 2 should always be less positive than for Room 3, as it is closer to the negative outcome, and should always be the same as the distance does not change. According to the spread-through-time account, self-reports are shaped by the *direction effect*. The self-reports will be less positive when the direction of movement is *away* from the trap than *toward* the trap because the temporal distance is lower in the *away* direction and the continuous valence has less time to recover from the recent visit to the trap room. Notably, because in our design, the structural distance is positively correlated with the temporal distance, we can expect to see some *structural effect* under the spread-through-time account, but not the other way around.

In our simulations, we observe an apparent *direction effect* in the state-independent model simulations (Fig. 2c) and a clear *structural effect* in the state-dependent model simulations (Fig. 2d). In addition, because in our design, the structural distance is positively correlated with the temporal distance, we can expect to see some *structural effect* under the spread-through-time account. These results indicate that our experimental design can capture meaningful differences in the contribution of the two accounts to the valence spread by examining the *structural* and *direction* effects on self-reports in Rooms 2 and 3.

2.2.3. Persistence of the valence representations

The final measure in which the two accounts are expected to differ is the post-maze valence representation of the different states. The spread-through-structure account predicts that post-maze valence self-reports will be similar to mid-maze valence self-reports because the structural distance of each room from the trap room, learned while moving through the maze, is sustained at the end of the maze. This prediction is indeed captured by the four independent valence representations for each state stored by the state-dependent model simulations. On the other hand, because the temporal distance of each room from the outcome room is not sustained at the end of the maze, the spread-through-time account predicts that the post-maze valence self-reports of all the rooms will be the same and not informative. Accordingly, the simulations of the state-independent model indeed include only one valence representation to report.

3. Experiment 1

The first experiment aimed to provide an initial answer to our research questions, that is, how does the time versus structure accounts shape the spread of valence across states, and does the dependency on these accounts and the valence spread rate differ between the two formats of valence representation, affective and semantic.

3.1. Method and material

3.1.1. Transparency and openness

We report how we determined our sample size, all data exclusions, all manipulations, and all measures in the study, and we follow JARS (Kazak, 2018). All data and analysis codes are available at https://osf.io/6b2u9. Data were analyzed using R, version 4.2.2 (R Core Team, 2022). For a complete list of the packages used, see the supplementary material. This study's design and its analysis were pre-registered. Pre-registration of Experiment 1 can be viewed at https://aspredicted.org/HX8_488. Pre-registration of Experiment 2 can be viewed at https://aspredicted.org/ZYN_TN4.[4] Any discrepancies from the pre-registered analysis were marked as such. A complete analysis according to the pre-registration can be found in the supplementary material.

3.1.2. Participants

Participants included 367 UK residents (paid GBP 3.25 for 20 min of task performance) aged 18 to 65 ($M = 41.4$, $SD = 12.6$) recruited via the Prolific platform. The study was approved by the required ethics committee (Project ID Number: 206/22). We excluded 102 participants based on performance according to pre-registered exclusion criterion. Specifically, we calculated the standard deviation for the self-reports in 2 (affective valence/semantic valence) * 3 (first maze, second maze, third maze) * 2 (mid-maze self-reports, post-maze self-reports) = 12 clusters of self-reports. We required a standard deviation of 5 (10% of the total scale of 50 units) or more in each cluster. We excluded participants who failed to reach the required variability in their self-reports in two or more clusters (i.e., the standard deviation of self-reports <5). Notably, we checked the robustness of the reported results to the exclusion of participants by contrasting them to those obtained with a full sample of 367 participants, that is, when no participants were excluded. As Section 9 of the Supplementary material details, excluding participants primarily strengthened the obtained effects. Crucially, it did not create any effect that did not exist in the full, no-exclusion sample. We recruited participants until reaching the final pre-registered sample size of $n = 265$. This sample size was expected to detect a small interaction effect, that is, partial eta square of 0.03, found in pilot studies at a power of 80% and an α of 0.05 (Erdfelder, Faul, Buchner, & Lang, 2009)). The final sample's demographics are detailed in Table S1. The main characteristics of the pilot experiments are detailed in Section 7 of the supplementary material.

3.1.3. Measures

While moving through the maze, in 33% of the trials the participants were asked to give valence self-reports on the room they visited. Half the self-reports were semantic, and the other half were affective (Fig. 1c). At the end of each maze, the participants were asked to give affective and semantic self-reports on all four rooms that appeared in this maze (random order of rooms).

The participants reported the affective valence of the rooms by answering the question: "How do you feel in this room?". The self-reports were given on a bi-polar continuous slider scale ranging from 0 ("Very bad") to 50 (Very good").[5] Notably, the participants received specific instructions on reporting their experienced feelings to ensure the measurement was not contaminated with semantic knowledge (Hamzani, Mazar, Itkes, Petranker, & Kron, 2019; see the supplementary material for details). As part of these instructions, at the end of the maze the participants were instructed to first imagine being in the room reported and only then to answer how they felt.

The participants reported the semantic valence of the rooms by answering the question: "How close are you to the trap?". The self-reports were given on a bi-polar continuous slider scale ranging from 0 ("Very close") to 50 (Very far").[6] Notably, the question we ask here, that is, "How close are you to the trap?" deviates from the standard question used to measure semantic valence, that is, "How positive/

[4] There was a typo in the pre-registration regarding the percentage of the trials in which the participants would be asked to give a self-report. The pre-registration indicates 25% whereas the actual percentage used was as in Experiment 1, that is, 33%.

[5] Because the neutral point of this bi-polar scale is not intuitive, that is, 25, in the analysis we subtracted 25 from all the self-reports so the new neutral point is 0.

[6] Because the neutral point of this bi-polar scale is not intuitive, that is, 25, in the analysis we subtracted 25 from all the self-reports, so the new neutral point is 0.

negative is the stimulus?". The difference in the questions is due to the nature of the measured stimuli. Usually, the researchers are interested in measuring the semantic valence of stimuli with an a priori valence, that is, Unconditioned Stimulus like an image of a car accident or a puppy dog. In this case, the standard question of "How positive/negative is the stimulus?" can be answered based on the stimulus a priori valence. In our case, the stimuli are unique abstract images that, on purpose, have no a priori valence. Only during the associative learning task, when the trap room is constantly associated with a negative outcome, the stimuli acquire valence. This valence is solely based on the stimuli distance from the trap room. Therefore, in this study, semantic knowledge is best measured by the question: "How close are you to the trap?".[7]

3.2. Results

We will start by reporting the spread rate of the two formats of valence representation while moving through the maze. We will then check how the actual self-reports agree with the predictions and simulations of the two accounts.

3.2.1. The rate of valence spread

To check the rate of valence spread across states, we used a pre-registered repeated measures ANOVA with the mid-maze self-reports as a dependent variable and the distance from the trap room (0/1/2/3) and type of valence self-reports (Affective/Semantic) and their interactions as main effects (Fig. 3a). We found a significant effect of distance, $F(3, 792) = 335.2$, $p < 0.0001$, $\eta_p^2 = 0.56$, 95% CI [0.52,1.00], indicating that the participants discriminated between the valence of the different states in the task. We also found a significant effect of the type of valence self-report, $F(1, 264) = 173.4$, $p < 0.0001$, $\eta_p^2 = 0.40$, 95% CI [0.32,1.00], indicating that the participants used the self-reports scale differently in the two types of self-reports. We also found a significant interaction effect, $F(3, 792) = 19.0$, $p < 0.0001$, $\eta_p^2 = 0.07$, 95% CI [0.04,1.00].

To further check the interaction, we performed post hoc contrasts with p-value adjustment using the Tukey method for comparing a family of 8 estimates. There was a significant difference between the semantic self-reports of distance 2 ($M = 0.22$, $SD = 6.76$) and distance 3 ($M = 2.13$, $SD = 8.83$), $t(264) = 3.5$, $p = 0.013$, $d = 0.43$, 95% CI [0.19, 0.67]. However, the difference between the affective self-reports of distance 2 ($M = 4.19$, $SD = 6.24$) and distance 3 ($M = 3.97$, $SD = 8.02$) was not significant, $t(264) = 0.49$, $p = .99$. In addition, we directly compared the differences between the ratings of distances 3 and 2 in the two types of valence self-reports. The difference between the semantic self-reports of distances 3 and 2 ($M = 1.91$, $SD = 8.87$) was larger than the difference between the affective self-reports of distance 3 and distance 2 ($M = -0.22$, $SD = 7.39$), $t(264) = 4.06$, $p < 0.001$, $d = 0.5$, 95% CI [0.25,0.74]. This last finding indicates that the affective valence representation while moving through the maze spreads less than the semantic valence representation.

3.2.1.1. Fit of linear and quadratic mixed effects models. To further explore the possibility that the affective valence representation spreads less than the semantic valence representation, we compared the fit of two mixed effects regression models with the mid-maze affective (/semantic) valence self-reports as the dependent variable. The first assumes that the self-reports are a linear function of the distance from the trap room, that is, a high spread rate. The second assumes that the self-reports are a quadratic function of the distance from the trap room, that is, low spread rate (see also Fig. 2a). This analysis was not pre-registered.

We used group-level coefficients (fixed effects) to model population-level effects and individual-level coefficients (random effects) to capture average individual responses (Gelman & Hill, 2006). We report standardized coefficients, which represent the partial correlation between the dependent and independent variables and are, therefore, effect size indicators. We compared the model fitting scores BIC and AIC between the models using ANOVA (BIC – Bayesian information criterion, Schwarz, 1978; AIC – Akaike information criterion, Akaike, 1974).

As shown in Table 1, the quadratic model better explains both the affective and semantic self-reports. However, in the affective self-reports, the percentage of improvement in terms of AIC/BIC as a result of adding the quadratic term is more considerable, that is, 1.1% versus 0.1%. See also Figs. S1 and S2 in Section 8.1 of the supplementary material, showing the fitting of the linear and quadratic models to the affective and semantic self-reports. This result further supports the claim that the affective valence representation spreads less than the semantic valence representation.

3.2.2. Dissociative direction and structural effects

To dissociate the influence of the direction of movement and structure on the affective self-reports, we ran an unregistered repeated measures ANOVA, with the affective mid-maze self-reports of Rooms 2 and 3 as a dependent variable, and the Room number (2/3), the direction (*away* from the trap/*toward* the trap) and their interactions as main effects (Fig. 3c). We found a significant effect of room number, that is, a structural effect, $F(1, 244$[8]$) = 139.3$, $p < 0.0001$, $\eta_p^2 = 0.36$, 95% CI [0.29,1.00], supporting the spread-through-structure account. We also found a significant direction effect, $F(1, 244) = 6.74$, $p = 0.01$, $\eta_p^2 = 0.03$, 95% CI [0,1], supporting the spread-through-time account, which predicts that the same room will receive more positive self-reports when the direction is toward the trap. The interaction effect was not significant, $F(1, 244) = 0.02$, $p = .89$.

We performed the same analysis for the semantic mid-maze self-reports (Fig. 3d). We found a significant effect of room number, that is, a structural effect, $F(1, 250$[9]$) = 108.2$, $p < 0.0001$, $\eta_p^2 = 0.39$, 95% CI [0.31,1.00], supporting the spread-through-structure account. The direction effect was not significant, $F(1, 250) = 0.24$, $p = .62$. We also found a significant interaction effect, $F(1, 250) = 7.3$, $p = 0.007$, $\eta_p^2 = 0.03$, 95% CI [0,1]. This result indicates a localized direction effect, limited to Room 2. The self-reports immediately after visiting the trap room were less positive than those immediately preceding a visit to the trap room. However, the direction effect was not observed in Room 3, which is further from the trap.

3.2.3. Persistence of the valence representations

We examined the post-maze self-reports to further check whether the participants based their valence reports on one continuous representation or on four different representations, one for each state (time versus structure accounts, respectively). The post-maze self-reports were the dependent variable in a pre-registered repeated-measure ANOVA, and the distance from the trap room (0/1/2/3), type of self-reports (Affective/Semantic), and their interactions were the main effects (Fig. 3.B). We found a significant effect for distance from the trap, $F(3, 792) = 170.7$, $p < 0.0001$, $\eta_p^2 = 0.39$, 95% CI [0.35,1.00], indicating that when both types of self-reports are taken together, the participants discriminated between the valence of the different states in the task during the post-maze phase. Notably, as detailed below, this discrimination is limited to semantic self-reports. We also found a significant effect for the type of valence self-reports, $F(1, 264) = 74.1$, $p < 0.0001$, $\eta_p^2 = 0.22$,

[7] For similar deviation from the standard version of instructions to measure semantic valence, see Section 2.1.4 of Heimer et al. (2023).

[8] 20 participants had some missing data and were therefore excluded **only** from this analysis.

[9] 14 participants had some missing data and were therefore excluded **only** from this analysis.

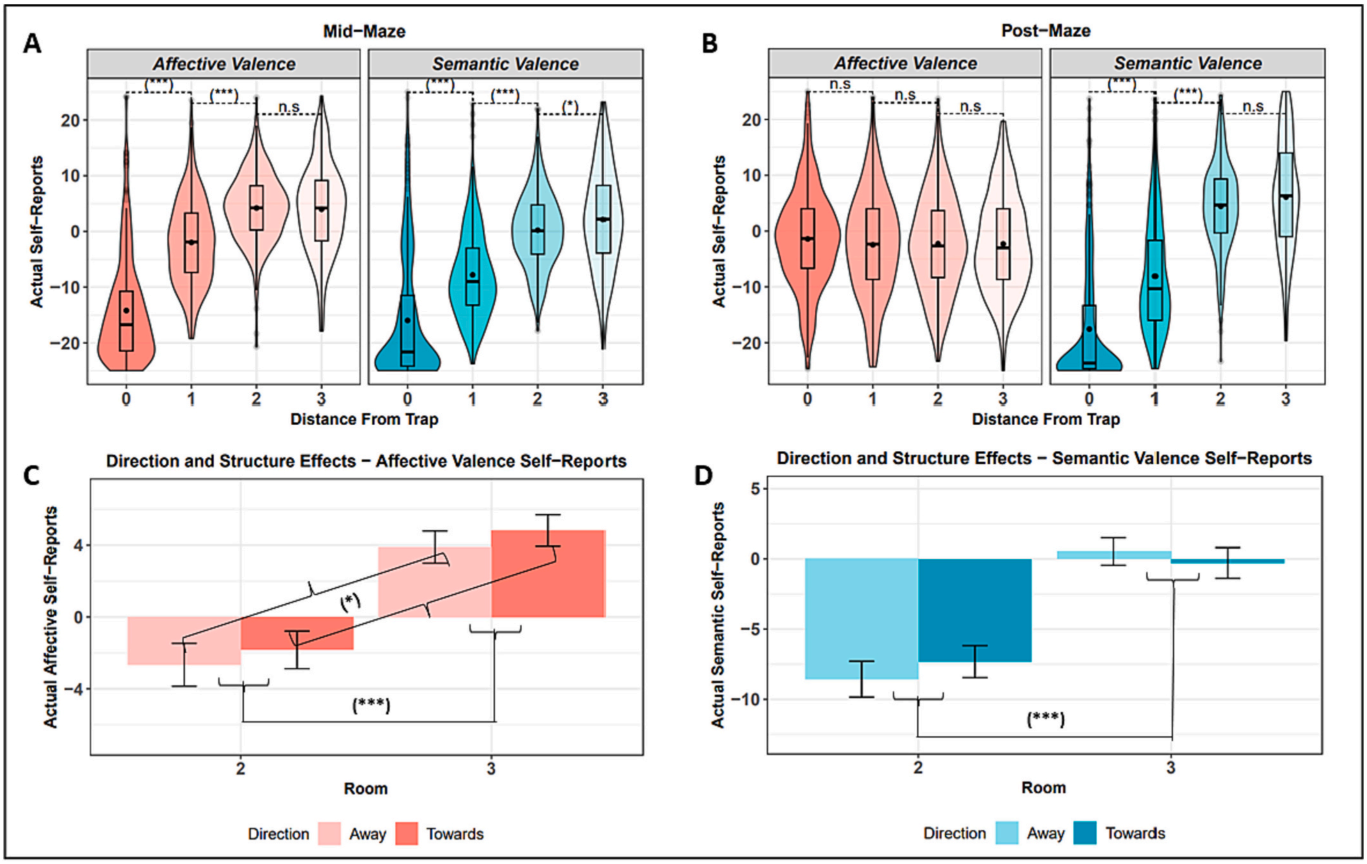

Fig. 3. Actual self-reports of Experiment 1. (*) < 0.05, (**) < 0.01, (***) < 0.001 (a) Affective and Semantic Valence Mid-Maze Self-Reports According to the Distance from the Trap. A significant difference exists between the semantic self-reports of distances 2 and 3. However, the difference between the affective self-reports of distances 2 and 3 was not significant. In addition, the difference between the semantic self-reports of distances 3 and 2 was larger than the difference between the affective self-reports of distances 3 and 2, indicating that the affective valence representation while moving through the maze spreads less than the semantic valence representation. (b) Affective and Semantic Valence Post-Maze Self-Reports According to the Distance from the Trap. There was no evidence of differences between the various distances from the trap (e.g., distance 0 versus distance 1) in the affective valence self-reports, indicating that the affective valence representations faded away at the end of the maze, as predicted by the spread-through-time account. On the other hand, as predicted by the spread-through-structure account, there were significant differences between the semantic self-reports of distances 0, 1, and 2. (c) The Effect of Direction and Structure on Affective Valence Self-Reports. The direction effect was significant, supporting the spread-through-time account, which predicts that the same room will receive more positive self-reports when the direction is toward the trap. (d) The Effect of Direction and Structure on Semantic Self-Reports. The direction effect was not significant, supporting the spread-through-structure account, which predicts that the same room will receive the same self-reports, regardless of the previous room visited.

Table 1
The results of the linear and quadratic mixed effects in Experiment 1.

Dependent variable	Affective self -reports			Semantic self -reports		
Model	Linear	Quadratic	Difference	Linear	Quadratic	Difference
Intercept standardized coefficient	0.001	−0.003	0.00	−0.003	−0.003	0.00
Intercept Significance level	***	***		***	***	
Distance from trap – standardized coefficient	0.51	1.29	0.78	0.46	0.75	0.29
Distance from trap – Significance level	***	***		***	***	
Distance from trap ˆ 2 – standardized coefficient	NA	−0.81	NA	NA	−0.31	NA
Distance from trap ˆ 2 – Significance level	***	***		***	***	
AIC	65,191	**64,471**	−720	66,963	**66,870**	−93
BIC	65,219	**64,506**	−713	66,991	**66,905**	−86

* <0.05, ** < 0.01, *** < 0.001.

95% CI [0.15,1.00], indicating that, on average, across all rooms, the semantic self-reports are less positive than the affective self-reports. Notably, we also found a significant interaction effect, $F(3, 792) = 182.8$, $p < 0.0001$, $\eta_p^2 = 0.41$, 95% CI [0.37,1.00], which is considerably larger than the interaction effect found in the mid-maze self-reports ($\eta_p^2 = 0.07$).

To further check the interaction, we performed post hoc contrasts with *p*-value adjustment using the Tukey method for comparing a family of 8 estimates. Post-maze affective self-reports were similar for all rooms, regardless of their distance from the trap (e.g., distance 0 versus distance 1) in the affective valence self-reports, indicating that the affective valence representations faded away at the end of the maze, as predicted by the spread-through-time account. On the other hand, as predicted by the spread-through-structure account, there were significant differences between the semantic self-reports of the distance 0 ($M = -17.56$, $SD = 11.01$) and distance 1 ($M = -8.09$, $SD = 9.78$), $t(264) =$

12.6, $p < 0.0001 = 1.55$, 95% CI [1.27, 1.82], and between distance 1 ($M = -8.09$, $SD = 9.78$) and distance 2 ($M = 4.44$, $SD = 7.93$), $t(264) = 16.28$, $p < 0.0001$, $d = 2.00$, 95% CI [1.71, 2.3]. The difference between distance 2 ($M = 4.44$, $SD = 7.93$) and distance 3 ($M = 6.08$, $SD = 10.1$) was not significant $t(264) = 2.58$, $p = .17$.

3.3. Discussion

In the first experiment, we used a four-state association task in which the leftmost state had negative valence to investigate whether different formats of valence representation, that is, semantic and affective, are different in their dependency on the two accounts, that is, time and structure, and their spread rate.

To check for possible differences in the spread rate of the affective and semantic formats of valence representation, we analyzed the actual self-reports given by the participants. The results of the ratings made by participants during the maze task indicate that the affective format of the trap's negativity, that is, negative feelings, spread less to the other states than the semantic format of the trap's negativity, that is, knowledge of the distance between each room and the trap. In addition, by examining direction and structural effects on ratings, we found that both affective and semantic valence representations demonstrated a *structural effect*, supporting the spread-through-structure account. However, the affective valence representation also demonstrated a *direction effect*, supporting the spread-through-time account. The self-reports made when participants moved away from the trap (in close temporal proximity to the trap) were less positive than when moving toward the trap (further in time from the trap experience). Finally, self-reports at the maze's end clearly demonstrated a dissociation between the two valence representation formats. The distinction between the affective valence representations of the different states fades away, as predicted by the spread-through-time account. In contrast, the semantic valence post-maze representations were similar to the mid-maze representations, as predicted by the spread-through-structure account.

4. Experiment 2

The aim of Experiment 2 was to extend Experiment 1 and study how outcome characteristics affect valence spread. First, we wanted to replicate the results of Experiment 1 using a different type of outcome – monetary reward instead of waiting time. Second, we wanted to examine the effect of the outcome's valence, that is, positive or negative, on our research questions. Specifically, we aimed to check whether the outcome's valence influences the dependency of the two formats of valence representation (i.e., semantic and affective) on the two mechanisms for valence spread (i.e., time and structure). We also aimed to check whether the outcome's valence influences the spread rate of the affective and semantic valence. For this purpose, we used the same four-state associative learning task as in Experiment 1. However, this time we manipulated the valence of the outcome as a between-participants variable. In the gain condition, the participants gained points in the outcome room (i.e., the leftmost room), whereas in the loss condition, the participants lost points in the outcome room. We also shortened the task to two mazes of 65 trials each instead of the three mazes used in Experiment 1.

4.1. Method and material

4.1.1. Participants

Participants included 403 UK residents (paid GBP 2.00 for a 15-min task plus an average bonus of GBP 1.00) aged 18 to 65 ($M = 39.0$, $SD = 12.7$) recruited via the Prolific platform. The study was approved by the required ethics committee (Project ID Number: 206/22). Eighty-three participants were excluded based on performance according to a pre-registered exclusion criterion. Specifically, we calculated the standard deviation for the self-reports in 2 (affective valence/semantic valence) * 2 (first maze, second maze) * 2 (mid-maze self-reports, post-maze self-reports) = 8 clusters of self-reports. We required the standard deviation in each cluster to be 5 (10% of the total scale of 50 units) or more. We excluded participants who failed to reach the required variability in their self-reports in two or more clusters (i.e., the standard deviation of self-reports <5). Notably, we checked the robustness of the reported results to the exclusion of participants by contrasting them to those obtained with a full sample of 403 participants, that is, when no participants were excluded. As Section 9 of the Supplementary file details, excluding participants primarily strengthened the obtained effects but did not create effects that did not exist in the full sample. In addition, out of the 34 effects checked for, two only exist when no participants are excluded. We added notifications on these effects in the relevant places in our results.

We recruited participants until we reached a total sample of 320 participants, at least 135 participants in each condition, that passed our exclusion criteria as pre-registered. This sample size was expected to detect a small interaction effect, that is, partial eta square of 0.05, at a power of 80% and an α of 0.05 ((Erdfelder, Faul, Buchner, & Lang, 2009). One hundred and seventy-two participants were randomly allocated to the loss condition, in which the participants started with 200 points, each equivalent to 1 penny, and lost 5 points in each visit to the outcome room. One hundred and forty-eight participants were randomly allocated to the gain condition, in which the participants started with 0 points and gained 5 points in each visit to the outcome room. The final samples of the two conditions did not differ in their main demographic characteristics. See Table S1 for details.

4.2. Results

4.2.1. The rate of valence spread

4.2.1.1. Gain condition. To check the spread across states of the semantic and affective valence representations in the gain condition, we used a pre-registered repeated measures ANOVA, with the mid-maze self-reports of the gain condition as a dependent variable and the distance from the outcome room (0/1/2/3) and type of self-reports (Affective/Semantic) and their interactions as main effects (Fig. 4a). Notably, as pre-registered, we transformed the affective self-reports only in the gain condition using the formula: Transformed self-reports = |50 – Original Self-reports|. This transformation enables a comparison between the affective self-reports in the loss and gain conditions.

We found a significant effect of distance, $F(3, 426) = 97.5$, $p < 0.0001$, $\eta_p^2 = 0.41$, 95% CI [0.35,1.00], indicating that the participants discriminated between the valence of the different states in the task. The type of valence self-report was not significant, $F(1, 142) = 0.7$, $p = .41$. We found a significant interaction effect, $F(3, 426) = 2.7$, $p = 0.05$, $\eta_p^2 = 0.02$, 95% CI [0,1], indicating that the valence spread rate of the types of self-reports is different. To further check the interaction, we performed post hoc contrasts with p-value adjustment using the Tukey method for comparing a family of 8 estimates. Notably, there was a significant difference between the semantic self-reports of distance 2 ($M = 0.35$, $SE = 0.63$) and distance 3 ($M = 3.1$, $SE = 0.75$), $t(142) = 3.2$, $p = 0.03$, $d = 0.54$, 95% CI [0.2, 0.87]. However, the difference between the affective self-reports of distance 2 ($M = -0.52$, $SE = 0.65$) and distance 3 ($M = 1.19$, $SE = 0.64$) was not significant, $t(142) = 2.3$, $p = .28$. In addition, we directly compared the differences between the ratings of distance 3 and distance 2 in the two types of valence self-reports. The difference between the semantic self-reports of distance 3 and distance 2 ($M = 2.74$, $SD = 10.3$) was larger than the difference between the affective self-reports of distance 3 and distance 2 ($M = 1.88$, $SD = 8.8$), t(144[10]) =

[10] 3 participants had some missing data and were therefore excluded **only** from this analysis.

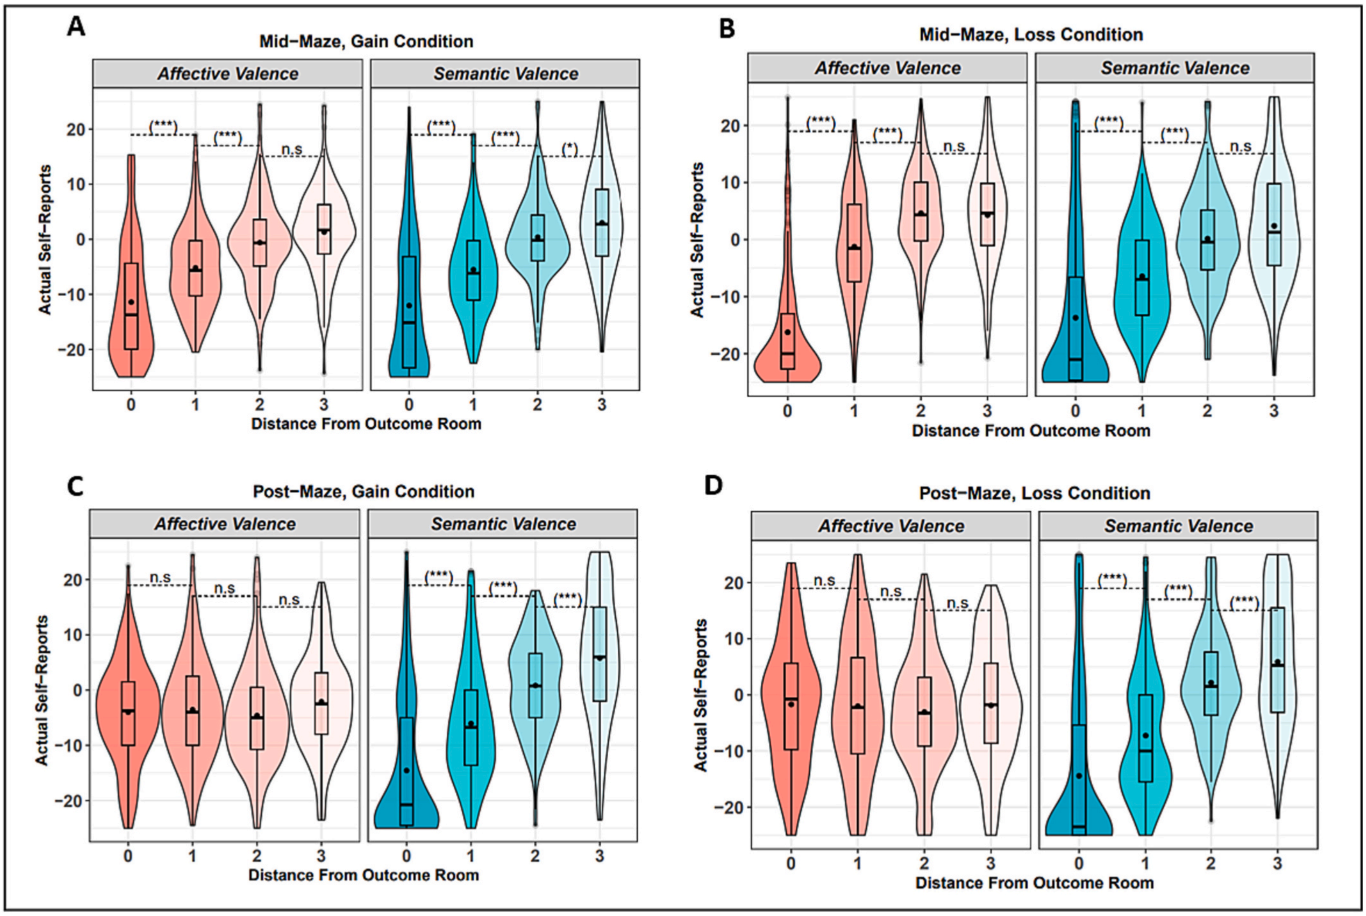

Fig. 4. Experiment 2 Results. (*) < 0.05, (**) < 0.01, (***) < 0.001. (a) Affective and Semantic Mid-Maze Self-Reports According to the Distance from the Outcome Room in the Gain Condition. A significant difference exists between the semantic self-reports of distance 2 and distance 3. However, the difference between the affective self-reports of distances 2 and 3 was not significant. In addition, The difference between the semantic self-reports of distances 3 and 2 was larger than the difference between the affective self-reports of distances 3 and 2 indicating that the affective valence representation while moving through the maze spreads less than the semantic valence representation. (b) Affective and Semantic Mid-Maze Self-Reports According to the Distance from the Outcome Room in the Loss Condition. A significant interaction effect results from a greater spread of the semantic valence versus the affective valence. (c) Affective and Semantic Post-Maze Self-Reports According to the Distance from the Outcome Room in the Gain Condition. There was a significant interaction effect. In the affective self-reports, the differences between the various rooms faded away at the end of the maze. In the semantic valence self-reports, all the differences between the various distances from the outcome room were significant. (d) Affective and Semantic Post-Maze Self-Reports According to the Distance From the Outcome Room in the Loss Condition. As in the Gain condition, there was a significant interaction effect. In the affective self-reports, the differences between the various rooms faded away at the end of the maze. On the other hand, in the semantic valence self-reports, all the differences between the various distances from the outcome rooms were significant.

3.9, $p < 0.001$, $d = 0.65$, 95% CI [0.31,0.98]. This last finding indicates that the affective valence representation while moving through the maze spreads less, that is, was more quadratic, than the semantic valence representation.

4.2.1.2. Loss condition. We repeated the analysis with the mid-maze self-reports of the loss condition as a dependent variable (Fig. 4b). We found a significant effect for the distance from the outcome room, $F(3, 513) = 201.5$, $p < 0.0001$, $\eta_p^2 = 0.54$, 95% CI [0.5,1]. The type of valence self-report was also significant, $F(1, 171) = 33.04$, $p < 0.0001$, $\eta_p^2 = 0.16$, 95% CI [0.09,1]. We also found a significant interaction effect, $F(3, 513) = 16.9$, $p < 0.0001$, $\eta_p^2 = 0.09$, 95% CI [0.05,1], larger than the gain condition's similar interaction effect ($\eta_p^2 = 0.02$). To further check the interaction, we performed post hoc contrasts with p-value adjustment using the Tukey method for comparing a family of 8 estimates. The difference between the semantic self-reports of distance 2 ($M = 0.15$, SE $= 0.65$) and distance 3 ($M = 2.39$, SE $=0.76$), $t(171) = 3.03$, $p = 0.06$ was not significant, but approached significance.[11] The difference between the affective self-reports of distance 2 ($M = 4.6$, SE $= 0.57$) and distance 3 (M $= 4.3$, SE $= 0.67$) was not significant, t(171) $= 0.47$, $p = .99$. In addition, we directly compared the differences between the ratings of distance 3 and distance 2 in the two types of valence self-reports. The difference between the semantic self-reports of distance 3 and distance 2 ($M = 2.24$, SD $= 9.73$) was larger than the difference between the affective self-reports of distance 3 and distance 2 ($M = –0.32$, SD $= 8.97$), $t(171) = 2.59$, $p = 0.01$, $d = 0.4$, 95% CI [0.09,0.7]. This last finding indicates that the affective valence representation while moving through the maze spreads less, that is, is more quadratic, than the semantic valence representation.

4.2.1.3. Fit of linear and quadratic mixed effects models. As in Experiment 1, to further explore the possibility that the affective valence representation spreads less than the semantic valence representation, we compared the fit of two mixed effects regressions for the self-reports (see

[11] In the full sample scenario, there is significant effect, $p = .02$. $\eta_p^2 = 0.38$.

Section 3.2.1.1 for details). The first assumes that the self-reports are a linear function of the distance from the outcome room, that is, a high spread rate. The second assumes that the self-reports are a quadratic function of the distance from the trap room, that is, a low spread rate (see also Fig. 2a). This analysis was not pre-registered.

In the gain condition (see Table 2), the quadratic model fit to the semantic valence self-reports was almost identical to the linear model fit. In the affective valence self-reports, the percentage of improvement in terms of AIC/BIC as a result of adding the quadratic term is less considerable than in the affective valence self-reports of Experiment 1 (0.2% in the current experiment versus 1.1% in Experiment 1). See also Figs. S3 and S4 in Section 8.2 of the supplementary material, showing the fitting of the linear and quadratic models to the affective and semantic self-reports in the gain condition. We conclude that in the gain condition, the difference between the spread rate of both types of valence narrows. Both types are characterized by a relatively high rate of spread, that is, a linear trend.

The quadratic model better explains the affective self-reports in the loss condition (see Table 3). The percentage of improvement in terms of AIC/BIC as a result of adding the quadratic term is 1.3%. However, in the case of semantic self-reports, the improvement that results from adding the quadratic term is less considerable, that is, 0.1%. See also Figs. S5 and S6 in Section 8.3 of the supplementary material, showing the fitting of the linear and quadratic models to the affective and semantic self-reports in the loss condition. This result further supports the claim that the affective valence representation spreads less than the semantic valence representation in the loss condition.

4.2.2. Dissociative direction and structural effects

4.2.2.1. Gain condition. To dissociate the effect of the direction of movement and structure on the affective self-reports, we ran an unregistered repeated measures ANOVA, with the affective mid-maze self-reports of Rooms 2 and 3 as a dependent variable, and room number (2/3), direction (away from the trap/toward the trap) and their interactions as main effects (Fig. 5a). We found a significant effect of room number, that is, a *structural effect*, $F(1, 106$[12]$) = 27.01$, $p < 0.0001$, $\eta_p^2 = 0.2$, 95% CI [0.1,1.00], supporting the spread-through-structure account. The direction effect was not significant, $F(1, 106) = 0.91$, $p = .34$, as was the interaction effect, $F(1, 106) = 1.63$, $p = .2$.

We repeated the analysis for the semantic mid-maze self-reports (Fig. 5b). We found a significant effect of room number, $F(1, 110$[13]$) = 42.57$, $p < 0.0001$, $\eta_p^2 = 0.28$, 95% CI [0.17,1.00], that is, a *structural effect*, supporting the spread-through-structure account. The direction effect was insignificant, $F(1,110) = 0.54$, $p = .47$. We found a significant interaction effect, $F(1, 119) = 11.79$, $p = 0.0008$, $\eta_p^2 = 0.1$, 95% CI [0.03,1]. This result indicates a localized direction effect, limited to Room 2. The self-reports immediately after visiting the trap room were less positive than those immediately preceding a visit to the trap room. However, the direction effect was not observed in Room 3, which is further from the trap.

4.2.2.2. Loss condition. We repeated the unregistered analysis for the affective mid-maze self-reports in the loss condition (Fig. 5c). We found a significant effect of room number, that is, a *structural effect*, $F(1, 126$[14]$) = 39.68$, $p < 0.0001$, $\eta_p^2 = 0.24$, 95% CI [0.14,1.00], supporting the spread-through-structure account. The direction effect was insignificant, $F(1, 126) = 1.71$, $p = .19$. Notably, we failed to replicate the significant direction effect found in the affective mid-maze self-reports of Experiment 1, where the outcome was also negative ($\eta_p^2 = 0.03$). The interaction effect, $F(1, 126) = 0.04$, $p = .82$, was also not significant.

We repeated the analysis for the semantic mid-maze self-reports (Fig. 5d). We found a significant effect of room number, that is, a *structural effect*, $F(1, 127$[15]$) = 52.18$, $p < 0.0001$, $\eta_p^2 = 0.29$, 95% CI [0.19,1.00], supporting the spread-through-structure account. The direction effect was marginal but not significant, $F(1, 127) = 3.25$, $p = 0.07$, as was the interaction effect, $F(1, 127) = 3.54$, $p = 0.06$.[16]

4.2.3. Persistence of the valence representations

4.2.3.1. Gain condition. To check whether the participants based their valence reports on one continuous representation or on four different representations, one for each state (time versus structure accounts, respectively), we used a pre-registered repeated measures ANOVA. This time the post-maze self-reports were the dependent variable, and the distance from the outcome room (0/1/2/3), the type of self-reports (Affective/Semantic), and their interactions were the main effects (Fig. 4c). We found a significant effect for the distance from the outcome room, $F(3, 441) = 63.7$, $p < 0.0001$, $\eta_p^2 = 0.3$, 95% CI [0.24,1.00], indicating that when both types of self-reports are taken together, the participants discriminated between the valence of the different states in the task during the post-maze phase. Notably, as detailed below, this discrimination is limited to semantic self-reports. The type of self-reports was not significant, $F(1, 147) = 0.03$, $p = 0.85$. We also found a significant interaction effect, $F(3, 441) = 48.3$, $p < 0.0001$, $\eta_p^2 = 0.25$, 95% CI [0.19,1].

To further check the interaction, we performed post hoc contrasts with *p*-value adjustment using the Tukey method for comparing a family of 8 estimates. Notably, there was no evidence of differences between the various distances from the outcome room (e.g., distance 0 versus distance 1) in the affective valence self-reports, indicating that the affective valence representations faded away at the end of the maze, as predicted by the spread-through-time account. On the other hand, as predicted by the spread-through-structure account, there were significant differences between the various distances from the outcome room in the semantic self-reports. Specifically, there was a significant difference between distance 0 ($M = -14.6$, $SE = 1.03$) and distance 1 ($M = -6.1$, $SE = 0.86$), $t(147) = 8.1$, $p < 0.0001$, $d = 1.34$, 95% CI [0.98,1.69]. There was also a significant difference between distance 1 ($M = -6.1$, $SE = 0.86$) and distance 2 ($M = 0.82$, $SE = 0.68$), $t(147) = 5.92$, $p < 0.0001$, $d = 0.86$, 95% CI [0.5,1.19] and between distance 2 ($M = 0.82$, $SE = 0.68$) and distance 3 ($M = 5.79$, $SE = 0.97$), $t(147) = 5.06$, $p < 0.0001$, $d = 0.83$, 95% CI [0.5,1.17].

4.2.3.2. Loss condition. In a similar pre-registered ANOVA, we found a significant effect for the distance from the outcome room, $F(3, 513) = 49.4$, $p < 0.0001$, $\eta_p^2 = 0.22$, 95% CI [0.17,1.00] (Fig. 4d), indicating that when both types of self-reports are taken together, the participants discriminated between the valence of the different states in the task during the post-maze phase. Notably, as detailed below, this discrimination is limited to semantic self-reports. The type of self-reports was also significant, $F(1, 171) = 9.25$, $p = 0.003$, $\eta_p^2 = 0.05$, 95% CI [0.01,1.00], indicating that, on average, across all rooms, the semantic self-reports are less positive than the affective self-reports. We also found a significant interaction effect, $F(3, 513) = 55.1$, $p < 0.0001$, $\eta_p^2 = 0.24$, 95% CI [0.19,1]. To further check the interaction, we performed

[12] 41 participants had some missing data and were therefore excluded **only** from this analysis.

[13] 37 participants had some missing data and were therefore excluded **only** from this analysis.

[14] 41 participants had some missing data and were therefore excluded **only** from this analysis.

[15] 40 participants had some missing data and were therefore excluded **only** from this analysis.

[16] In the full sample scenario, there is significant effect, $p = .03$. $\eta_p^2 = 0.03$.

Table 2
The results of the linear and quadratic mixed effects in Experiment 2, gain condition.

Dependent variable	Affective self-reports			Semantic self-reports		
Model	Linear	Quadratic	Difference	Linear	Quadratic	Difference
Intercept standardized coefficient	0.003	0.003	0.00	0.004	0.004	0.00
Intercept significance level	***	***		***	***	
Distance from trap – Standardized coefficient	0.41	0.78	0.37	0.39	0.55	0.16
Distance from trap – Significance level	***	***		***	***	
Distance from trap ˆ 2 – Standardized coefficient	NA	−0.38	NA	NA	−0.17	NA
Distance from trap ˆ 2 – Significance level	***	***		***	***	
AIC	23,858	**23,810**	−48	25,595	**25,588**	−7
BIC	23,882	**23,840**	−42	25,620	**25,618**	−2

* <0.05, ** < 0.01, *** < 0.001.

Table 3
The results of the linear and quadratic mixed effects in Experiment 2, loss condition.

Dependent variable	Affective self-reports			Semantic self-reports		
Model	Linear	Quadratic	Difference	Linear	Quadratic	Difference
Intercept standardized coefficient	0.001	0.000	0.00	0.002	0.002	0.00
Intercept significance level	***	***		***	***	
Distance from trap – Standardized coefficient	0.52	1.38	0.86	0.39	0.68	0.29
Distance from trap – Significance level	***	***		***	***	
Distance from trap ˆ 2 – tandardized coefficient	NA	−0.9	NA	NA	−0.3	NA
Distance from trap ˆ 2 – Significance level	***	***		***	***	
AIC	28,198	**27,809**	−389	29,833	**29,798**	−35
BIC	28,223	**27,840**	−383	29,858	**29,830**	−28

* <0.05, ** < 0.01, *** < 0.001.

post hoc contrasts with p-value adjustment using the Tukey method for comparing a family of 8 estimates. As found in the gain condition, there was no evidence for differences between the various distances from the outcome in the affective valence self-reports, supporting the spread-through-time account. On the other hand, in the semantic valence self-reports, all the differences between the various distances from the outcome room were significant, as predicted by the spread-through-structure account. Specifically, there was a significant difference between distance 0 ($M = -14.4$, $SE = 1.13$) and distance 1 ($M = -7.26$, $SE = 0.83$), $t(171) = 5.9$, $p < 0.0001$, $d = 0.9$, 95% CI [0.59,1.22]. There was also a significant difference between distance 1 ($M = -7.26$, $SE = 0.83$) and distance 2 ($M = 2.17$, $SE = 0.7$), $t(171) = 9.36$, $p < 0.0001$, $d = 1.43$, 95% CI [1.09,1.77] and between distance 2 ($M = 2.17$, $SE = 0.7$) and distance 3 ($M = 5.88$, $SE = 0.9$), $t(171) = 4.48$, $p = 0.0004$, $d = 0.69$, 95% CI [0.38,0.99].

4.3. Discussion

In the second experiment, we aimed to examine how the outcome's valence affects the spread of affective and semantic valence across states. To this end, we used the same four-state associative learning task as in Experiment 1, but this time, we manipulated the valence of the outcome as a between-participants variable. In one condition, the participants lost points when visiting the leftmost room, that is, the outcome room. This condition was similar to Experiment 1, where the participants lost time visiting the leftmost room. In the second condition, the participants gained points in the outcome room. This condition enabled us to extend our inquiry and check directly whether the valence of the outcome, that is, negativity versus positivity, affects its spread across states.

Most of the loss condition results replicated Experiment 1's results. As found in Experiment 1, while moving through the maze, the affective representation of the negative valence spread less than the semantic representation of the same valence. At the end of the maze, the dissociation between the two types of valence representations increased. The distinction between the affective valence representations of the different states fades away, as predicted by the spread-through-time account. On the other hand, the semantic valence representations are similar to the mid-maze representations, as predicted by the spread-through-structure account. Additional support for the reliance of the semantic valence representation on the spread-through-structure account is provided by the lack of a direction effect (i.e., moving away/toward the outcome room) on the self-reports of the same room. Notably, we could not find an effect of direction in the affective valence self-reports. However, we did find such an effect, which indicates reliance on the spread-through-time account, in Experiment 1.

Moving to the gain condition results, the positivity of the outcome narrows the differences between the two types of valence representations found in the case of a negative outcome. Specifically, the affective mid-maze self-reports in the gain condition still spread less than the semantic self-reports but to a lesser extent than in the Loss condition. Notably, the move to positive valence did not change the pattern of post-maze self-reports found in Experiment 1 and in the loss condition of Experiment 2. We again found strong evidence for the reliance of the affective valence representation on the spread-through-time account. Similarly, we found strong evidence for the reliance of the semantic valence representation on the spread-through-structure account. The move to positive valence also did not affect the direction versus structural effects analysis results. We again found evidence for the reliance of both types of valence representation on the spread-through-structure account.

5. General discussion

The spread of an outcome's valence, that is, its positivity or negativity, to the events (/states) that led to this outcome, is crucial to optimizing behavior. Behavioral and brain studies suggest that valence spread is an efficient algorithmic solution to sequential-state planning that is used by humans. We offered two mechanistic accounts of valence spread, time and structure, each supported by a unique computational model and predictions. In two pre-registered experiments ($N = 585$), we examined the contribution of these two accounts to the development of two different formats of valence representation for the different states, semantic and affective, and the spread rate of the outcome's valence

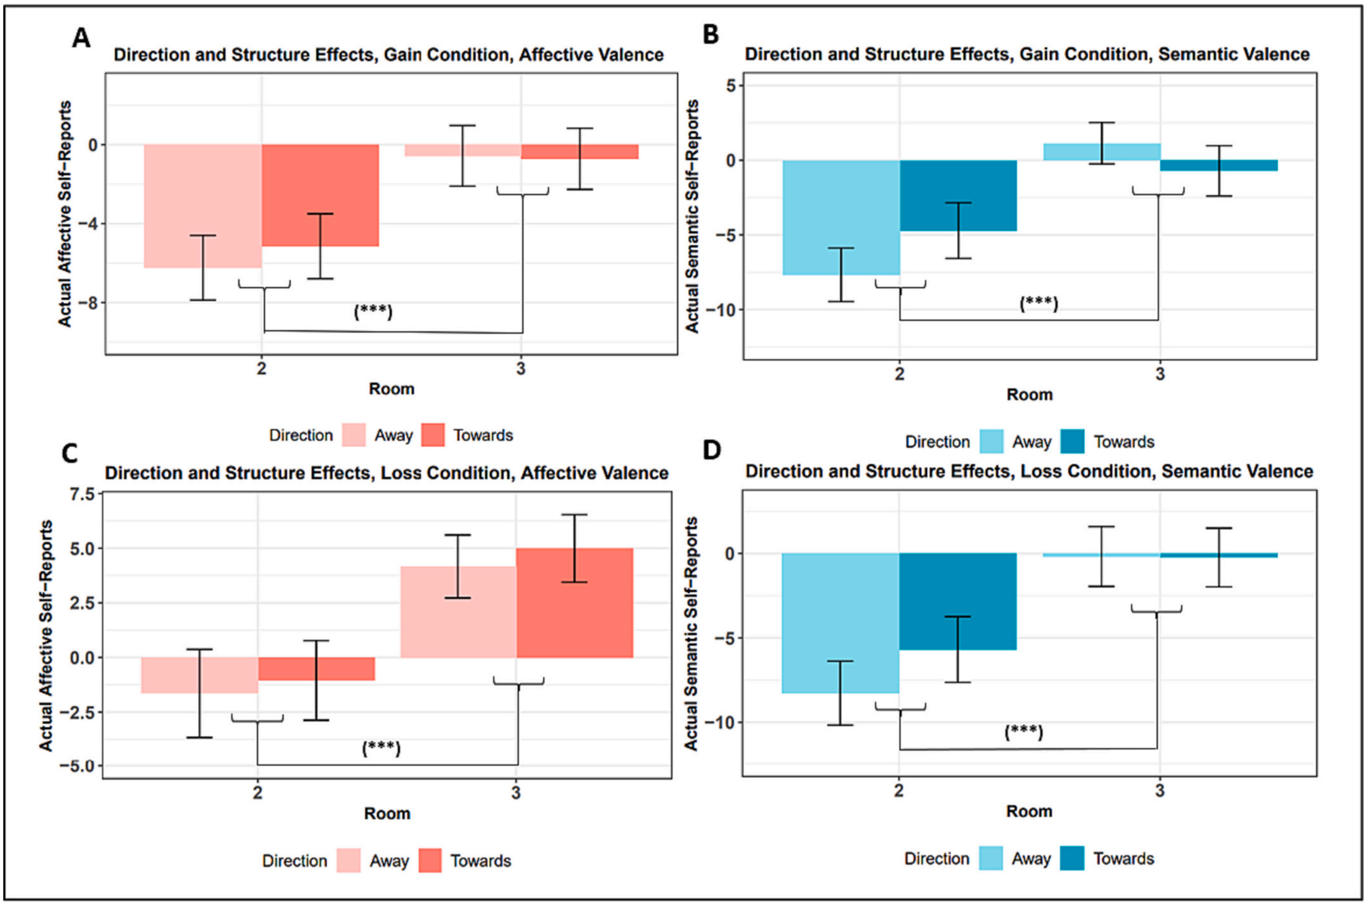

Fig. 5. Direction and Structural effects in Experiment 2. (a) Affective Self-Reports, Gain Condition. The direction effect was not significant. The structural effect was significant, supporting the spread-through-structure account, which predicts that the same room will receive the same self-reports regardless of the previous room visited. (b) Semantic Self-Reports, Gain Condition. The direction effect was not significant. The structural effect was significant, supporting the spread-through-structure account. We found a significant interaction effect resulting from the fact that in Room 2, the away self-reports were significantly less positive than the toward self-reports but in Room 3 there was no significant difference between the away and toward self-reports. (c) Affective Self-Reports, Loss Condition. The direction effect was not significant. The structural effect was significant, supporting the spread-through-structure account. (d) Semantic Self-Reports, Loss Condition. The direction effect was not significant. The structural effect was significant, supporting the spread-through-structure account.

across states in each format.

Our results show that structure and temporal dynamics shape the spread of valence representations across states and that valence representations with different formats follow somewhat different spread dynamics.

5.1. Summary of findings

In both experiments, we found that while moving through the maze, the affective representation of the outcome's negativity, that is, the conscious negative feelings associated with the negative experience, spread less to the other states than did the semantic representation, that is, the knowledge about the distance of each state from the outcome room. Notably, the dissociation between the spread rates of the affective and semantic valence representations is considerably moderated when the outcome is positive. In this case, both formats of valence representation exhibit a relatively high spread of the outcome's valence. To summarize, the existence of a dissociation between the affective and semantic valence mid-maze spread dynamics depends on the specific condition of a negative outcome.

The post-maze self-reports in both experiments strongly supported the relative dependency of the affective valence representation on the spread-through-time account. We found that the semantic valence representations at the end of the maze were similar to those given while moving through the maze, as predicted by the spread-through-structure account. On the other hand, the affective valence representations of the different states were essentially the same across states. They did not resemble the pattern observed in the mid-maze affective ratings, as predicted by the spread-through-time account. These last findings align with the theoretical claim on the difference between the two formats of valence representation, affective and semantic, in their dependence on temporal proximity to experience. At the end of the maze, the participants encountered difficulty reconstructing their experienced mid-maze feelings in the various states of the task. Therefore, they could not give meaningful reports on the differences between the states based on their feelings. Nonetheless, they maintained their knowledge regarding the differences between the states, that is, the semantic valence representation was maintained.

Finally, using an analysis aimed at dissociating the unique contribution of the spread-through-time and spread-through-structure accounts, we found that both formats of valence representation were dependent on the structural distance between the outcome and the other states. We also observed a direction effect in the affective representation in Experiment 1, indicating that affective representations were also affected by a state-independent process and the temporal distance from the outcome. However, we did not observe this effect in the affective ratings in Experiment 2, where the negative outcome was a loss of points (and monetary bonus) and not having to wait extra time to move. Maybe the fact that the delay is a primary reinforcer, whereas the loss of points is a secondary reinforcer, contributed to this result (but see Delgado, Labouliere, & Phelps, 2006). Future studies can use a similar task to further explore this possibility.

5.2. Affective valence representations reflect "Here and Now" dynamics

Our study is part of a growing literature that seeks to dissociate affective and semantic representations of valence (for review, see Itkes & Kron, 2019). Such dissociation provides a new window into studying what is affect, what is semantic knowledge, the differences between the two, and their role in the affective response. Previous works found evidence for two formats of valence representation, affective and semantic, that differ in their sensitivity to the perceptual details of the stimulus (Olteanu, Salama, Kimchi, & Kron, 2022) and should be measured using different self-report instructions (Hamzani et al., 2019). Using dynamic learning paradigms, it was shown that affective valence representations are attenuated over multiple exposures (Itkes et al., 2017) and are updated faster, that is, are more localized in time to the experience only in case of variable outcomes (Heimer et al., 2023). Here we expand on these works and demonstrate that the affective representations of negative valence are also more localized in state-space, that is, show less spread to other states than the semantic representations, only in case of an adverse outcome.

Combining the results of Heimer et al. (2023) on the temporal dynamics of the affective and semantic valence and the current study results, the negative affective valence tends to have a more localized effect than the semantic valence. In other words, the negative affective valence gives more weight to events close in time or space, that is, focusing on the "Here and Now" than the semantic valence (see Olteanu et al., 2022 for a similar conclusion). The "Here and Now" focus of the affective valence can be explained by appraisal theories (for review, see Moors, Ellsworth, Scherer, & Frijda, 2013). According to these theories, the affective response is preceded by at least one evaluation stage, in which specific dimensions of an event are appraised (Olteanu, Golani, Eitam, & Kron, 2019). Specifically, one of the critical dimensions of the event's appraisal is its relevance to the observer's concerns. Higher relevance stimuli, like adverse stimuli that pose a threat to our well-being, are expected to yield a more intense affective response than lower relevance stimuli (e.g., N'diaye, Sander, & Vuilleumier, 2009; Olteanu et al., 2019; Sander, Grafman, & Zalla, 2003). Appraisal theories further claim that the relevance of a stimulus is influenced by its proximity in time and space to the observer, that is, its "Here and Now" characteristics. Therefore, the focus of the affective valence on the "Here and now" is a consequence of the higher relevance of proximal events to our concerns.

5.3. Implications and limitation

Our study provides novel direct evidence for valence spread across states. Instead of inferring about this spread from goal-directed reinforcement learning models (Juechems & Summerfield, 2019) or from behavioral and neuroimaging findings (e.g., Gershman et al., 2014; Keramati et al., 2016; Pauli et al., 2015), we explicitly asked the participants to rate the affective and semantic valence of the different states. The results provide robust and direct support for valence spread across states, both in the format of semantic knowledge and in the format of feelings. Notably, the direct evidence for the spread of negative feelings between states may potentially contribute to theoretical models of phenomena like avoidance, where individuals avoid situations, or states, that may lead to negative outcomes (for review, see Krypotos, Effting, Kindt, & Beckers, 2015). Understanding and characterizing the pattern of semantic and affective valence spread can explain why some individuals are more prone to avoidance behavior, and how structural and temporal relations between states support this behavior.

Our results also have implications for associative and reinforcement learning theories, by simultaneously considering the learning processes of the semantic and affective valence representations. Previous research focused on the mechanisms enabling people to learn the different states in the task and the expected transition probabilities between them, that is, learning factual, semantic knowledge about the states' valence (e.g., Peer, Brunec, Newcombe, & Epstein, 2021). However, our behavior is not solely guided by semantic knowledge but also by our emotional state (e.g., Lerner, Li, Valdesolo, & Kassam, 2015). Crucially, the two formats of valence representations, semantic and affective, might obey different spreading dynamics, leading to diverging guidance to planning and behavior, not taken into account by past work that focused only on the semantic format of valence representation.

Our results indeed show that the two formats' guidance to behavior diverges when the outcome is negative. Relying on the semantic format is expected to lead to the well-predicted, optimal sequential-state planning thoroughly described in reinforcement learning theories. However, relying on the affective format is expected to yield a sub-optimal planning behavior, biased toward the "Here and Now". This bias might have implications for studying why people repeat dangerous patterns of behavior, even after experiencing adverse outcomes in the past, such as gambling, irresponsible financial behavior, and health-threatening behaviors. As our study shows, this behavior may result from a bias toward the affective representations of the outcomes' negative valence, that is, an excessive focus on the "Here and Now". Notably, this "Here and Now" bias is combined with a relative neglection of the more spreadable representation of adverse valence across states, that is, the semantic representation. Specifically, these people may fail to recognize the expected state transitions, that is, the connections between their current behavior (smoking) and the adverse outcomes they will experience in the future (smoking-related illness). They may also be unaware of how close their current state (making a rush investment decision) is to an aversive state (bankruptcy). We suggest reframing these people's decisions in terms of a sequential-state planning problem, that is, identifying the state structure that these people are about to pass. Once framed as a sequential-state planning problem, we can develop interventions to better propagate the negative future valence to preceding states by emphasizing the state transitions.

One limitation of the current study is that our experimental design did not allow a reliable direct model fitting to the valence self-reports. The use of parameter estimation and model comparison analyses could have resulted in more insightful conclusions regarding the differences between the learning of affective and semantic valence representations of states. Unfortunately, our experimental design, which included passive movement across states and not active action selection, meant that we did not have extensive trial-by-trial choices to model, but only the relatively sparse self-reports of each format of valence (i.e., each format was reported in only 1/6 of the trials) (see Rutledge, Skandali, Dayan, & Dolan, 2014, for an active design with ratings). We chose this design as it allowed participants to experience all states, learn their transitions, and form a stable representation of their relationship to the outcome state, which an active task would not have allowed. In addition, we did not want to add too many self-reports, which may interfere with the spread-through-time account and make the experiment too long and tiresome. Future research could use a similar design with a better cover of self-reports to further explore the learning processes.

To conclude, the current study explored the spread of the two formats of valence representation, affective and semantic, across states. Our study provides the first evidence for a dissociative effect of feelings versus knowledge on sequential-state planning aimed to avoid adverse outcomes. Reliance on the negative feelings associated with each state, that is, the affective valence representation, will lead to a relatively localized spread of the outcome's negativity to the preceding states, primarily based on the temporal proximity of each state to the negative experience. On the other hand, reliance on the knowledge regarding the negativity of each state, that is, the semantic valence representation, will lead to a greater spread of the outcome's negativity to the preceding states, mainly based on each state's spatial proximity to the outcome. Future research should acknowledge the differences between the two formats of valence and their unique influence on behavior.

Funding

OH and UH were supported by the Israel Science Foundation (1532/20).

Links to pre-registrations, data, and code

Experiment 1 – Pre-registration can be viewed at https://aspredicted.org/HX8_488.

Data and code can be viewed at https://osf.io/6b2u9.

Experiment 2 – Pre-registration can be viewed at https://aspredicted.org/ZYN_TN4.

Data and code can be viewed at https://osf.io/6b2u9/.

CRediT authorship contribution statement

Orit Heimer: Conceptualization, Formal analysis, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. **Uri Hertz:** Conceptualization, Funding acquisition, Methodology, Resources, Supervision, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare no conflict of interest concerning the publication or the authorship of this article.

Data availability

Links to all data and code are detailed in the title page and the manuscript

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2023.105714.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19*(6), 716–723.

Barrett, L. F. (2006). Valence is a basic building block of emotional life. *Journal of Research in Personality, 40*(1), 35–55. https://doi.org/10.1016/j.jrp.2005.08.006

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron, 100*(2), 490–509. https://doi.org/10.1016/j.neuron.2018.10.002

Bellmund, J. L. S., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science, 362*(6415). https://doi.org/10.1126/science.aat6766

Bradley, M. M., Lang, P. J., & Cuthbert, B. N. (1993). Emotion, novelty, and the startle reflex: Habituation in humans. *Behavioral Neuroscience, 107*(6), 970.

Cardinal, R. N., Parkinson, J. A., Hall, J., & Everitt, B. J. (2002). Emotion and motivation: The role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience and Biobehavioral Reviews, 26*(3), 321–352. https://doi.org/10.1016/S0149-7634(02)00007-6

Carver, C. S. (2006). Approach, avoidance, and the self-regulation of affect and action. *Motivation and Emotion, 30*(2), 105–110. https://doi.org/10.1007/s11031-006-9044-7

Dael, N., Mortillaro, M., & Scherer, K. R. (2012). Emotion expression in body action and posture. *Emotion, 12*(5), 1085–1101. https://doi.org/10.1037/a0025737

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron, 69*(6), 1204–1215. http://www.ncbi.nlm.nih.gov/pubmed/21435563%0Ahttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3077926.

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation, 5*(4), 613–624. https://doi.org/10.1162/neco.1993.5.4.613

Delgado, M. R., Labouliere, C. D., & Phelps, E. A. (2006). Fear of losing money? Aversive conditioning with secondary reinforcers. *Social Cognitive and Affective Neuroscience, 1*(3), 250–259. https://doi.org/10.1093/scan/nsl025

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron, 80*(2), 312–325. http://www.ncbi.nlm.nih.gov/pubmed/24139036%0Ahttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC3807793.

Du, S., Tao, Y., & Martinez, A. M. (2014). Compound facial expressions of emotion. *Proceedings of the National Academy of Sciences of the United States of America, 111*(15). https://doi.org/10.1073/pnas.1322355111

Eldar, E., Rutledge, R. B., Dolan, R. J., & Niv, Y. (2016). Mood as representation of momentum. *Trends in Cognitive Sciences, 20*(1), 15–24. https://doi.org/10.1016/j.tics.2015.07.010

Epstein, R. A., Patai, E. Z., Julian, J. B., & Spiers, H. J. (2017). The cognitive map in humans: Spatial navigation and beyond. *Nature Neuroscience, 20*(11), 1504–1513. https://doi.org/10.1038/nn.4656

Erdfelder, E., Faul, F., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods, 41*(4), 1149–1160. https://doi.org/10.3758/BRM.41.4.1149

Garvert, M. M., Dolan, R. J., & Behrens, T. E. J. (2017). A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *ELife, 6*, 1–20. https://doi.org/10.7554/eLife.17086

Gazzaniga, M. S., Ivry, R. B., & Mangun, G. R. (2014). *Cognitive neuroscience: The biology of the mind* (4th ed.). W.W. Norton.

Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.

Gershman, S. J., Markman, A. B., & Otto, A. R. (2014). Retrospective revaluation in sequential decision making: A tale of two systems. *Journal of Experimental Psychology: General, 143*(1), 182–194. https://doi.org/10.1037/a0030844

Givon, E., Itzhak-Raz, A., Karmon-Presser, A., Danieli, G., & Meiran, N. (2019). How does the emotional experience evolve? Feeling generation as evidence accumlation. *Emotion, 20*(2), 271–285. https://doi.org/10.1037/emo0000537

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron, 66*(4), 585–595. https://doi.org/10.1016/j.neuron.2010.04.016

Hamzani, O., Mazar, T., Itkes, O., Petranker, R., & Kron, A. (2019). Semantic and affective representations of valence: Prediction of autonomic and facial responses from feelings-focused and knowledge-focused self-reports. *Emotion, 20*(3), 486–500. https://doi.org/10.1037/emo0000567

Heimer, O., Kron, A., & Hertz, U. (2023). Temporal dynamics of the semantic versus affective representations of valence during reversal learning. *Cognition, 236*(July). https://doi.org/10.1016/j.cognition.2023.105423

Hodes, R. L., Cook, E. W., III, & Lang, P. J. (1985). Individual differences in autonomic response: Conditioned association or conditioned fear? *Psychophysiology, 22*(5), 545–560.

Itkes, O., Kimchi, R., Haj-Ali, H., Shapiro, A., & Kron, A. (2017). Dissociating affective and semantic valence. *Journal of Experimental Psychology: General, 146*(7), 924–942. https://doi.org/10.1037/xge0000291

Itkes, O., & Kron, A. (2019). Affective and semantic representations of valence: A conceptual framework. *Emotion Review, 11*(4), 283–293. https://doi.org/10.1177/1754073919868759

Izquierdo, A., Brigman, J. L., Radke, A. K., Rudebeck, P. H., & Holmes, A. (2017). The neural basis of reversal learning: An updated perspective. *Neuroscience, 345*, 12–26. https://doi.org/10.1016/j.neuroscience.2016.03.021

Juechems, K., & Summerfield, C. (2019). Where Does Value Come From? *Trends in Cognitive Sciences, 23*(10), 836–850. https://doi.org/10.1016/j.tics.2019.07.012

Kazak, A. E. (2018). Editorial: Journal article reporting standards. *American Psychologist, 73*(1), 1–2. https://doi.org/10.1037/amp0000263

Keramati, M., Smittenaar, P., Dolan, R. J., & Dayan, P. (2016). Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum. *Proceedings of the National Academy of Sciences of the United States of America, 113*(45), 12868–12873. https://doi.org/10.1073/pnas.1609094113

Keren, H., Zheng, C., Jangraw, D. C., Chang, K., Vitale, A., Rutledge, R. B., … Stringaris, A. (2021). The temporal representation of experience in subjective mood. *ELife, 10*, 1–24. https://doi.org/10.7554/eLife.62051

Klein-Flügge, M. C., Wittmann, M. K., Shpektor, A., Jensen, D. E. A., & Rushworth, M. F. S. (2019). Multiple associative structures created by reinforcement and incidental statistical learning mechanisms. *Nature Communications, 10*(1), 1–15. https://doi.org/10.1038/s41467-019-12557-z

Krypotos, A. M., Effting, M., Kindt, M., & Beckers, T. (2015). Avoidance learning: A review of theoretical models and recent developments. *Frontiers in Behavioral Neuroscience, 9*. https://doi.org/10.3389/fnbeh.2015.00189

Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. *Annual Review of Psychology, 66*, 799–823. https://doi.org/10.1146/annurev-psych-010213-115043

Momennejad, I. (2020). Learning structures: Predictive representations, replay, and generalization. *Current Opinion in Behavioral Sciences, 32*, 155–166. https://doi.org/10.1016/j.cobeha.2020.02.017

Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour, 1*(9), 680–692. https://doi.org/10.1038/s41562-017-0180-8

Moors, A., Ellsworth, P. C., Scherer, K. R., & Frijda, N. H. (2013). Appraisal theories of emotion: State of the art and future development. *Emotion Review, 5*(2), 119–124.

N'diaye, K., Sander, D., & Vuilleumier, P. (2009). Self-relevance processing in the human amygdala: gaze direction, facial expression, and emotion intensity. *Emotion, 9*(6), 798.

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology, 53*(3), 139–154. https://doi.org/10.1016/j.jmp.2008.12.005

Öhman, A., Hamm, A., & Hugdahl, K. (2000). Cognition and the autonomic nervous system: Orienting, anticipation, and conditioning. In J. T. Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds.), *Handbook of psychophysiology* (pp. 533–575). Cambridge University Press.

O'Donoghue, B., Osband, I., Munos, R., & Mnih, V. (2018). *The uncertainty Bellman equation and exploration, 9*, 6154–6173.

Olteanu, L., Golani, S., Eitam, B., & Kron, A. (2019). The effect of relevance appraisal on the emotional response. *Emotion, 19*(4), 715–725.

Olteanu, L., Salama, K., Kimchi, R., & Kron, A. (2022). A divergent effect of stimulus perceptual details on affective and semantic representations of valence. *Emotion.*. https://doi.org/10.1037/emo0001148

Osgood, C. E. (1952). The nature and measurement of meaning. *Psychological Bulletin, 49* (3), 197.

Pauli, W. M., Gentile, G., Collette, S., Tyszka, J. M., & O'Doherty, J. P. (2019). Evidence for model-based encoding of Pavlovian contingencies in the human brain. *Nature Communications, 10*(1), 1–11. https://doi.org/10.1038/s41467-019-08922-7

Pauli, W. M., Larsen, T., Collette, S., Tyszka, J. M., Seymour, B., & O'doherty, J. P. (2015). Distinct contributions of ventromedial and dorsolateral subregions of the human substantia nigra to appetitive and aversive learning. *Journal of Neuroscience, 35*(42), 14220–14233. https://doi.org/10.1523/JNEUROSCI.2277-15.2015

Pavlov, P. I. (2010). Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. *Annals of Neurosciences, 17*(3), 136. https://doi.org/10.5214/ans.0972-7531.1017309

Peer, M., Brunec, I. K., Newcombe, N. S., & Epstein, R. A. (2021). Structuring knowledge with cognitive maps and cognitive graphs. *Trends in Cognitive Sciences, 25*(1), 37–54. https://doi.org/10.1016/j.tics.2020.10.004

Prével, A., & Krebs, R. M. (2021). Higher-order conditioning with simultaneous and backward conditioned stimulus: Implications for models of Pavlovian conditioning. *Frontiers in Behavioral Neuroscience, 15*(November), 1–7. https://doi.org/10.3389/fnbeh.2021.749517

R Core Team. (2022). *R: A language and environment for statistical computing.* Vienna, Austria: R Foundation for Statistical Computing.

Rmus, M., Ritz, H., Hunter, L. E., Bornstein, A. M., & Shenhav, A. (2022). Humans can navigate complex graph structures acquired during latent learning. *Cognition, 225* (July 2021), Article 105103. https://doi.org/10.1016/j.cognition.2022.105103

Robinson, M. D., & Clore, G. L. (2002). Episodic and semantic knowledge in emotional self-report: Evidence for two judgment processes. *Journal of Personality and Social Psychology, 83*(1), 198–215. https://doi.org/10.1037/0022-3514.83.1.198

Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2017). Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Computational Biology, 13*(9), 1–35. https://doi.org/10.1371/journal.pcbi.1005768

Russell, J. A. (1983). Pancultural aspects of the human conceptual organization of emotions. *Journal of Personality and Social Psychology, 45*(6), 1281–1288. https://doi.org/10.1037/0022-3514.45.6.1281

Rutledge, R. B., Skandali, N., Dayan, P., & Dolan, R. J. (2014). A computational and neural model of momentary subjective well-being. *Proceedings of the National Academy of Sciences of the United States of America, 111*(33), 12252–12257. https://doi.org/10.1073/pnas.1407535111

Sander, D., Grafman, J., & Zalla, T. (2003). The human amygdala: an evolved system for relevance detection. *Reviews in the Neurosciences, 14*(4), 303–316.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 461–464.

Seymour, B., O'Doherty, J. P., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., … Frackowiak, R. S. (2004). Temporal difference models describe higher-order learning in humans. *Nature, 429*(6992), 664–667. https://doi.org/10.1038/nature02581

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis.* Cambridge, Massachusetts: B.F. Skinner Foundation.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). The MIT Press.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review, 55*(4), 189.

Wang, L., Li, X., Pi, Z., Xiang, S., Yao, X., & Qi, S. (2021). Spatiotemporal dynamics of affective and semantic valence among women. *Frontiers in Human Neuroscience, 15*, Article 602192. https://doi.org/10.3389/fnhum.2021.602192