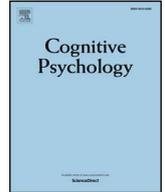


Contents lists available at [ScienceDirect](#)

Cognitive Psychology

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

The detour problem in a stochastic environment: Tolman revisited

Pegah Fakhari*, Arash Khodadadi, Jerome R. Busemeyer

Indiana University, Department of Psychological and Brain Sciences, Bloomington, IN, United States



A B S T R A C T

We designed a grid world task to study human planning and re-planning behavior in an unknown stochastic environment. In our grid world, participants were asked to travel from a random starting point to a random goal position while maximizing their reward. Because they were not familiar with the environment, they needed to learn its characteristics from experience to plan optimally. Later in the task, we randomly blocked the optimal path to investigate whether and how people adjust their original plans to find a detour. To this end, we developed and compared 12 different models. These models were different on how they learned and represented the environment and how they planned to reach the goal. The majority of our participants were able to plan optimally. We also showed that people were capable of revising their plans when an unexpected event occurred. The result from the model comparison showed that the model-based reinforcement learning approach provided the best account for the data and outperformed heuristics in explaining the behavioral data in the re-planning trials.

1. Introduction

Humans deal with planning problems in their everyday situations. One of the very familiar situations is to navigate from one place to another in a neighborhood or city. In this scenario, usually there is more than one path to choose, and, depending on the goal, one might select the shortest path, the city roads, a bypass/highway outside the traffic area or a path with the minimal traffic lights. Once he chooses the highway, he still needs to decide whether to take the toll line and/or where to exit. On the other hand, if he chooses the city roads, he would need to decide which intersection to go, to use the main street or to use shortcuts, etc. In other words, after selecting a general path (plan), there are still small paths decisions.

This is an example of a more general problem in which one needs to optimally plan a sequence of interdependent choices to accomplish a goal. In some cases, the shortest path is the optimal path and in others his goal might be to avoid the traffic at all costs.

In multistage decision making, unlike isolated choices, the focus is on how people analyze the interrelated choices to make an *optimal* sequence of choices, [Hotaling, Fakhari, and Busemeyer \(2015\)](#) and [Gonzalez, Fakhari, and Busemeyer \(2017\)](#). Usually, sequential decisions are represented in a decision tree in which the result of an action at one stage (e.g. a decision node) will be fed into the next stage, which might be another decision node or possibly an output node. Consider a decision tree with two decision nodes (yellow circles) and five possible (green) paths that is represented in the upper panel of [Fig. 1](#). Given the starting position, **S** and the goal, **G**, the paths are Path B, Path A1A2, Path A1C2, Path C1C2 and Path C1A2. The black dashed line in the grid separates path B and path A1A2.

In order to make optimal choices, one should know the actual output of each decision node and the uncertainty of each transition. For instance, although the number of steps (or actions) between the starting position and the next decision node (or the goal) is not depicted in [Fig. 1](#), based on the expected losses (*EL*), we know that path B is the best path to go to the goal position. In an experience-

* Corresponding author at: Indiana University, Department of Psychological and Brain Sciences, 1101 E. 10th Street, 47405-7007 Bloomington, IN, United States.
E-mail address: pegah.fakhari@gmail.com (P. Fakhari).

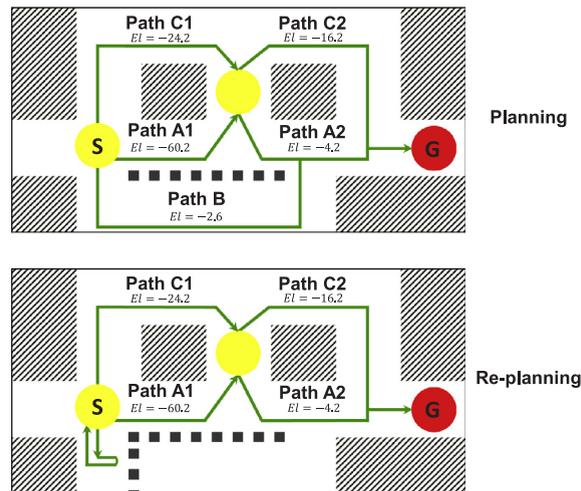


Fig. 1. General grid world used in our experiments. The decision nodes are represented by yellow circles and the goal is depicted by red circle. The green paths are available to participants. Note that participants cannot see the obstacles depicted by shadow areas. The black dashed line in the grid separates path B and path A1A2. Top: In the planning trials, for this current start and goal positions, 5 paths are available to participants and they need to find the optimal path. Down: In the re-planning trials, the optimal path, path B is blocked and participants need to find the detour. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

based decision tree, this knowledge is established from (an individual's) experience, Daw, Gershman, Seymour, Dayan, and Dolan (2011), Momennejad et al. (2016), Huys et al. (2012), Dezfouli and Balleine (2012) and Keramati, Smittenaar, Dolan, and Dayan (2016) while in a description-based version, it is provided by the experimenter (available during the task), Hotaling and Busemeyer (2012), Hey and Knoll (2011), Drner and Schaub (1994), Johnson and Busemeyer (2001) and Johnson and Busemeyer (2005). However, having this knowledge cannot guarantee the optimal behavior, Hey (2005), Sims, Neth, Jacobs, and Gray (2013), Yechiam, Erev, Yehene, and Gopher (2003), Hotaling and Busemeyer (2012), Huys et al. (2015), Momennejad et al. (2016), Huys et al. (2015), Botvinick, Niv, and Barto (2009) and Keramati, Dezfouli, and Piray (2011).

1.1. Examining the optimality of re-planning in sequential decision making tasks

In many of the previous studies on planning in decision trees, the environment is not dynamically changing, Nassar, Wilson, Heasly, and Gold (2010). Once the participant learns about the risky sequential decision environment, she can *optimally* plan her actions and does not need to update her knowledge later. But in real life, our environment is always changing and unexpected events happen. Usually, we have two strategies to deal with these situations: reevaluate our plans with the new information (which is also known as re-planning) or ignore the new information and stick to our original plan.

In this article, we extend previous planning experimental designs to situations in which the participants experience random changes in the environment and need to modify their original plans to get to the goal position. We test learning to plan and re-planning in one unique framework: a 4 by 7 grid world with stochastic losses. In the learning phase of our experiments, we look into the planning behavior and how participants can learn to find the optimal sequence of choices. Then, in the test phase, we block the optimal path randomly and ask our participants to find a detour path (re-planning behavior) based on what they have learned during the learning phase as illustrated in the bottom panel of Fig. 1. On 30% of the trials, path B (the *optimal* path) is randomly blocked and becomes unavailable¹. The vertical black dashed line that is added to the grid shows the wall that makes path B unavailable. In this situation, path C1A2 is the best path to choose (optimal re-planning behavior).

It is important to emphasize that the starting and goal positions are not fixed in our design. We randomize these pairs for three reasons: (1) to make sure that our participants have a fair exposure to different aspects of the grid world environment; (2) a random starting point and goal makes it more challenging to discriminate different models and their predictions; (3) to examine human planning behavior from different arbitrary decision nodes located in different layers of a decision tree.

In Section 4.1, we discuss how we fit different models and compare the results in detail. However, it is important to highlight that our design includes a generalization test that fits model parameters to the planning phase (with no blocks in the optimal path), and then subsequently uses these same parameters to predict re-planning in a generalization test when blocks are introduced, Busemeyer and Wang (2000). This provides a very strong test of the competing models that vary in number of parameter and model complexity. In other words, our model comparison is not restricted to how different models can learn the model of the environment and whether or not they can predict planning (which they have been trained for), but involves a more rigorous test on how they can perform in an unexperienced environment.

¹ Participants do not experience the blockage unless they check the optimal path.

1.2. Computational models for learning sequential decision making from experience

In most everyday situations, it is almost impossible to have access to all of the information required to make sequential decisions from the start, and much of this information must be learned from experience. While it is necessary to take advantage of previous experiences to make better choices in the future, it is also essential to evaluate actions and assign (proper) credit to the earlier choices as well as to the later actions in a sequence (temporal credit assignment). These two criteria are the heart of many reinforcement learning (RL) and unsupervised learning algorithms in computer science (Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998; Walsh & Anderson, 2014).

In RL, the goal is to choose appropriate action(s) which maximize the expected sum of future rewards. In this set of problems, the agent initially does not know the correct correspondence between states and actions and receives feedback following an action or a sequence of actions. This framework can address learning mechanisms, ranging from the very basic stimulus response (habitual behavior) to the more complicated goal-directed behavior (including, but not limited to, planning). Whether people conduct themselves in accordance with the RL predictions in sequential choices is still under investigation and debate (for a detailed review, see Walsh & Anderson, 2014). For instance, Huys et al. (2012), asked their participants to plan a sequence of 2- choices by traveling through 6 states (represented as boxes on the screen) to maximize their score. Different RL models were fitted to explain the underlying mechanism. They found that people extensively pruned sub-trees with large losses in order to reduce the size of the decision problem into a computationally manageable problem (Huys et al., 2012)².

In general, there are two classes of RL models that can provide a solution to our sequential decision problem. The first class finds the optimal policy by learning the model of the environment and is called the model-based approach. The second class, the model-free approach, is able to maximize the expected sum of future rewards *without* knowing the characteristics of the environment, Sutton and Barto (1998) and Daw, Niv, and Dayan (2005). Therefore, the second approach is computationally efficient. But as we explain later, it only incorporates one-step rewards (plus the expected future reward from the next step) from a particular state into planning which limits its ability to plan when the start and goal positions change.

There is a third RL algorithm, called successor representation, that learns a rough representation of the environment by storing the expected future visits of each state, Gershman, Moore, Todd, Norman, and Sederberg (2012) and Dayan (1993). It is computationally less expensive than the model-based RL and can easily explain the planning behavior in our experiments. It has limited ability to re-plan in a changing environment (it performs differently with regard to reward and transition devaluation).

Last but not least, we investigate how heuristics and simple strategies perform in our risky sequential choice environment. In Section 3.2, we show the predictions of 3 different heuristic based models. Our design was able to show that these heuristic models fail, both quantitatively and qualitatively.

2. Material and methods

2.1. Overview

Learning the characteristics of the environment to plan optimally (and possibly re-plan in case of the unexpected changes) was first examined by Tolman and Honzik with rats, Tolman (1948) and Tolman and Honzik (1930). In this task, which is known as the detour experiment, first, rats were exposed to three different paths with different lengths to a goal/food location similar to Fig. 2, a. Then in the test phase, they blocked the shortest path and checked the rats' choice in finding the second shortest path (Voicu & Schmajuk, 2002). Tolman (1948) and Tolman and Honzik (1930) summarized that acquiring skills like turning right at specific positions to do a particular task (e.g. reaching a goal/food) could be explained by a stimulus-response learning system. However, only goal directed behavior could explain the rapid learning curve and fewer errors in rats' behavioral data in finding the detour paths, Tolman, Ritchie, and Kalish (1946).

Inspired by these findings, one of our very basic questions in this study is whether people can plan in a stochastic environment and if so how far they can modify their plans in order to maximize their earnings. There are two phases in our experiment: a training phase which allows participants to explore and learn about the grid world (without monetary reward) and a test phase that goes beyond optimal planning and requires finding the second best path to maximize their money. One of the main differences between our design and Tolman's detour problem is that in Tolman (1948), the environment is deterministic but in our grid world experiment, there are multiple cells with stochastic rewards (2 cells in experiment 1 and 5 cells in experiments 2 and 3). Therefore, the *best* path is not simply the *shortest* path but is the *optimal* path which requires planning. Recognizing the difference between the shortest and optimal paths plays a key role on the (amount of) money that participants can earn.

We should also note that in our design, we used random pairs in each trial to make sure that participants explore the entire grid world. However, in the original detour problem, Tolman (1948) exposed rats to different paths using fixed starting and goal positions.

Table 1 highlights the differences and similarities among the three experiments. It is important to emphasize that participants can only see a plain grid on the screen along with their current position (yellow circle) and their destination (red circle, Fig. 2) with no sign/cue of the obstacles or the stochastic losses. They do not have access to papers nor calculators/cellphones to do any computations or to take notes. In the instruction, they are told to consider a scenario that they move to a new city (a 4 by 7 grid) and need to get from one place to another (determined by the yellow and red circles). In the test phase, participants are told that a random

² They used two different discounting parameters to differentiate the greatest loss from the other losses.

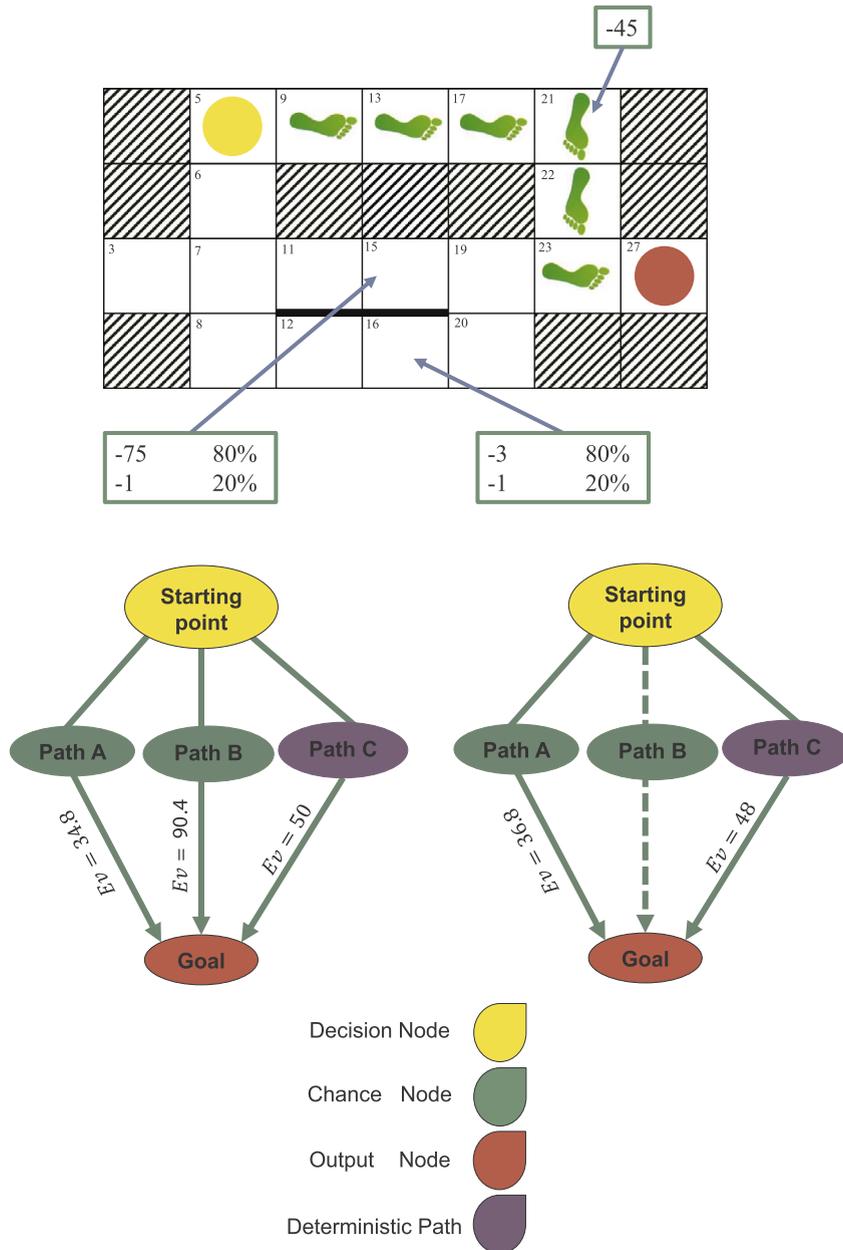


Fig. 2. Experiment 1. Up: A 4 by 7 grid world experiment. Green footsteps indicate one possible path from starting point to the goal (path C). Expected values (Ev) are calculated based on the current starting point. Down: Corresponding decision tree for planning (left) and re-planning (right). There are 3 possible paths between the starting position and the goal. Expected value in the re-planning phase is different from planning phase and is calculated from cell 7. Cell 7 is the position that participants choose their detour path in the re-planning trials. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

accident might happen in one of the possible routes and block that path. If they see the accident they need to find a detour path.

In the following section, first we summarize the behavioral results for these three experiments and then we try to characterize participants' strategies in planning and re-planning in the pretest and the test blocks respectively. Presumably any candidate model must incorporate environmental changes in its planning approach in order to make sure that a new information is propagated to the action selection module promptly.

2.2. Experiment 1

2.2.1. Participants

Twenty healthy (11 females) participants performed experiment 1 for payment. Minimum payment was "\$9 and participants

Table 1
Summary of experiments.

Design	Experiment 1	Experiment 2	Experiment 3
Number of learning blocks	6	6	6
Number of test blocks	2	2	3
Number of pretest blocks	1	1	1
Number of probabilistic losses	2	5	5
Fixed G and S in the test blocks	Yes	Yes	No
Fixed G and S in the learning blocks	No	No	No
Fixed G and S in the pretest block	No	No	No
Cells with stochastic losses	15, 16	9, 11, 16, 19, 17	9, 11, 16, 19, 17
Number of re-planning trials in the test trials	13 out of 40	13 out of 40	20 out of 60

Experiment 1 has two stochastic losses with one sure loss at cell 21 with the loss of -45 . In experiment 2 and 3, there are 5 stochastic losses (similar environment). The rate of re-planning trials in the test blocks is fixed in all three experiments (33%). Unlike experiments 1 and 2, the pairs in the test blocks of the experiment 3 are random. In all three experiments, there is one pretest block and six learning blocks. G and S stand for Goal and Starting positions.

could earn up to "\$"16 based on their performance during the task. Informed consent was collected from all participants and the study was approved by Indiana University Institutional Review Boards.

2.2.2. Task

Participants learned to *correctly* travel through a grid world in as few moves as possible while maximizing their reward. In experiment 1, there were 9 blocks and each has 20 trials. In each trial participants viewed a typical 4 by 7 grid world example and were asked to start from a *random* starting point, represented by a yellow circle, and reach to a random goal position (in red circle) using arrow keys (up, right, down and left), as shown Fig. 2, a. For instance in Fig. 2, the start and goal points are arbitrarily located at cell 5 (start) and cell 27 (goal).

Except for the goal point, participants received either a fixed or stochastic punishment for any movement in the grid world. If they caught the goal, they earned + 100 points (later exchanged for money at an exchange rate equal to 0.01). The fixed *regular* punishment was delivered in transition to most of the cells and it was equal to -1 . At cell 21, there was another deterministic (but not regular) punishment. If participants entered this cell, they received -45 . Finally, cells 15 and 16 had stochastic punishments. If transmitted to these cells, participants received a different punishment (-75 , -3 respectively) with probability of 0.8 and a regular punishment (-1) with probability of 0.2. It has been suggested by Hertwig, Barron, Weber, and Erev (2004) that the existence of rare events, specially outcomes that occur with probability less than 0.15, can lead to non-optimal behavior. Using this current settings for the loss structure, we have one optimal, one sub-optimal and one non-optimal path (three distinct paths). For instance, starting at cell 3, path A through 3, 7, 11, 15, 19, 23, 27 is the non-optimal path. Path B through 3, 7, 8, 12, 16, 20, 19, 23, 27 and path C through 3, 7, 6, 5, 9, 13, 17, 21, 22, 23, 27 are the optimal and sub-optimal paths. Note that the probabilities, the amount and the positions of the punishments were not known to participants.

In each attempt to catch the goal, a participant had 15 moves. There were ten hidden fixed obstacles which prevented participants to go to the selected direction while traveling in the grid world (e.g. selecting up at cell 11 was blocked). Hidden obstacles are shown with upward diagonal texture in Fig. 2, a. Each time a participant hit an obstacle, she received a punishment of -1 and used up one of her 15 moves. The first six blocks were designed to help participants learn about the obstacles, punishments and the details of the grid world (learning phase). During this phase participants were not getting paid.

The 7th block was designed to check optimal planning. Similar to previous blocks, it consisted of 20 random pairs of starting and goal points but the difference was that participants got paid based on their score (we called it the pretest block). Finally in the last two blocks (8th and 9th), optimal path, Path B, was blocked in one-third of the trials. The random blockage was to further examine optimal planning and also re-planning (finding the second optimal path) in our dynamic environment. The starting and goal positions were fixed at cell 3 and cell 27 in 8th and 9th blocks (the test blocks).

In Fig. 2, bottom panel, we present two decision trees. One on the left, all three paths are available (planning phase) and their expected values were calculated based on the current positions of the start (cell 5) and goal (cell 27). For instance if one chooses path A (through cell 6, 7, 11, 15, 19, 23, 27), the expected value is 34.8, if he chooses path B (through 6, 7, 8, 12, 16, 20, 19, 23, 27), the expected value is 90.4 and finally for path C (through 9, 13, 17, 21, 22, 23, 27), the exact value is 50 since all punishments are deterministic. The decision tree on the right is related to re-planning phase where path B is not available (dashed line). The expected values are calculated from cell 7 because this cell is the starting point to find the detour path after the participant experiences the random blockage in the optimal path (path B) at cell 8. Then he goes up to cell 7 to choose one of the available paths (path A and path C). If he chooses path A, then the expected value is 36.8 while choosing path C leads to the value of 48.

Note that in experiment 1, the corresponding decision trees in planning are directly related to the grid world on the top panel of Fig. 2 and where the starting point and goal are located. For another pair, the expected values of the paths are different. In the test blocks, though, the starting and goal positions are fixed and the corresponding decision tree (in re-planning) is the same.

The hallmark of our design was that the grid environment was not deterministic. Thus optimal planning was not *necessarily* reduced to finding the shortest path but to select the path with the maximum expected value through experience. For instance, in Fig. 2, top panel, if the starting point was at cell 3 and the goal was at cell 27, the shortest path, path A, had the lowest expected value compared to path B and path C. In the re-planning trials when path B was blocked, the second optimal path was path C which was

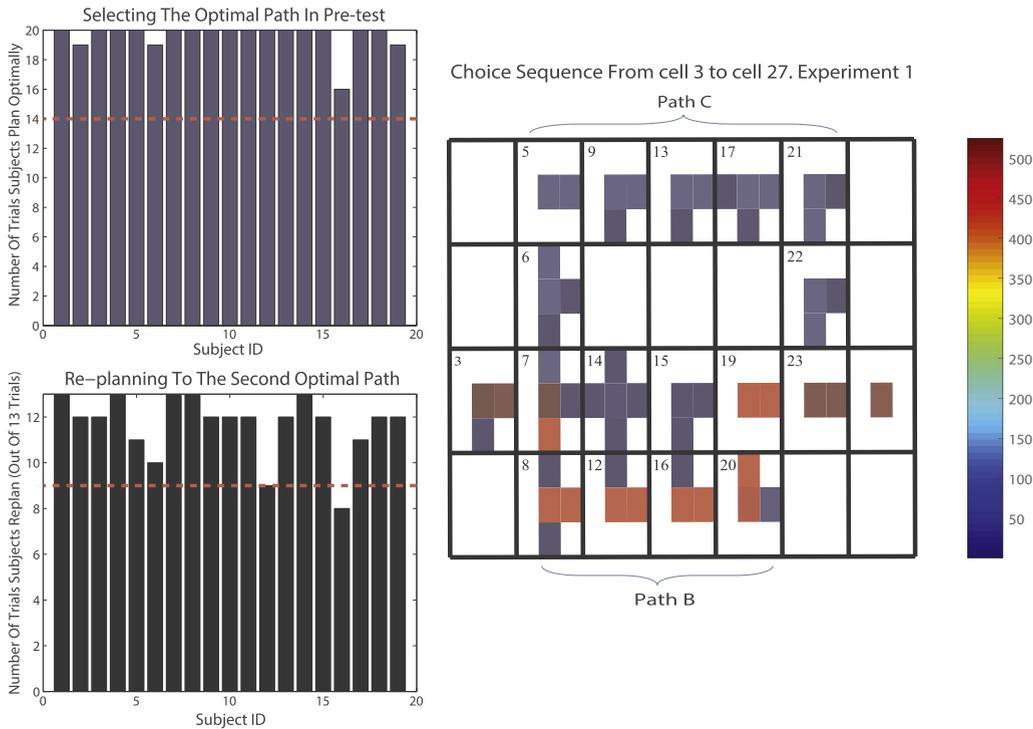


Fig. 3. Experiment 1. Up left: All participants (total 19), chose the optimal path ($\alpha = 0.05$) in the pretest block with 20 different pairs of goal and starting points. Down left: 17 participants re-planned to path C (the second optimal path) when they encountered with the blocked path B ($\alpha = 0.05$). The optimal path was blocked in 13 trials (out of 40 trials) in the test blocks. Right: Participants' choice in planning trials in the test blocks in experiment 1. The heatmap shows the frequency of choices in the planning trials when path B was not blocked. Path B was the optimal path and selected frequently (depicted in red). The related cells were 8, 12, 16, 20, 19 and 23. The white cells with no cell number are obstacles. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

longer than path A.

2.2.3. Results

All three experiments that we designed had 6 blocks of learning without payment (training phase). This allowed participants to explore the grid world without any concerns about their score. So if their scores became negative - as long as they were in the first 6 blocks - they were not punished. At the beginning of these blocks, participants did more exploring and less planning but as they moved forward, they should be acquainted with the environment and able to find the optimal path. In order to examine this, we fixed the starting position at cell 3 and the goal at cell 27 for the last five trials of the 6th block. With this design, we were able to recognize preliminary evidence of learning in the final blocks of the learning phase, specifically in block 6. A participant who learned the grid world correctly, should be able to select path B (the optimal path) in these final trials otherwise we excluded his data. With this criterion, one of the participants' data was not included in our analysis.

Analyzing the pretest block's data (seventh block), all 19 participants learned to choose the optimal path to reach the goal, Fig. 3, up left. We used the binomial distribution to compute the minimum number of trials in which the participant must plan optimally to be able to reject the null hypothesis that the optimal path was selected randomly, ($\alpha = 0.05$). Participants who chose the optimal path in only 14 trials or less (out of 20 trials) did not pass this test and failed to perform optimally. The red dashed line in top left panel of Fig. 3 shows this minimum.

In the test phase, path B was blocked in 13 trials (out of 40 trials). The minimum number of trials required for an optimal performance which could reject the null hypothesis was 10. The red line showed this threshold. Participants who selected the (optimal) detour path less than or equal to 9 trials out of 13 trials were not optimal in re-planning trials. Only two of the participants failed to re-plan correctly to the second optimal path (path C), Fig. 3, down left. Mixed-effect regression analysis with block number as regressor and participants as the random effect, showed that the average reward learning curve (as a function of block) was significant ($t(113) = 12.138, p < 0.0001$). Fig. 3, right, shows that participants (on average) took path B, the optimal path, in planning from cell 3 to cell 27 more frequently than path A or path C. Note that path C was the second optimal path if path B was blocked. In the heatmap, the red color shows the more frequent directions or state occupancies. For each cell, the number of times participants visited a cell and the direction they took are counted. In the test blocks of experiment 1, there are 27 trials that re-planning did not occur and the starting and goal positions are fixed at 3 and 27. Ideally, if all participants choose the optimal path (path B) in all trials, cell 7,8,12,16,20,19 and 23 should be visited $19 \times 27 = 513$ times while path A and path C should not be selected. As the heatmap shows, there are few trials (less than 50 trials in total) that path B is not selected (indicated by dark blue cells in path A and path C).

2.3. Experiment 2

2.3.1. Participants

Thirty-six healthy participants (18 males), performed experiment 2 for payment of "\$9 to "\$16 based on their performance in the task. All participants signed the informed consent forms. Indiana University Institutional Review Boards has approved the study.

2.3.2. Task

In the first experiment the best and worst paths (path B and path A) included a cell with a stochastic reward. In order to select the second optimal path in the re-planning trials, participants needed to compare a deterministic path (path C) with a stochastic path (path A). This might raise this question that our findings in the previous section only reflected participants' preference for a sure thing over a gamble. Therefore, we changed the environment in experiment 2 and designed a more complicated environment as depicted in Fig. 4. Similar to the previous experiment most cells had a regular punishment of -1 , but there were 5 cells with stochastic reward. These cells were 9,11,16,17,19 and with probability of 0.8 subjects received -20 , -75 , -3 , -30 and -5 respectively. The number of available paths was increased from 3 to 5 and they were: (e.g. starting at cell 7) Path A1A2 through 7,11,15,19,23,27; Path A1C2 through 7,11,15,14,13,17,21,22,23,27; Path B through 7,8,12,16,20,19,23,27; path C1C2 through 7,6,5,9,13,17,21,22,23,27 and path C1A2 through 7,6,5,9,13,14,15,19,23,27. Path B which includes cell 16 with expected loss of -2.6 is the optimal path. Choosing any path that contains cell 11 with expected loss of -60.2 has the worst effect on participants' monetary reward (path A1A2 and A1C2). The losses at cell 9, cell 17 and cell 19 were adjusted to create two sub-optimal paths: path C1C2 and C1A2. Based on the expected losses of these five critical cells, these paths could be prioritized as follows: path B, path C1C2, path C1A2, path A1A2 and path A1C2.

In this design the corresponding decision tree had two decision nodes and five chance nodes, the bottom panel of Fig. 4. Similar to experiment 1, in the test phase, the starting point and the goal position were fixed at 3 and 27 respectively. On 33% of the trials in the test phase, when the optimal path is randomly blocked, path C1A2 was the optimal path among the remaining paths. Fig. 4, b shows the corresponding decision tree.

2.3.3. Results

To analyze participants' learning curve, we fitted a mixed-effect regression to average reward that each participant earned during the first 7 blocks (the learning and the pretest blocks). We found that the average reward learning curve as a function of block was significant ($t(215) = 9.441, p < 0.0001$). Note that the number of available paths increased but similar to experiment 1, first, we checked whether the selected path was an optimal path or not. This is a simple straight forward test to check the optimal behavior in our experiment. As we did in experiment 1, we computed the minimum number of trials in which the participant needed to choose the optimal path to reject the null hypothesis. Since the number of trials in the pretest block and the test blocks has not changed, these thresholds were similar to experiment 1 which was 14 in the pretest block and 10 in the test blocks. Using this analysis, 30 out of 36 subjects found the optimal path in the pretest block ($\alpha = 0.05$). When it came to find the detour option, 27 participants found the second optimal path correctly; 7 chose the third optimal path and two did not re-plan, as shown in Fig. 5, panel on the left. In our analysis, the optimal strategy was to first check the original optimal path (path B) and if it was blocked then switch to the second optimal path (path C1A2). Fig. 5, right panel, captured this result from participants' choice. The heatmap shows the number of times participants visited each cell during trials when optimal path is not blocked in the test blocks (27 planning trials). Since we have 36 participants, the maximum number of times that a cell could be occupied is $27 \times 36 = 972$ (represented in dark red). In few trials, non-optimal paths (path A1, path C1 and path C2) are selected (shaded in dark blue).

2.4. Experiment 3

In the instructions of experiment 1 and 2, participants were told that only one new obstacle would be added to the environment and it would block one path. Also in the test phase of the first and second experiments, the starting point was at 3 and the goal point was at 27 in all trials. But having a fixed pair in the test phase might indirectly lead participants to perform optimally in the re-planning trials. For instance, learning to go to cell 27 from cell 3 was easier than learning to go to cell 13 from cell 5 in pair (5, 13). In the latter pair, participants not only needed to know the stochastic payoffs but also paths' length to choose the optimal path. One should know that path 5, 9, 13 had only 2 steps with expected loss of -17.2 and path B had 10 steps with two smaller losses (the total expected loss was -13.8). However in pair (3, 27), comparing the reward structures was enough to select the optimal path because path A was the shortest with minimum expected value and paths C1A2, C1C2, A1C2, A1A2 had smaller expected value and were longer than path B. In short, learning to choose the optimal path in some pairs was harder; consequently re-planning in these pairs when the environment changed (e.g. blockage in the optimal path) was more challenging.

Thus in experiment 3, we changed the test phase (8th and 9th block) to investigate this new question: how people modify their decisions when the change in the environment is less informative. Specifically, instead of having a fixed pair of starting point and goal position in the 8th and 9th blocks, we randomized these pairs (similar to what participants got used to do in the first 7 blocks).

2.4.1. Participants

Among thirty-two healthy subjects who participated in experiment 3, there were 19 males. Experiment 3 has three blocks of test and thus the upper limit of payment is increased to "\$20 but still depends on their performance in the task. All participants signed the informed consent forms. Indiana University Institutional Review Boards has approved the study.

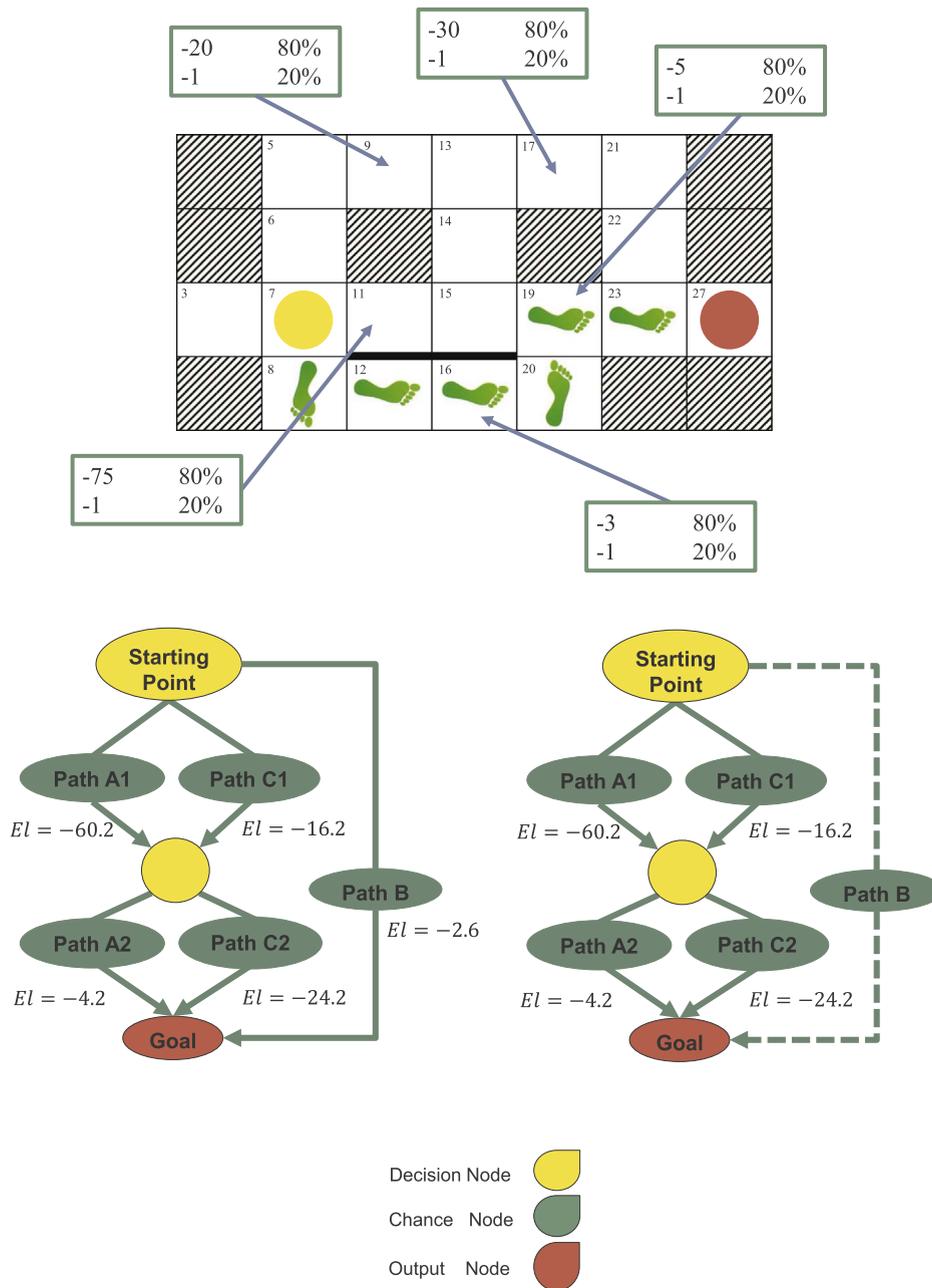


Fig. 4. Experiment 2. Up: A 4 by 7 grid world experiment. The starting position is the yellow circle (cell 7) and the goal is located at the red circle (cell 27). Green footsteps indicate the optimal path from the starting point to the goal in the environment (which is path B). There are 5 possible paths between the starting point and goal position. Path A1A2 through 7, 11, 15, 19, 23, 27; Path A1C2 through 7, 11, 15, 14, 13, 17, 21, 22, 23, 27; Path B through 7, 8, 12, 16, 20, 19, 23, 27; path C1C2 through 7, 6, 5, 9, 13, 17, 21, 22, 23, 27 and path C1A2 through 7, 6, 5, 9, 13, 14, 15, 19, 23, 27. Down: Corresponding decision tree for planning (left) and re-planning (right) with the exact 5 possible paths between the start and goal positions. In the re-planning trials (on the right), path B is blocked in 30% of the trials (dashed line). Expected losses are calculated based on the current starting point (cell 7). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.4.2. Task

In sum, we kept the same design as previous experiments but participants saw random pairs in the test phase. In order to keep the rate of changes in the environment (in the test phase) to 33% (as it was in experiment 1 and 2), we increased the number of re-planning blocks to three blocks (8th, 9th and 10th) and therefore in 20 out of 60 trials the optimal path was blocked.

2.4.3. Results

Recall that in this experiment, in the test phase, the goal and starting point were not fixed, and thus for each pair we determined

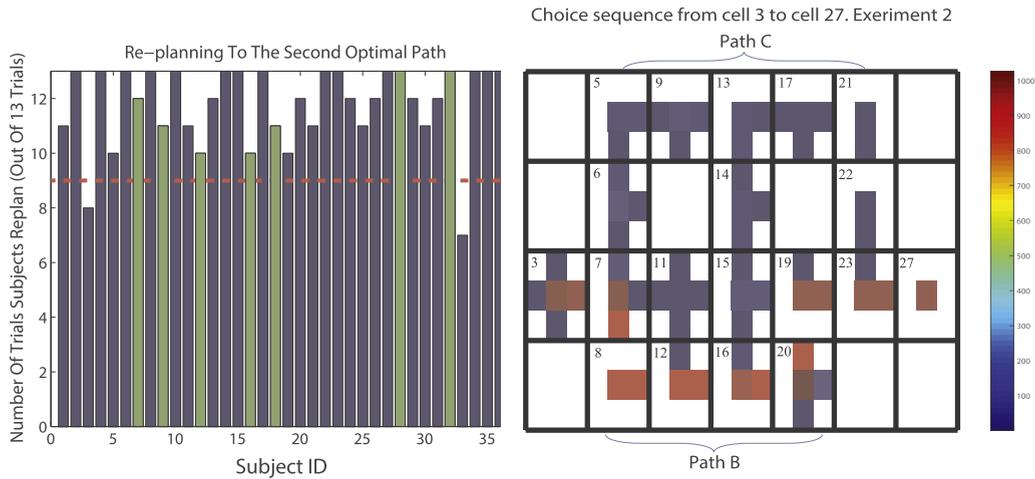


Fig. 5. Experiment 2. Left: Blue bars show participants who correctly re-planned to the second optimal path which was path C1A2 (except the one below the red threshold) and green bars are the group of participants who re-planned to C1C2, the third optimal path, ($\alpha = 0.05$). The optimal path was blocked in 13 trials (out of 40 trials) in the test blocks. The goal and starting points were fixed but the detour path contained stochastic punishments. Right: (Heatmap) Participants’ choice in planning trials in the test blocks in experiment 2. The heatmap shows the number of times participants selected the optimal path in the planning trials when path B was not blocked. Path B went through the following cells: 7, 8, 12, 16, 20, 19, 23, 27. The blue cells shows path C1C2 and C1A2. A cell with darker red color is visited more often. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the optimal and non-optimal paths separately. We found that out of 32 participants, there were five who were able to find the optimal path in the pretest and the test trials, but not able to do so in the re-planning trials. Also there are two other participants that were able to plan optimally in the pretest and re-plan in the test phase but failed to choose the optimal paths in the test phase where there is no blockage. And the rest behaved optimally in both pretest and test blocks for planning and re-planning trials, see Fig. 6, left panel.

We had 20 different pairs in the pretest block (pair (7, 13) was different from pair (13, 7)). Only two participants did not plan optimally in this block. For some pairs, e.g. pair (13, 21), finding the optimal path was more challenging. In this particular pair, the non-optimal path was 13, 17, 21 with stochastic punishment at cell 17 with expected loss of -24.2 and the optimal path was the longer path of 13, 14, 15, 19, 23, 22, 21 with expected loss of -4.2 at cell 19. Participants needed to compare the longer path and greater expected value with the shorter path and smaller expected value. At pair (5, 13), the optimal path, path B, was even longer and less appealing. On average participants showed mediocre performance (below chance) in this pair. At pair (14, 22), though, they had no problem finding the optimal path probably due to the fact that path 14, 13, 17, 21, 22 and path 14, 15, 19, 23, 22 were symmetric with equal length and participants just needed to compare each path’s expected reward.

Fig. 6, right panel, shows how often participants chose the optimal path in the test phase when it was not blocked (inspecting the planning behavior) averaged across all participants from cell 5 to cell 13. The heat map shows that participants chose to go down, heading to the optimal path, in spite of the fact that the shorter path requires only two steps. The panel at the bottom demonstrates re-planning behavior from cell 3 to cell 27 in the test phase with fewer data points. Note that in the test phase of experiment 3, each pair was presented 6 times and in only 2 of them the optimal path was blocked. Thus the maximum number of times to choose the optimal path for all participants in experiment 3 for each pair is $32 \times 6 \times \frac{2}{3} = 128$.

The bottom panel of Fig. 6 shows participants’ choice in the re-planning trials for pair (3, 27) summarized as follows: (a) Participants chose path C1A2 more than path C1C2 when path B was blocked. (b) Cell 8 should have been visited $2 \times 32 = 64$ times (light green) if all participants behaved optimally. But participants were not optimal in all trials, cell 8 was visited less frequently (dark blue). (c) Similarly, the total number of times that cell 7 should have been visited is 128 (going from 3 to 7 and then going from 8 to 7). In our data, it is 124 due to the fact that there were trials that participants selected path C1 without checking path B and thus visited this cell once. (d) Cell 6 and cell 14 have been visited more than 64 times (75 and 87 times respectively). We found that some participants failed to remember the correct transition matrix when they experienced a blockage in the optimal path and hit the surrounding obstacles two or three times before selecting the correct direction.

With the salient loss of -75 (with probability of 0.8) occurring at cell 11, we compared the rate that participants received this loss before and after the pretest block. The rate of experiencing -75 before the pretest block was 12.5 while this rate went down to 1.25 during the pretest and test blocks in experiment 3 ($t(31) = 9.241, p < 0.0001$). We should emphasize that in most of the trials before the pretest block, the path containing the -75 punishment was one of the accessible paths from the starting point to the goal position and more importantly in 60% of these trials (73 out of 120 trials) this path was the shortest path or one of the two equally shortest paths. The fact that the rate of selecting this path was reduced significantly from the learning blocks to the test phase showed that participants learned to avoid the great losses as they entered to the pretest and test blocks. We did the same analysis on the cells that deliver losses greater than -5 . The rate of hitting cell 17 with -30 loss in the learning blocks was significantly different from the test blocks ($t(31) = 6.253, p < 0.0001$), but for cell 9 with loss of -20 , there was no significant changes. Note that all of these losses are probabilistic. From the mixed-effect regression analysis on the average reward learning curve (in the learning blocks), we found that the block number as a predictor was significant ($t(191) = 6.478, p < 0.0001$).

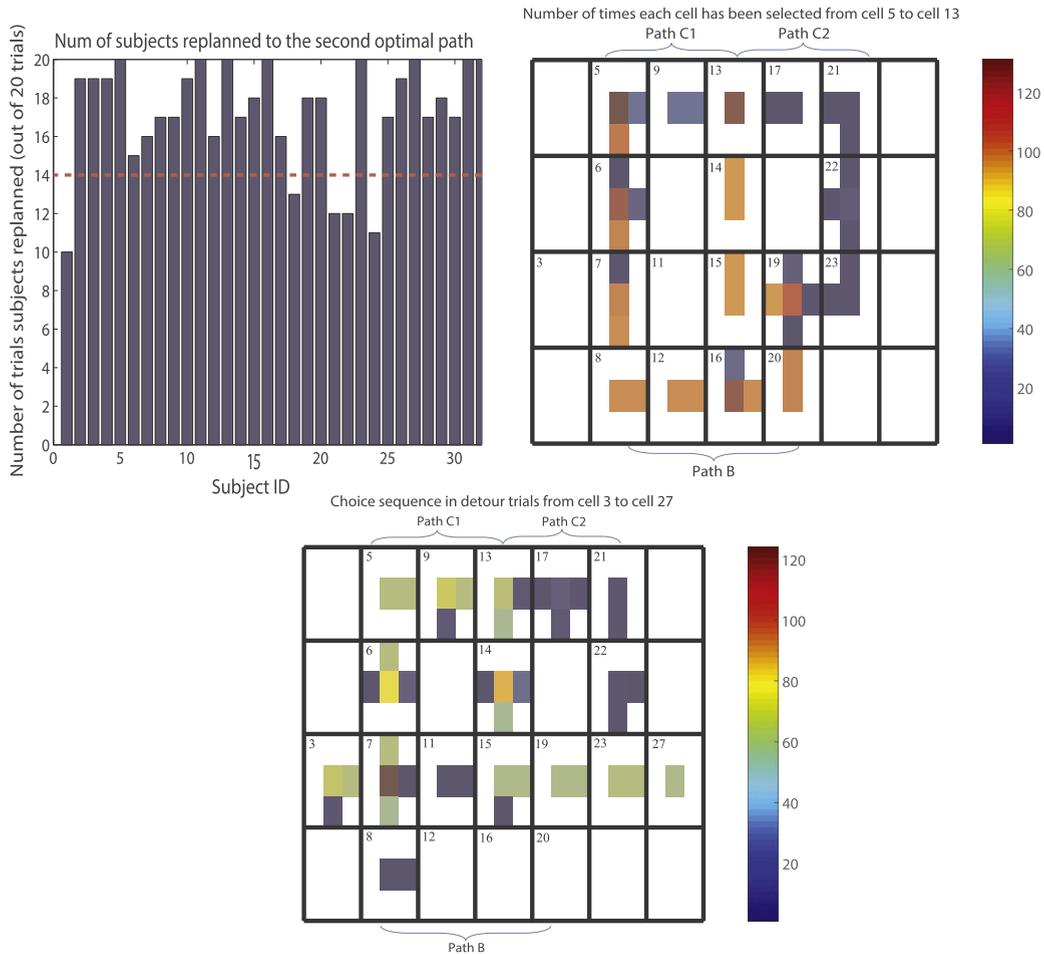


Fig. 6. Experiment 3. Left: Blue bars above the red line indicate participants who could re-plan optimally. The optimal path was blocked in 20 trials (out of 60 trials) in the test blocks. Right: Heatmap of participants' choice in the planning trials in the test blocks in experiment 3 when the starting position was at 5 and the goal was at 13. Pair (5, 13) was one of the challenging pairs in our design, which ruled out suggested models. The heatmap shows that majority of participants chose to take path B instead of going directly to cell 13 using path C1. There was a (possible) decision node when they reach cell 19. A few chose to take path C2 and were faced with the loss at cell 17. Bottom: The number of times participants visited each cell in the grid when path B was blocked for pair (3, 27). Cell 12, 16 and 20 are part of path B and that is why they never get selected in re-planning trials. The second optimal path went through 7, 6, 5, 9, 13, 14, 15, 19, 23 and 27 and was selected more frequently. Pair (3, 27) was repeated 6 times. Note that white cells with no number are obstacles. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In the test blocks, participants were faced with random blockage of the optimal path, which might weaken their performance in recalling the obstacles.³ We used the mixed-effect regression to fit the number of times a participant hit an obstacle as a function of block number in the pretest and test blocks (considering participants as the random effect) and we did not find any significant difference between the rate of hitting obstacles in the pretest block vs. the test blocks ($t(2527) = 1.426, p = 0.154$).⁴

3. Theoretical analysis

Reinforcement learning (RL) algorithms have been previously used to explain human behavior during learning of dynamic decision making tasks (Huys et al., 2012, 2015; Simon & Daw, 2011; Sutton & Barto, 1998) and can be formulated within a Markov decision process (MDP) framework. In a MDP, the communication between an agent and the stochastic environment (at time step, k) is through three sets defined by state $\{s^k\}$, action $\{a^k\}$ and (probabilistic) reward r_k . After taking action a_k the environment transfers to a new state s_{k+1} with probability $\mathbf{T}_{ij}^u(k) = Pr(s_{k+1} = s^j | s_k = s^i, a_k = a^u)$ and the agent receives a probabilistic reward $r_k = r$ with probability $\mathbf{R}_{ij}^u(r, k) = Pr(r_k = r | s_k = s^i, s_{k+1} = s^j, a_k = a^u)$, Khodadadi, Fakhari, and Busemeyer (2014). The transition function and the reward function have the Markov property meaning that the transition probability and the reward probability are independent of the

³ Besides the burden of recalling grid world's configuration, starting from the 7th block, participants actually earned monetary rewards based on their choices and this might impair their performance and memory recollection.

⁴ Note that the obstacles were deterministic (in contrast with probabilistic punishments).

history of states and actions $\{s_1, a_1, \dots, s_k, a_k\}$ and merely depend on the state and the action at time step k . In other words:

$$Pr(s_{k+1} = s^j | s_k, a_k, \dots, s_1, a_1) = Pr(s_{k+1} = s^j | s_k, a_k) \quad (1)$$

$$Pr(r_k = r | s_{k+1}, s_k, a_k, \dots, s_1, a_1) = Pr(r_k = r | s_{k+1}, s_k, a_k) \quad (2)$$

The main goal is to find the optimal policy, π^* , which determines the probability of selecting action a in state s while maximizing a desired function of accumulated rewards called *return*. A widely used form of the return function is the *expected sum of future reward discounted by γ* , Sutton and Barto (1998):

$$E \left[\sum_{k=0}^{\infty} \gamma^k r_k \right] \quad (3)$$

where discounting factor, γ , weighs the future reward relative to the immediate reward. Using the return function in Eq. (3), *state-action* function, $Q(s, a)$, can be defined as the expected discounted sum of rewards received in state i and action j at time step k given the policy, π :

$$Q_{\pi}(s^i, a^j, k) = E \left[\sum_{u=k}^{\infty} \gamma^{u-k} r_u \mid s_k = s^i, a_k = a^j, \pi \right] \quad (4)$$

$$Q_{\pi}(s_k, a_k) = E[r_k + \gamma Q_{\pi}(s_{k+1}, a_{k+1})] \quad (5)$$

This equation is known as the Bellman equation. The expectation on the right hand side of this equation depends on the functions \mathbf{T} and \mathbf{R} . Notice that the Bellman equation provides one equation for each state and so can be considered as a set of recursive equations or more formally a system of equations. If the transition and reward functions are known, the Bellman equation can be solved by dynamic programming method, Bertsekas and Tsitsiklis (1996). However, in most cases, including our experiments, the agent does not know the model of the system.

The model-based method tries to estimate \mathbf{T} and \mathbf{R} functions and solve the Bellman equation with *value iteration*, Sutton and Barto (1998) and Daw et al. (2005). In this approach, the model-based agent learns the model of the system and uses this model to find the optimal policy. However, the model-free approach such as temporal difference learning algorithm, uses an estimate of the difference between the two sides of Bellman Eq. (5) which is called the temporal difference (TD) error. Therefore, *without* knowing the \mathbf{T} and \mathbf{R} functions, TD is able to learn the optimal policy which can maximize the return function. It is also computationally simple and efficient. However, its ascendancy is limited to a stationary environment. Since it *locally* updates the state-action function (Eq. (5)), if there is a change in the environment (e.g. re-planning in our experiments or reward devaluation) it takes a long time to propagate that (new) information to update all the state-action estimations. In addition, in the training blocks of our experiments, the TD algorithm cannot learn the correct state-action function or learn the optimal path because the goal and starting point are not fixed. At best it can learn different sets of state-action function for each pair and for each pair retrieve that relevant information (which is reduced to a look up table). Because of these reasons, the model-free RL (model 1) is not a good candidate to analyze our results, but to complete our model comparison benchmark, we documented the quantitative fit.

The simplest form of the temporal difference (TD) error is defined in Eq. (6) and is used to update the state-action function which is called the Q-learning model (Eq. (7)):

$$\delta_k = r_k + \gamma \cdot \max_a \widehat{Q}_{\pi, k}(s^i, a) - \widehat{Q}_{\pi, k}(s^i, a^j) \quad (6)$$

$$\widehat{Q}_{\pi, k+1}(s^i, a^i) = \widehat{Q}_{\pi, k}(s^i, a^i) + \alpha_c \cdot \delta_k \quad (7)$$

Note that α_c is the learning rate. Because the goal and starting points could be in any side of the grid, the associated $Q(s, a)$ is initially set to Q_0 for all cells in the grid, Simon and Daw (2011). We could have specified different Q_0 for each cell and for each action according to their position in the grid and the starting position (and especially if they were next to the grid's borders), but this would add more free parameters to the model with negligible benefit. Using the Q-values, an action selection module, e.g. Softmax decision rule, can be used to choose the appropriate action with a probability of:

$$p_k(a) = \exp(\beta \cdot Q_k(s, a)) / \sum \exp((\beta \cdot Q_k(s, a))) \quad (8)$$

where β is the temperature. If $\beta \rightarrow 0$, the algorithm selects a random action (in our case, we have 4 actions: up, down, right and left, thus each action is selected with probability of 0.25) and if $\beta \rightarrow \infty$, the action with the maximum Q-value is selected (as in ϵ -greedy policy).

We also used a baseline model, with no learning. A baseline model (model 1) randomly chooses an action at each state and does not consider the previous experiences. Unlike the two models that we explained above, modeling participants' behavior in our task is not trivial. A candidate model should be able to plan ahead and estimate the reward function simultaneously and appropriately. With these goals in mind, the baseline model has no planning at all and fails to learn the payoffs in the grid world and the model-free RL has problems in learning the correct reward function and planning because of the dynamically changing environment. Again, the model-free RL plans for one step-reward and with changing environment this strategy is not useful. In the following, we explain how different strategies of estimating the reward function and various manipulations of the planning module can change models' predictions in the learning and test phases.

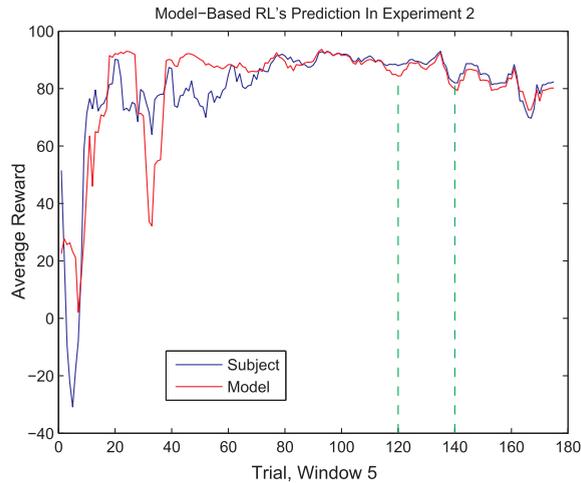


Fig. 7. Model-based RL's prediction in (red) for experiment 2. Green lines indicate the beginning of the pre-test block (7th block) and the test blocks (8th and 9th). Data related to subject 14. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.1. Model based reinforcement learning

As discussed earlier, the model-based RL solves the Bellman equation by learning the model or representation of the environment (cognitive map in Tolman (1948)). Unlike TD learning algorithm, the model-based RL's estimates and updates are *global* and because of this distinction, when the environment changes or the reward is devalued, the model-based RL can respond to these changes quickly. We compute the state-action function with q-value iteration (Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998) as follows:

$$\hat{Q}_{k+1}(s^i, a^u) \leftarrow \sum_{s^j} T(s^i, a, s^j) [R(s^i, a, s^j) + \gamma \max_a (\hat{Q}_k(s^j, a))] \tag{9}$$

Generally estimating **T** and **R** function depends on the environment with which the agent is dealing. Since the obstacles in our grid world are deterministic, with only one experience, participants can learn the transition function. We can also assume that each time an action is selected in a specific state, the transition function is updated by a learning rate respectively. In order to learn the probabilistic reward function, one possible method is to take the average of what the participant has earned for each state-action pair so far (which will be discussed later in Eqs. (10) and (11)). Using these two pieces of information, we can compute the state-action value function ($Q(s, a)$) at *each* state with q-value iteration (Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998). Once we have state-action function, we choose the appropriate action using Softmax decision rule as in Eq. (8).

The average reward predicted by model-based RL for experiment 2 is shown in Fig. 7. The model-based RL predicted similar behavior as our subject (No. 14). To show the qualitative fit, we chose a very basic version of the model-based RL that updates its estimation *at each step* using q-value iteration method. We assumed that the model knows the correct transition matrix and does not need to learn the **T** function.

Taking the average of previous samples requires a perfect memory (specifically at cells with stochastic rewards) and it is not practical for human participants. Instead, we use a linear filter to approximate the reward function dynamically using Eq. (10):

$$R(k + 1) = \alpha_1 \cdot R(k) + (1 - \alpha_1) \cdot r \tag{10}$$

where r is the reward that is delivered upon transition and $R(k)$ is the averaged reward up to time step, k . α_1 determines the relative importance of the averaged (observed) reward versus the current reward. We also include a forgetting factor, α_2 , which plays a decaying role in reward estimation (Eq. (11)) when the greater loss is not experienced (e.g. in 20% of the time, these cells have a regular loss of -1):

$$R(k + 1) = \alpha_2 \cdot R(k) \tag{11}$$

Assuming that people are only sensitive to great losses, learning and estimating the reward function only occur when a great loss is experienced. Whenever a participant experiences a great loss (anything but -1) at a cell, the model updates its estimation using Eq. (10). We tested models with different assumptions regarding α_1 , α_2 by categorizing the losses into three levels of saliency (low, medium and high): Cell 16 and cell 19 with possible losses of -3, -5 (low saliency, LS), cell 9 and cell 17 with probable losses of -20, -30 (medium saliency, MS) and cell 11 with loss of -75 in 80% of the time (high saliency, HS). Similar pattern can be formed for experiment 1 with -45 at cell 21 with medium salient loss, very salient loss of -75 at cell 15 and a small loss of -3 at cell 16. Table 1 in Supplementary materials summarized these models along with their parameters. Based on Bayesian Information Criterion (BIC), we found that the model with $\alpha_{1HS}, \alpha_{1MS}, \alpha_{1LS}$ and no decay rate was the best to use.⁵ Cell 16 and cell 19 with smaller losses are parts of

⁵ Table 2 in Supplementary materials. Also check Fig. 1 in Supplementary materials which shows the estimated reward at different cells with stochastic losses (cells 11, 9, 17, 16 and 19) with no decay rate.

Table 2
Summary of models and their free parameters.

Model	Description	Parameters	No. of free parameters
1	Baseline (random model)	–	0
2	Q-learning	$\beta, \gamma, q_{u0}, q_{r0}, q_{d0}, q_{l0}$	6
3	Model-based RL	$\beta, \gamma, \alpha_{HS}, \alpha_{MS}, \alpha_{LS}$	5
4	Avoids Salient Loss-MBRL	β, γ	2
5	Remembers The Last R-MBRL	β, γ	2
6	Finds The Shortest Path-MBRL	β, γ	2
7	Cubed Model-Based RL	$\beta, \gamma, \alpha_{HS}, \alpha_{MS}, \alpha_{LS}, \omega$	6
8	Successor Representation	$\beta, \gamma, \alpha_l, \alpha_{HS}, \alpha_{MS}, \alpha_{LS}$	6
9	Avoids Salient Loss-SR	β, γ, λ	3
10	Remembers The Last R-SR	β, γ, λ	3
11	Finds The Shortest Path-SR	β, γ, λ	3
12	Hybrid SR-MB	$\beta, \gamma, \alpha_l, \omega_{hb}, \alpha_{HS}, \alpha_{MS}, \alpha_{LS}$	7

MB = Model-based, SR = Successor Representation.

optimal path in most of the pairs in the pretest block and participants visit those cells more often. Thus the algorithm sets a smaller α_1 for these cells as if the expected reward is basically the current reward. The losses at cells 17 and 9 are less extreme than cell 11 and it is not always optimal to select these paths. Therefore, the algorithm estimates their rewards with different rates.

In addition to linear filter, there are other methods to estimate reward function as well (e.g. simple strategies or rules). In the next section, we introduce 3 different heuristics that we use to learn the reward structure. Note that we are still using the model-based RL framework but each strategy estimates the **R** function differently with small or no computational cost (model 4, 5 and 6).

3.2. Heuristic-based models

We use 3 different heuristic-based models to fit to our data (training dataset). These heuristics are: avoiding the great loss, remembering the last reward and finding the shortest path. These heuristics along with previously described models are summarized in Table 2.

3.2.1. Avoiding the salient loss

As we discussed above, participants were sensitive to the punishment at cell 11 with expected value of -60.2 . This cell had the greatest and the most *salient* loss. A candidate model which avoids the salient loss can explain the behavior of our participants who chose not to enter to cell 11 (model 4). This model does not estimate the reward function but only avoids cell 11 in experiment 2 and 3 (and cell 15 in experiment 1) with -75 loss. Note that this behavior does not indicate that these participants chose the optimal path necessarily. For instance, in pair (5, 13), we had two groups of participants: The ones who failed to find the optimal path by choosing path C1 (and facing the loss of -20 at cell 9) and not path B (with the minimum loss) and participants who learned about other punishments which are not as salient as cell 11 and at the end chose path B (which is optimal).

3.2.2. Remembering the last reward

The fifth model (in Table 2) is categorized as a model-based RL, but only remembers the last reward and does not estimate the **R** function. This model saves the last value of the reward at each cell. Thus with one bad experience (e.g. receiving -75 at cell 11), it is less probable for the model to enter that cell. The last value of any experience is saved for the rest of the computations and action selection. This model does not need the full memory of reward samples because it does not need to estimate or learn the stochastic reward structure. This simplicity can cause sub-optimal behavior in re-planning trials during the test phase. Consider the following scenario: In the 8th block at pair (7, 13), the optimal path (path B) is blocked, the model chooses path C1, assuming that the last experience of path A1 was with the great loss at cell 11. Then at cell 9, it would receive -1 instead of -20 . This would be saved in the model's memory. For the next pair, e.g. pair (5, 13), which includes cell 9, the model selects path C1 over path B simply because its last experience at cell 9 is with a regular punishment. Moreover, it only takes two steps to catch the goal with the expected value of 99. Path B, however, requires participants to take 10 steps and its expected value is 88.4. This strategy would be very useful when the environment is deterministic.

3.2.3. Finding the shortest path

There are few participants that chose the shortest path (A1A2) in the trials with a blocked optimal path, but were able to find the optimal path in regular trials. This suggests that these participants did not care about maximizing their reward when they encountered an obstacle in the optimal path (against the goal of the experiment), and it is only important for them to catch the goal using the remaining available moves. This approach (model 6), with no re-planning, finds the shortest path instead of choosing a path with minimum loss. Also note that, choosing the shortest path, is not necessarily distinguishable from other strategies since there are so many pairs in these experiments that have two paths of equal length. For instance, in experiment 1, in the pretest block, 12 out of 20 pairs have this characteristic: a path with a regular punishment and one with a great loss but both have equal length. Therefore,

model 6 can choose any of these paths with equal probability.

In Experiment 1, there were two people who chose the shortest path in those trials. In experiment 2, there were 6 participants who failed to choose the second optimal path and were not sensitive to loss -30 and one person who chose the shortest path ignoring all the losses. Finally in the third experiment, there were 7 participants with similar pattern in their behavior. What surprised us in these results was that these participants learned to find the optimal path when they were asked in the pretest block but with experiencing the blockage in the optimal path, they ignored or forgot their findings on reward structure and switched their strategy from choosing the optimal path to the shortest path.

3.3. Cubed model-based RL: Planning on a smaller grid

Earlier, we explained how the shortest path and the optimal path are different in our design. Based on this distinction, we define the *length* of each path as the number of steps in the shortest optimal path. Thus the length of the path between cell 5 and cell 13 is 10 because the optimal path contains path B but it only takes two steps to reach cell 13 from cell 9 (disregarding the stochastic punishments). Now assume that at the beginning of the learning phase a participant is at cell 5 and she wants to reach cell 13. At the first glance, path C1 is the shortest and path B is the longest. But if she only takes one step, she would probably (80% of the time) faced with the great loss at cell 9. Next time she would divert her decision to either path A1 or path B. Assuming she chose path A1, after three steps, she would experience the maximum loss at cell 11 with probability of 0.8. These two results would lead her to take path B in the next trial. In this scenario one could use q-value iteration for three steps ahead (instead of the whole path as in model 3) and prune the rest of the paths and lower the computational cost (model 7). This model, which we called “cubed model-based RL”, is fundamentally different from the previous heuristic-based models because the planning module has been changed. In fact, Model 7 plans for a small window ahead of the current position (using the shortest length for each pair).

The minimum number of steps to prune the decision tree can be an extra free parameter in cubed model-based RL. While pruning at three steps is good for pair (5, 13), it is not a good stopping point for pair (3, 17). In order to compare path C1 and B correctly, one should plan for at least four steps ahead from cell 3 to be able to see the losses at cell 12 and cell 9. So for different pairs, the model needs to use different pruning depths. The diversity of pairs in each experiment (e.g. experiment 3 has 65 pairs) forces this model to pick the maximum number of available steps for planning (which is 7 in our grid world) and thus with one extra parameter, cubed model-based RL cannot be distinguished from model 3 for many participants at many pairs as we report in the model fit result.

So far we have discussed about the first 7 models in Table 2. These models are different in their planning and learning approaches. In traditional model-free RL, the Q-values are computed using TD errors. But since the goal is not fixed and the starting points are randomized, this approach cannot learn the correct Q-values at each states. Moreover, TD mechanism is unable to plan ahead when the environment is not deterministic. The baseline model lacks learning and planning regardless of the environment. We proposed 3 different heuristics (avoiding salient losses, finding the shortest path and remembering the last punishment) based on our data. Although these strategies are simpler to simulate and require less computations (because they do not estimate the R function), they can *only* explain a small portion of our data and mainly have wrong predictions. The (full) model-based RL, model 3, learns the model of the environment and estimates the reward structure. Having the model of the environment, which is independent of the goal or the starting position, helps the agent to predict the consequence of each action before taking them.

3.4. Successor representation

In addition to the main RL families that have been explained so far, there is another algorithm called successor representation (SR) that was first introduced by Dayan (1993) which has the flexibility of the model-based RL and the simplicity of the TD learning. In this algorithm, the Q-values are decomposed into a reward matrix and a successor map which predicts the (discounted) future occupancy of all states, M (which is called SR matrix). Given any initial state, s^i , the SR matrix counts the number of times that a subsequent state (in a trajectory) is visited later:

$$M(s^i, s^u, a^j) = \mathbb{E} \left[\sum_{\nu=0}^{\infty} \gamma^{\nu} \mathbb{I}[s_{k+\nu} = s^u] | s_k = s^i, a_k = a^j \right] \quad (12)$$

where $\mathbb{I}[\cdot] = 1$ when its argument is true and zero otherwise, Kulkarni, Saeedi, Gautam, and Gershman (2016). While it is possible to compute the SR matrix from the transition matrix (Gershman et al., 2012), it is more common (and less expensive) to estimate SR matrix via Bellman equation as we did in Q-learning and the updating rule is as follows (α_l is a learning rate):

$$M(s^i, s^u, a^j) \leftarrow M(s^i, s^u, a^j) + \alpha_l [\mathbb{I}[s_{k+1} = s^u] + \gamma M(s_{k+1}, s^u, a_{k+1}) - M(s^i, s^u, a^j)] \quad (13)$$

In order to calculate the Q-values given a policy π , we need to compute the inner product of the reward and SR matrix using Eq. (14), Kulkarni et al. (2016):

$$Q_{\pi}(s, a) = \sum_{s^u} M(s^i, s^u, a^j) R(s^u) \quad (14)$$

As we discussed before, there are multiple methods to estimate reward function (three different heuristics and one linear filter, Eq. (10)). Consistent with our model-based RL fit, we use the linear filter and heuristics to learn the reward function (avoiding the great loss, remembering the last R and finding the shortest path related to models 9, 10 and 11 respectively).

3.5. Hybrid SR-MB

It is important to emphasize that the SR matrix, M , is different from the transition matrix. Thus, if there are abrupt changes in the environment, i.e. transition revaluation, the SR model fails to adapt to these changes without learning the new trajectories. In our experiments, in the test blocks, the random blockage in the optimal path is one example of transition revaluation. In order to choose the correct action, the SR algorithm needs to relearn the (new) trajectories. One possible solution to this problem is to combine the SR model with the model-based RL model (hybrid SR-MB), Momennejad et al. (2016) and Russek, Momennejad, Botvinick, Gershman, and Daw (2017). For instance, the probability of each action can be a weighted average of the model-based RL's predicted probability and the SR's predicted probability. By assigning a greater weight to the model-based RL's prediction, the hybrid SR-MB can be more flexible toward the sudden changes of the transition matrix.

4. Model fit results

Although we have 12 different models, some of them share common characteristics in planning or reward estimation. For instance, model 5 is different from model 4 in estimating the reward function (model 5 only remembers the last reward while model 4 avoids the greatest loss), but they both use model-based decision making in planning. Because of these differences and similarities, we divide these models into 6 unique categories: (a) the random choice model with no planning and no learning (model 1). (b) the model-free RL which does not learn the structure of the environment but does plan based on one-step rewards and the next step's expected future reward (model 2). (c) the model-based RL and/or with heuristics (including models 3, 4, 5 and 6). These models are similar in their planning but differ on how they estimate the reward structure. (d) the cubed model-based RL (model 7) with a constrained view of the environment. (e) the SR model and/or with heuristics (including models 8, 9, 10 and 11). These models learn a similar multi-step representation of the environment, but their estimations of the reward function are different. Note that the SR models are not as good as model-based RL in learning the environment but better than the model-free RL. (f) the hybrid SR-MB (model 12) which blends the SR algorithm with the model-based RL to improve its performance in transition devaluation.

In some of these categories, i.e. category c, there are four different models and in some, i.e. category f, there is only one model. Our strategy to compare these models is to first nominate one candidate in each category (within group comparison) and then measure how good they are in contrast to other categories. Note that we first compare the performances of these models in the learning blocks and then evaluate their predictions in the test blocks.

4.1. Model fit results in the learning (training) blocks

There are multiple ways to evaluate models. The first that we chose to report is based on BIC comparison (Schwarz, 1978). To this end we fitted these models to the data from learning blocks (block 1–6 with $n = 120$ trials) using one-step-ahead prediction method, Eq. (15). In this method, the model predicts (each) participant's choice on the next trial using the sequence of choices and payoffs that she has taken and experienced respectively. Specifically, the candidate model calculates the probability of selecting a particular action, that the participant selected, at each cell using the history of choices and rewards/punishments ($\{s_1, a_1, r_1, \dots, s_k, a_k, r_k\}$).

$$LL_i = \sum_1^{n-1} \ln(\text{Prob}(a_{k+1} | \{s_1, a_1, r_1, \dots, s_k, a_k, r_k\})) \quad (15)$$

In order to find the best fitted parameters, the maximum likelihood estimation is applied (simplex search algorithm) and BIC is computed as in Eq. (16):

$$BIC = -2 \cdot \ln(LL_i) + k \cdot \ln(N) \quad (16)$$

where N is the number of observation and LL_i is the likelihood of the i th participant.

First, we compared the models in the model-based RL category (model-based RL and heuristics, category c). The result of our model fit is summarized in Table 3 of the supplementary materials. Model 5 which remembers the last reward wins in experiment 1, but model 3, model-based RL with a linear estimation of reward, has the lowest BIC in experiment 2 and 3. Simplicity of the environment in experiment 1 plays a critical role in this result. In experiment 1 there were only three paths (rather than 5) and the detour path (path C) had no stochastic punishment. The loss at path B, the optimal path, was very small and even if the participant was unlucky, the worse case loss was -3 . In path A, however, the difference between the regular punishment and the worse case loss was considerable. In 80% of the times, the environment delivered a loss of -75 and thus with one such experience the model with the strategy of remembering the last punishment was no longer likely to choose this path. Therefore, model 5 was able to find the optimal path with only two free parameters (temperature and discount) without learning the reward structure. In the other two experiments because there was 20% chance that participants receive a regular punishment of -1 at cell 9 and cell 17, the strategy of recording *only* the last punishment might mislead the model to choose the wrong path.

Similarly, the (within category) model comparison for SR models (in category e) is outlined in Table 4 of the supplementary materials. The full SR model had the better BIC in experiments 2 and 3 while the heuristic-based SR model which ignores the rewards (and chooses the shortest path) won in experiment 1.⁶

⁶ Important to mention that in experiments 2 and 3, we had similar AIC results but not in experiment 1. In fact the full SR model had a better AIC.

Table 3
Model fit result- BIC comparison.

BIC	Baseline	Q-learning	MBRL	MB-RemR	SR	SR-ShPath	CMBRL	Hybrid SR-MB
Exp 1	2571	2266	–	1891	–	1518	1917	1520
Exp 2	2638	2479	1853	–	1896	–	1997	1848
Exp 3	2626	2451	1832	–	1809	–	2255	1738

MBRL = Model-based RL; RemR = Remembers the Last R; ShPath = Find the Shortest Path; SR = Successor Representation; CMBRL = Cubed Model-based RL.

Next, we compared the remaining models (the winner in each category) using BIC for each experiment, [Table 3](#).

In experiment 1, the SR model which finds the shortest path (with 3 free parameters) had the lowest BIC. Note that the hybrid SR-MB model had also a very low BIC (1520) although it had 7 free parameters. In experiments 2 and 3, the best model was the hybrid SR-MB. Baseline model, Q-learning and Cubed model-based RL had the worst fit as we expected. The full model-based RL was among the top three models in all experiments. In experiments 1 and 3, it was the third best model but in experiments 2, it held the second place. Note that all of these results were computed by the data in the learning blocks.

One way to compare these models is to treat BIC/AIC⁷ as the log model evidence for each participant and investigate if there is a significant difference between the BIC (AIC) for a pair of models, [Khodadadi, Fakhari, and Busemeyer \(2017\)](#). However, this needs many pairwise comparisons. Moreover, the results based on AIC and BIC are not consistent to some extent⁸: based on AIC, there is much stronger evidence that the hybrid SR-MB model is the best model in experiment 1. A more sophisticated approach for comparing this number of models is to compare their Exceedance Probability (EP) which have been proposed by [Stephan, Penny, Daunizeau, Moran, and Friston \(2009\)](#). In this algorithm, each model was treated as a random variable and the probability of generating (participants') data by a model was defined by a multinomial distribution described by a Dirichlet distribution. Stephan et al. used a variational Bayes method to estimate the Dirichlet distribution's parameters based on marginal likelihood of each model, [Stephan et al. \(2009\)](#).⁹ Given the parameters of the Dirichlet distribution, it is possible to compute the probability that a model is more likely than any other model.¹⁰ The best model has the highest EP (probability closer to 1). [Table 4](#) shows the EP of six models that we discussed above.

In experiment 1, the hybrid SR-MB model and the SR-ShPath model were competing with each other and as we expected, the EPs calculated by AIC and BIC were not consistent. In experiment 2, the full model-based RL was more or less similar to the hybrid SR-MB model while in experiment 3, the superior fit was provided by the hybrid SR-MB model. These results confirmed our previous findings that the heuristic-based models were not able to explain the data in a more complicated environment. The hybrid SR-MB model was the best model in all three experiments if we only considered the AIC results (and AIC-based EP). For the next section, we chose the best three models in each experiment and compared their predictions in the test blocks.

4.2. Model prediction in the test blocks

So far, we used the learning phase data to train our models and then used their quantitative fit to evaluate them. As we presented above, the hybrid SR-MB model was dominantly the best model (among our proposed heuristics and models) in explaining the data in all three experiments using AIC. These results were not so clear when we compare the BIC and the EP of these models. The SR-ShPath and the MBRL had a better fit in experiment 1 and 2 respectively. Therefore, we took our model comparison to another level and investigated their predictions in the test phase. Using the knowledge of the training phase, we examined whether these models can find the optimal path in an environment where they have not experienced yet. We evaluated these predictions using two approaches: (a) calculate the log-likelihood of each model during the test phase and compare these results (similar to what we have done in the training phase); (b) using the best fitted parameters, compare models' prediction at critical cells in the test phase.

[Table 5](#) summarizes the EP of the three most successful models discussed in model fitting section. The heuristic-based SR model (which finds the shortest path) had the greatest EP in experiment 1. In experiments 2 and 3, the model-based RL was the best model in explaining the test trials' data. To our surprise, the hybrid SR-MB model failed in all of these experiments. One possibility that can explain the failure of hybrid SR-MB model is that in our model fitting, we only used the learning blocks to train these models. When the environment was not changing, even a pure SR model (with or without heuristics) provided a better fit than a model-based RL ([Table 4](#) and [Table 4 of Supplementary materials](#)). But the SR model needs to experience new trajectories when the transition matrix changes. When we linearly combined the SR algorithm with the model-based RL, greater weights were assigned to the SR model's prediction.¹¹ Therefore, the model-based RL had small or zero contribution in predicting the re-planning trials' data in the test phase.

To investigate this issue, we fitted the hybrid SR-MB model to all data. Since there are 13 trials out of 180 trials in experiment 1 and 2, and 20 out of 200 trials in experiment 3 that the re-planning occurs, the fit did not change the results. The model-based RL in the hybrid SR-MB model still has *negligible* coefficient to contribute in all experiments. Specifically, these coefficients were 0.035, 0.08

⁷ Akaike Information Criterion, AIC, is calculated using Eq. (16), but $\ln(N)$ is substituted by 2.

⁸ Probably because the comparisons are affected more by possible outlier values in BIC or AIC.

⁹ We used the MATLAB code developed by [Gershman \(2016\)](#).

¹⁰ To compute these probabilities, it is necessary to have an estimate of the marginal likelihood of each model m for the data set D , i.e. $p(D|m)$.

¹¹ The average weights for the SR model in our experiments were close to 1 and almost 0 for the model-based RL.

Table 4

Exceedance probability of different models in 3 experiments using BIC and AIC.

EP	Baseline	Q-learning	MBRL	MB-RemR	SR	SR-ShPath	CMBRL	Hybrid SR-MB
EP-BIC-Exp 1	0	0	–	0	–	0.999	0	0.0001
EP-AIC-Exp 1	0	0	–	0	–	0.247	0	0.753
EP-BIC-Exp 2	0	0	0.5281	–	0.0003	–	0.0003	0.4711
EP-AIC-Exp 2	0	0	0.4314	–	0	–	0.0002	0.5683
EP-BIC-Exp 3	0	0	0.1412	–	0	–	0	0.8588
EP-AIC-Exp 3	0	0	0.1074	–	0	–	0	0.8926

MBRL = Model-based RL; RemR = Remembers the Last R; ShPath = Find the Shortest Path; SR = Successor Representation; CMBRL = Cubed Model-based RL.

Table 5

Exceedance probability of different models in the test blocks.

EP	MBRL	MB-RemR	SR	SR-ShPath	Hybrid SR-MB
EP-Exp 1	–	0	–	0.6775	0.3225
EP-Exp 2	0.568	–	0.2018	–	0.23
EP-Exp 3	1	–	0	–	0

EP calculated in the test phase. As in the training blocks, the model with simple heuristic win in experiment 1 but in experiment 2 and 3 the model-based RL which estimates the reward structure and uses that in planning provides a better fit to data.

and 0.06 for experiment 1, 2 and 3 (averaged across participants).

One solution to this problem is to use a mixture model which switches from the SR model in the planning trials to the model-based RL mechanism in the re-planning trials. We, also, fitted this model. With an extra free parameter, the fit allows the model-based RL (in the hybrid SR-MB model) to contribute in the re-planning trials. Note that the ratio of re-planning trials to planning trials is $\frac{1}{3}$. This time the coefficients increased to 0.32 in experiment 1, 0.12 in experiment 2 and 0.14 in experiment 3.

We also compared models' predictions at cell 8 after participants experienced a random blockage in the optimal path in experiment 1. In these trials, the optimal behavior is to go up to cell 7 and then reroute for the second optimal path (path C) to reach the goal. Note that in experiment 1, 13 out of 40 trials in the test blocks were re-planning trials. Also, that the starting position and the goal were fixed at 3 and 27. As we expected, the SR model which finds the shortest path (or basically any SR model) was not sensitive to the changes of the transition structure in the environment. The probability of selecting the right action at cell 8 was still greater than other actions (0.63), although path B was not available anymore. The predicted probability of selecting the up action, by SR model, was 0.08. However, in the model-based RL, the probability of selecting the right action is the lowest, 0.03. Similar patterns were found in experiment 2 and 3.

5. Discussion

Planning in a stochastic environment is challenging. It becomes even more challenging when the environment is unknown to us. No matter how complicated these problems are, we mainly use our previous experiences to deal with them. Sometimes the environment changes and forces us to change or modify our plan. As a result, we update our plan every now and then to make sure our plan becomes a success. In this article, we tried to study real life planning problems in a simplified situation using our grid world experiment. Using this framework, we developed three experiments to investigate planning and re-planning in humans while learning an unknown environment. After 6 blocks of training in a 4 by 7 grid, participants' planning skill was tested. At the beginning of the 7th block (the pretest block), we warned them that their score will be converted to monetary reward (the exchange rate was 0.01). The majority of our participants (19 out of 19 in experiment 1, 30 out of 36 in experiment 2 and 30 out of 32 in experiment 3) were able to find the optimal path. This is in contrast to previous studies that showed people are more likely to be sub-optimal in the description-based decision trees involving a probabilistic reward, [Hey and Knoll \(2011\)](#) and [Hotaling and Busemeyer \(2012\)](#). It has been proposed by [Erev, Ert, Plonsky, Cohen, and Cohen \(2017\)](#) that people show planning biases in description-based decision trees problems, but not in experience based learning. In what follows, we discuss our findings and contributions, separately.

5.1. Optimal planning in a grid world

[Simon and Daw \(2011\)](#) investigated the neural bases of planning in a dynamic maze comparing two well established RL models (TD learning and model-based theory). Participants navigated through a virtual maze (with 16 rooms) to earn the maximum possible reward. They found that the BOLD signal related to the value and the choice in the striatum is correlated with the prediction of model-based theory ([Simon & Daw, 2011](#)). In their design, participants did *not* need to estimate the expected values of the rewards because rewards were deterministic and the maze, although was constantly (and randomly) changing, was known to participants. Therefore, planning was equivalent to finding the new *shortest* path to reach the goal at each state by relearning the (new) environment while knowing the general structure of the maze. However, in our design, with the stochastic reward structure and

deterministic obstacles, the shortest path was not optimal and participants did not need to relearn the configuration.

The basic configuration of our grid world was similar to the detour problem in Tolman et al. (1946), but we modified the grid world into a stochastic environment where finding the shortest path was not optimal anymore. It is important to note that our grid world had two distinctive features that encourage goal-directed behavior: first, the (hidden) punishments were probabilistic, and second, the starting and goal positions were randomly located in different cells. In order to find the optimal path, participants needed to learn and compare the expected values of different paths. Employing probabilistic rewards enabled us to represent this problem in a decision tree framework to study (optimal) planning similar to experience-based decision-making tasks, Fig. 2.

5.2. Re-planning in humans

Through simple error-driven learning rules (e.g. TD learning), the model-free RL selects the action that leads to a greater reward (outcome) more frequently. The model-free RL stores the (action-) state values without learning the model of the system and therefore it is computationally simple. However, when the environment changes similarly to what we have in our grid world, a blockage in the optimal path, these state-action values are useless and need to be learned again! Note that the changes in the environment are not necessarily limited to a transition function, but also could be related to reward (e.g. reward devaluation), Dolan and Dayan (2013). This inability to be adapted to varying circumstances makes the model-free RL error-prone when the structure of the task is changing, Daw et al. (2005).

On the other hand, the model-based RL can update the state-action values globally due to its knowledge (or representation) of the system. For instance, in our grid world experiment, when the optimal path is blocked, the model-based RL only changes the transition function for that path. Consequently, this update changes the state-action values (Q-values) using a dynamic programming algorithm. This flexibility enables the model-based RL to pursue the goal without the need to experience and learn the environment again (goal-directed approach). A large amount of research has focused on how these two approaches cooperate or compete with each other in different areas (from clinical studies to neuroscience) using different tasks, Dolan and Dayan (2013), Decker, Otto, Daw, and Hartley (2016), Akam, Costa, and Dayan (2015), Boureau, Sokol-Hessner, and Daw (2015), Gillan, Otto, Phelps, and Daw (2015), Keramati et al. (2011), Lee, Shimojo, and O'Doherty (2014), Schach et al. (2014), Kool, Gershman, and Cushman (2017), Skatova, Chan, and Daw (2013), Sebold et al. (2014), Momennejad et al. (2016), Russek et al. (2017) and Keramati et al. (2016).

When a change occurs in our environment (e.g. grid world), we cannot rely on our habits anymore. We need to update our knowledge and modify our original plans to accomplish our goal (re-planning). To the best of our knowledge, re-planning in humans has been rarely studied, Momennejad et al. (2016) and Gershman et al. (2012) with only two-step decisions. We were able to analyze re-planning behavior by randomly blocking the optimal path in our test blocks. In experiment 1 and 2, the starting and goal positions were fixed during the test phase, but, in experiment 3, we randomized these pairs. Because of the stochastic nature of the payoffs, our results are not comparable directly with the rats used by Tolman (1948), but we showed that humans are capable of modifying their plans when there is a change in the environmental circumstances.

5.3. Computational modeling

We used 12 different models to fit the choice data. The baseline model selects a random action at each cell regardless of what participants have experienced. This model was the simplest model in our benchmark. The traditional Q-learning algorithm was not able to explain the data because the starting and goal positions in the environment in our experiments were not fixed. In addition, when a change happens in the environment, the whole set of the action-state values needs to be updated (again by extensive amount of learning and exposure to the new environment).

The (full) model-based RL tries to learn the model of the environment by estimating the transition and reward functions. This knowledge (of the environment) is later used to generate Q-value for each action. The model-based RL had the best predictions in the test blocks. In model 7, we restrict the model-based RL in its spatial search. Instead of a complete tree search that is commonly used in value iteration, we confine the model's planning depth to its k th nearest neighbors (k is a free parameter). While this modification can decrease the computational costs, for many pairs the best fitted k leads to a full tree search.

In addition to model-based and model-free RL, there is another alternative, SR, which is more flexible than model-free RL and computationally simpler than model-based RL, Dayan (1993). SR calculates the state values using both reward and a successor map which stores the expected and discounted future states' occupancies. In case of reward devaluation, SR's behavior is similar to model-based RL but when there is an alteration in the transition structure, it fails to adapt to the change (similar to model-free RL), Gershman et al. (2012), Kulkarni et al. (2016) and Momennejad et al. (2016).

Although the hybrid SR-MB model provided a better account for participants' choices in the learning blocks, it failed to predict the re-planning behavior in the test blocks. Since the candidate models were not trained by the test blocks' data, they needed to generalize their knowledge (from the learning blocks) to perform optimally in the re-planning trials when the optimal path was randomly blocked. One solution to this problem is to use a mixture model which switches from the SR model in the planning trials to the model-based RL mechanism in the re-planning trials. It can post hoc fit all the data the best, because it uses the best fitting model for the training blocks and then switches to the best predicting model for the test blocks.

Another possible candidate to explain the data can be a multi-step planning model proposed by Sutton, Precup, and Singh (1999) which introduces the concept of temporal abstraction of actions that are interrelated. So instead of selecting an action for each state, a sequence of actions can be chunked and used. This sequence is called an *option* and this framework is called hierarchical reinforcement learning (HRL), Sutton et al. (1999) and Botvinick et al. (2009). Multi-step planning models (MSP) can be implemented

in different ways. For instance, a MSP model can update the predetermined plan if any surprising or unexpected event happens (any changes in the transition function or unexpected punishment). So the model modifies the original plan more often at the beginning of learning and less when it builds the representation of the task (with fewer surprises). Also, the model can plan for a couple of steps until it gets to predefined subgoals (like a hub in the grid world) and then updates its estimates before proceeding to the final goal (or next subgoal), [Huys et al. \(2015\)](#). The current design did not allow us to specify a unique subgoal for different pairs but it can be a direction for the future research.

5.4. Applying heuristic-based models

Our heuristic-based models are inspired by some of our participants' data. Avoiding the greatest loss is useful in curtailing the decision tree search, but it is not optimal. It only guarantees not taking the worst action in the cells surrounding the greatest loss, though, it does not provide any policy in other cells. In other words, it is reduced to the baseline model at cells that are not adjacent to cell 11 and cell 15 and fails to explain our data. In [Huys et al. \(2012\)](#), they examined whether a big loss at early stages of a decision tree could interfere with finding the optimal path later or not. They found that participants are reluctant to take the paths with the greater loss although it might be counterproductive and they could earn a bigger reward later (on that path). In our experiment, we did not ask this question directly, but we found that there is a great avoidance toward the great loss of -75 .¹² Also in the re-planning trials, there were only a few participants (three out of eighty-seven in all 3 experiments) who chose the path with the greatest punishment. We should clarify that in our experiments, participants received 100 points at the goal position regardless of their chosen path. Hence, avoiding large losses in our experiments could not be compared with pruning behavior in [Huys et al. \(2012\)](#).

One way to differentiate between the cells with regular punishments and the cells with great losses is to use different discount parameters to weight future and immediate rewards differently as it was suggested in [Huys et al. \(2012\)](#). Instead, in model 3, we used different rates (e.g. α_{1HS} for the great loss of -75) to estimate the reward at these cells ([Tables 1 and 2 in Supplementary materials](#)). One of the models in this category assigns one special rate, α_{s1} (and possibly one special forgetting factor, α_{s2}) for estimating the reward at cell 11 for a loss of -75 while treating the other cells, with regular or non-regular losses, equally. In [Table 1 of the Supplementary materials](#), we summed up the BIC results of these models. The model which captures participants' tendency to avoid large losses with only two rates is inferior to our proposed model with three distinctive rates for estimating reward function.

There are few participants who chose the shortest path. Because the rewards are probabilistic, this approach can be the worst strategy to apply. If path A1A2 is involved in catching the goal in a pair, it will be the shortest path but costs the greatest loss. This is not consistent with the observed pattern in the majority of our participants' data.

Our final simple heuristic-based model (which was more successful than the others) remembers the last punishment at each cell. Before experiencing the non-regular losses at critical cells (11, 9, 17, 16 and 19 in experiment 2 and 3; 15, 16 and 21 in experiment 1), the model cannot prioritize available paths optimally. But with one (loss) experience, it starts to select actions that lead to greater expected rewards. Remembering the last experience simplifies the reward estimation procedure, but it also promotes wrong predictions after experiencing a regular punishment at critical cells. Thus, this model fails except for the simple case of experiment 1.

5.5. Related research

Unpacking different aspects of planning problems has been an interest to many researchers in different fields. While economic theory provides a normative approach for these problems, results from psychological experiments showed that people used simple heuristics, which were not necessarily optimal to maximize their payoffs in a multi-stage decision tree problem, [Hey and Knoll \(2007\)](#), [Hey and Knoll \(2011\)](#) and [Hotaling and Busemeyer \(2012\)](#). In these experiments, the probability of transition at each choice node was also known to the subjects. They found that, except for a few people who planned ahead (e.g. rational planners), the majority of participants were no-planners or planners with mixed heuristics, [Hey and Knoll \(2011\)](#) and [Hey \(2005\)](#). The latter group tried to simplify the problem by using less information or mixing different strategies (i.e. a combination of local optimization at one level and random guessing at a different level) whereas the rational planners used computations similar to backward induction ([DeGroot, 2005](#)). Yet, some of the strategies mentioned above had certain predictions that were not consistent with the experimental data. For instance, one strong prediction in backward induction is dynamic consistency,¹³ which has been questioned before, [Busemeyer, Weg, Barkan, Li, and Ma \(2000\)](#) and [Johnson and Busemeyer \(2001\)](#). [Hotaling and Busemeyer \(2012\)](#) tried to address these concerns by proposing a new model, the Decision Field Theory-Dynamic model (DFT-D), which was able to distinguish the aforementioned patterns of participants' data and also provided a dynamic account of the decision making process at decision nodes. DFT (among with other sequential sampling models) assumes that the subject accumulates noisy information favoring each choice alternative until the evidence favoring one of the alternatives meets a decision threshold, [Hotaling and Busemeyer \(2012\)](#) and [Khodadadi et al. \(2014\)](#). The DFT-D model is a cognitive-dynamical model that extends DFT for multistage decisions (in planning).

Using decision trees for studying sequential choices has several drawbacks. First, when the number of decision nodes increases (in bigger decision trees), participants find it too difficult to consider all the potential future consequences (of each decision) and thus

¹² The hot stove effect ([Denrell & March, 2001](#)) suggests that avoiding bad outcomes can form information asymmetry between available options. We randomized the starting points and the goals to make sure that participants have unbiased exposure to the environment.

¹³ Dynamic consistency assumes that decision makers follow the original plan that they made for their future choices. In [Johnson and Busemeyer \(2001\)](#), they found that people usually change their choices as they move forward in a decision tree (rather than following the original plan) and this inconsistency increases when they deal with longer decision trees.

might start choosing at random or adopt simple heuristics (Hey, 2005; Huys et al., 2012). Thus, it remains unclear if humans' sub-optimal performance in these tasks is a reflection of their planning behavior or their poor understanding of the environment with which they are dealing. Second, the traditional design of decision trees restricts our design to fix the starting position at the top of the tree; Doing otherwise would result in planning problems being reduced to a single one-time decision. For instance, in a 2-step decision tree, the starting point cannot be placed in the second layer. There are few studies that used 3-step decision trees; however, they fixed the starting position at the top of the decision tree and allowed participants to use Notepad to write comments and remarks (Hey & Knoll, 2011).

To address these issues, many researchers started to use spatial framework because it is more natural for participants for planning and multi-step decision making rather than a decision tree for which they have little or no experience in real life. In our experiment, we minimized spatial reasoning by displaying the environment on the screen along with the participants' position, the goal and all the feasible paths between them. More importantly, unlike (Doeller, King, & Burgess, 2008; Gallistel & Cramer, 1996; Yoshida & Ishii, 2006), we *did not* manipulate viewpoints, locational uncertainty and external cues. Therefore, only decision making and learning theories were directly relevant to our problem.

5.6. Conclusion

There are many theoretical studies in reinforcement learning on how to train an agent in an unknown, complicated environment by focusing on reducing the computational costs and optimizing the search algorithms (Barto & Mahadevan, 2003; Doya, 2000; Ersson & Hu, 2001; Fakhari, Rajagopal, Balakrishnan, & Busemeyer, 2013; Koenig & Likhachev, 2002; Meyer & Filliat, 2003; Matari, 1997; Simmons & Koenig, 1995; Thrun, 1998; Tani, 1996). Few of them, however, have been tested with behavioral data in sequential-choice tasks, Walsh and Anderson (2014). In this study, we were able to simplify one of the real life situations (navigation between two places) in our grid world experiment and evaluated predictions of the reinforcement learning theories with respect to choice data. Our design integrated experience-based decision-making into a classical decision tree problem. We showed that people are capable of revising their plans when an unexpected event occurs and that optimal re-planning requires learning the model of the environment (as in Tolman et al. (1946)).

Authors contributions

P.F. and A.K. designed experiments 1, 2 and 3, conducted the statistical analyses and wrote the paper. J.B. conceived of the study, coordinated the study and helped draft the manuscript.

Funding

The research was funded by the Air Force Grants AFOSR FA 9550-12-1-0397 and FA9550-15-1-0343. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the funding agency.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.cogpsych.2017.12.002>.

References

- Akam, T., Costa, R., & Dayan, P. (2015). Simple plans or sophisticated habits? State, transition and learning interactions in the two-step task. *PLoS Computational Biology*, 11.
- Barto, A. G., & Mahadevan, S. (2003). Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13, 341–379.
- Bertsekas, D. P., & Tsitsiklis, J. N. (1996). (1st ed.). *Neuro-dynamic programming* Vol. 3. Belmont, Mass: Athena Scientific.
- Botvinick, M. M., Niv, Y., & Barto, A. C. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, 113, 262–280.
- Boureau, Y.-L., Sokol-Hessner, P., & Daw, N. D. (2015). Deciding how to decide: Self-control and meta-decision making. *Trends in Cognitive Sciences*, 19, 700–710.
- Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, 44, 171–189.
- Busemeyer, J. R., Weg, E., Barkan, R., Li, X., & Ma, Z. (2000). Dynamic and consequential consistency of choices between paths of decision trees. *Journal of Experimental Psychology: General*, 129, 530–545.
- 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, 1204–1215.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8, 1704–1711.
- Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5, 613–624.
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychological Science*, 27, 848–858.
- DeGroot, M. H. (2005). *Optimal statistical decisions*. John Wiley & Sons.
- Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science*, 12, 523–538.
- Dezfouli, A., & Balleine, B. W. (2012). Habits, action sequences and reinforcement learning. *European Journal of Neuroscience*, 35, 1036–1051.
- Doeller, C. F., King, J. A., & Burgess, N. (2008). Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences*, 105, 5915–5920.

- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, *80*, 312–325.
- Doya, K. (2000). Reinforcement learning in continuous time and space. *Neural Computation*, *12*, 219–245.
- Drner, D., & Schaub, H. (1994). Errors in planning and decision-making and the nature of human information processing. *Applied Psychology*, *43*, 433–453.
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*, 369–409.
- Ersson, T., & Hu, X. (2001). Path planning and navigation of mobile robots in unknown environments. *2001 IEEE/RSJ international conference on intelligent robots and systems, 2001. Proceedings: Vol. 2*, (pp. 858–864). IEEE.
- Fakhari, P., Rajagopal, K., Balakrishnan, S. N., & Busemeyer, J. R. (2013). Quantum inspired reinforcement learning in changing environment. *New Mathematics and Natural Computation*, *09*, 273–294.
- Gallistel, C. R., & Cramer, A. E. (1996). Computations on metric maps in mammals: getting oriented and choosing a multi-destination route. *Journal of Experimental Biology*, *199*, 211–217.
- Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, *71*, 1–6.
- Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A., & Sederberg, P. B. (2012). The successor representation and temporal context. *Neural Computation*, *24*, 1553–1568.
- Gillan, C. M., Otto, A. R., Phelps, E. A., & Daw, N. D. (2015). Model-based learning protects against forming habits. *Cognitive, Affective, & Behavioral Neuroscience*, *15*, 523–536.
- Gonzalez, C., Fakhari, P., & Busemeyer, J. (2017). Dynamic decision making: Learning processes and new research directions. *Human Factors*, *59*(5), 713–721.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534–539.
- Hey, J. D. (2005). Do people (want to) plan? *Scottish Journal of Political Economy*, *52*, 122–138.
- Hey, J. D., & Knoll, J. A. (2007). How far ahead do people plan? *Economics Letters*, *96*, 8–13.
- Hey, J. D., & Knoll, J. A. (2011). Strategies in dynamic decision making: An experimental investigation of the rationality of decision behaviour. *Journal of Economic Psychology*, *32*, 399–409.
- Hotaling, J. M., & Busemeyer, J. R. (2012). DFT-D: A cognitive-dynamical model of dynamic decision making. *Synthese*, *189*, 67–80.
- Hotaling, J. M., Fakhari, P., & Busemeyer, J. R. (2015). Dynamic decision making A2 – Wright, James D. *International encyclopedia of the social & behavioral sciences* (pp. 708–713). (2nd ed.). Oxford: Elsevier.
- Huys, Q. J., Eshel, N., ONions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai trees in your head: How the Pavlovian system sculpts goal-directed choices by pruning decision trees. *PLoS Computational Biology*, *8*.
- Huys, Q. J., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., ... Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences*, *112*, 3098–3103.
- Johnson, J. G., & Busemeyer, J. R. (2001). Multiple-stage decision-making: The effect of planning horizon length on dynamic consistency. *Theory and Decision*, *51*, 217–246.
- Johnson, J. G., & Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Review*, *112*, 841.
- Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology*, *7*.
- Keramati, M., Smittenaar, P., Dolan, R. J., & Dayan, P. (2016). Adaptive integration of habits into depth-limited planning defines a habitual-goaldirected spectrum. *Proceedings of the National Academy of Sciences*, *113*, 12868–12873.
- Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2014). Learning to maximize reward rate: a model based on semi-Markov decision processes. *Frontiers in Neuroscience*, *8*.
- Khodadadi, A., Fakhari, P., & Busemeyer, J. R. (2017). Learning to allocate limited time to decisions with different expected outcomes. *Cognitive Psychology*, *95*, 17–49.
- Koenig, S., & Likhachev, M. (2002). Improved fast replanning for robot navigation in unknown terrain. *IEEE international conference on robotics and automation, 2002. Proceedings. ICRA'02: Vol. 1*, (pp. 968–975). IEEE.
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science*, *28*(9), 1321–1333.
- Kulkarni, T. D., Saedi, A., Gautam, S., & Gershman, S. J. (2016). *Deep successor reinforcement learning*. Available from: arXiv:1606.02396.
- Lee, S. W., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, *81*, 687–699.
- Matari, M. J. (1997). Reinforcement learning in the multi-robot domain. In R. C. Arkin, & G. A. Bekey (Eds.). *Robot colonies* (pp. 73–83). US: Springer.
- Meyer, J.-A., & Filliat, D. (2003). Map-based navigation in mobile robots: II. A review of map-learning and path-planning strategies. *Cognitive Systems Research*, *4*, 283–317.
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N., & Gershman, S. J. (2016). The successor representation in human reinforcement learning. *bioRxiv*.
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*, 12366–12378. <http://dx.doi.org/10.1523/JNEUROSCI.0822-10.2010>.
- 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. *bioRxiv*.
- Schad, D. J., Jnger, E., Sebold, M., Garbusow, M., Bernhardt, N., Javadi, A.-H., ... Rapp, M. A. (2014). Processing speed enhances model-based over model-free reinforcement learning in the presence of high working memory functioning. *Frontiers in Psychology*, *5*, 1450.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*, 461–464.
- Sebold, M., Deserno, L., Nebe, S., Schad, D. J., Garbusow, M., Hgele, C., ... Smolka, M. (2014). Model-based and model-free decisions in alcohol dependence. *Neuropsychobiology*, *70*, 122–131.
- Simmons, R., & Koenig, S. (1995). Probabilistic robot navigation in partially observable environments. In *IJCAI* (Vol. 95, pp. 1080–1087).
- Simon, D. A., & Daw, N. D. (2011). Neural correlates of forward planning in a spatial decision task in humans. *The Journal of Neuroscience*, *31*, 5526–5539.
- Sims, C. R., Neth, H., Jacobs, R. A., & Gray, W. D. (2013). Melioration as rational choice: Sequential decision making in uncertain environments. *Psychological Review*, *120*, 139.
- Skatova, A., Chan, P. A., & Daw, N. D. (2013). Extraversion differentiates between model-based and model-free strategies in a reinforcement learning task.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, *46*, 1004–1017.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction, Vol. 1*. Cambridge: MIT Press.
- Sutton, R. S., Precup, D., & Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, *112*, 181–211.
- Tani, J. (1996). Model-based learning for mobile robot navigation from the dynamical systems perspective. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, *26*, 421–436.
- Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, *99*, 21–71.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189–208.
- Tolman, E. C., & Honzik, C. H. (1930). *Introduction and removal of reward, and maze performance in rats*. University of California Publications in Psychology.
- Tolman, E. C., Ritchie, B., & Kalish, D. (1946). Studies in spatial learning. II. Place learning versus response learning. *Journal of Experimental Psychology*, *36*, 221–229.
- Voicu, H., & Schmajuk, N. (2002). Latent learning, shortcuts and detours: A computational model. *Behavioural Processes*, *59*, 67–86.
- Walsh, M. M., & Anderson, J. R. (2014). Navigating complex decision spaces: Problems and paradigms in sequential choice. *Psychological Bulletin*, *140*, 466.
- Yechiam, E., Erev, I., Yehene, V., & Gopher, D. (2003). Melioration and the transition from touch-typing training to everyday use. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *45*, 671–684.
- Yoshida, W., & Ishii, S. (2006). Resolution of uncertainty in prefrontal cortex. *Neuron*, *50*, 781–789.