

Model-based decision making and model-free learning: A primer

Nicole Drummond¹ | Yael Niv¹

Free will is anything but free. With it comes the onus of choice: not only what to do, but which inner voice to listen to – our “automatic” response system, which some consider “impulsive” or “irrational”, or our supposedly more rational deliberative one. Rather than a devil and angel sitting on our shoulders, research suggests that we have two decision making systems residing in the brain, in our basal ganglia. Neither system is the devil and neither is irrational. They both have our best interests at heart and aim to suggest the best course of action calculated through rational algorithms. However, the algorithms they use are qualitatively different and do not always agree on which action is optimal. The rivalry between habitual, fast action and deliberative, purposeful action is an ongoing one.

In this primer, we first expand on the theoretical and computational basis of each of the two systems and their neural substrates. Given the evidence that the two systems work in tandem, we then discuss how the brain arbitrates between them by asking what are the conditions under which each system is preferentially used. Finally, we discuss a task widely used to study the two modes of decision making in humans.

1 | GOAL DIRECTED VERSUS HABITUAL BEHAVIOR

The idea of two systems of control, one dependent on deliberation and one based on automatic reactions, has roots in the psychological literature on instrumental conditioning and the differentiation between goal directed and habitual behavior. In instrumental conditioning, animals learn which action (or series of actions) will bring about a desired outcome (e.g., rats learning to press a lever for food reward). Outcomes that increase the performance of an action are “reinforcers” – whether appetitive (food increasing lever-pressing) or aversive (a shock increasing escape behavior).

What do animals learn in instrumental conditioning? Early theorists suggested two main alternatives: according to Edward Thorndike, the core of instrumental learning is the formation of associations between stimuli and responses (“S-R” learning), with the reinforcer “stamping in” these associations, but not being part of the learned construct. In contrast with this behaviorist view, Edward Tolman, foreman of the cognitive revolution, argued that animals learn more elaborate cognitive maps. In these, the learned association is between actions and their outcomes, broadly construed (“A→O” learning), and the association is predicated on their causal contingency (which we denote by →).

As highlighted by Robert Rescorla, Ruth Colwill, Tony Dickinson and Bernard Balleine, these competing models make contradictory predictions: Tolman’s A→O model suggests that behavior should be sensitive to: 1) changes in the value of the outcome (e.g., if the outcome is no longer desired due to satiety, the rat will decrease its rate of pressing), and 2) changes in the contingency between the action and the outcome (e.g., if the action no longer leads to the outcome, the rat will cease to perform the action). Thorndike’s theory, on the other hand, predicts that learned

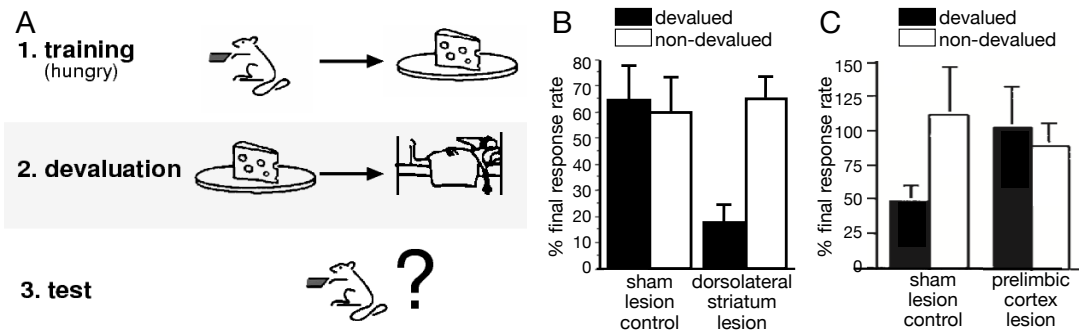


FIGURE 1 Outcome-devaluation test of goal directed versus habitual control. **A.** Rats are first trained to press a lever (the prototypical instrumental action) in order to obtain food. Training may be for a short duration (e.g., 4 daily session) or extensive (15 sessions). In phase 2 the outcome is devalued, for instance by injecting lithium chloride after consumption, which causes the animal to feel ill. The effects of the manipulation are then tested in a third phase by assessing how eagerly rats will press the lever. The test is performed in extinction (without new outcomes) so as to assess only the products of previous learning: will the value of leverpressing have changed, given the experience in the second phase? Habitual control predicts that it will not (as neither the stimulus nor the response from training were experienced in the devaluation phase, and therefore the S-R association should be intact). In contrast, if the animal is using goal-directed control, it will realize that the action leads to a no-longer-valuable outcome, and therefore the action is moot. Figure adapted from Dickinson & Balleine (1994). **B.** After extensive training (in the first phase), rats continue pressing the lever even after devaluation (sham lesion control). However, lesions of the dorsolateral striatum reinstate goal-directed control (lesion group). Blocking dopamine transmission to the dorsolateral striatum during training also shows a similar pattern (not shown here). Plotted is the rate of lever pressing at test, as compared to the last session of training. Figure adapted from Yin, Knowlton & Balleine (2004). **C.** In contrast, after only moderate training, rats' leverpressing still diminishes after devaluation (sham lesion control). This devaluation-sensitivity, however, requires an intact prelimbic cortex (lesion group), as well as other brain areas comprising the same cortico-basal-ganglia loop (not shown). Figure adapted from Killcross & Coutureau (2003).

actions will be impervious to such environmental (or internal) changes, as once they were stamped in, they will not be “stamped out” except by aversive outcomes. Indeed, not all decisions to press a lever are created equal: for an animal to make a deliberative, purposeful choice, it must choose the action because it knows it will result in a food pellet, and the food pellet is desirable to it (as in Tolman's model). However, as per Thorndike, this is not the only reason an animal may choose to press a lever.

What do animals do when the environment changes in such ways? It depends. In particular, early findings showed that rats that were trained extensively to press the lever indeed persisted in pressing for an undesirable outcome, as predicted by stimulus-response learning. However, rats who were only trained for a moderate amount decreased their lever-pressing substantially if the outcome the lever-presses led to was devalued, consistent with action→outcome learning (Figure 1). Indeed, the distinction between “goal-directed” and “habitual” instrumental actions hinges on whether said actions maintain adaptability in response to changes in outcome and/or contingency is considered, or are impervious to such changes, respectively. Mirroring the two behavioral patterns, in the brain there is also evidence for two separate systems for action control, as we will detail in section 4.

Humans, too, show both habitual and goal-directed modes of action: after moving house, we can flexibly plan a new route home from work; however, if we are talking to a friend, or listening to the radio, we might habitually make the wrong turn to our old address instead. As in rats, habits are bred through repetition – but they are not context independent or “irrational.” As stimulus-response associations, habits can be exquisitely context dependent (that is,

dependent on local stimuli; you can walk home habitually, but you will turn left at the correct intersection), and these actions have habitized precisely because they have consistently led to favorable outcomes in the past. In addition, as the example here suggests, goal-directed action requires more cognitively effortful deliberation, and action selection may give way to habits if those cognitive resources are otherwise occupied, as we discuss in section 3.

2 | MODEL-BASED VERSUS MODEL-FREE REINFORCEMENT LEARNING

Machine learning theory offers another, computational, framework for understanding these two decision-making systems. Specifically, Nathaniel Daw and colleagues used the framework of reinforcement learning (RL) to provide a formal description of the habitual vs deliberative choice mechanisms in terms of model-based and model-free action selection. RL defines the interaction between a learning agent and its environment in terms of decision making (i.e., action selection) within a Markov Decision Process (MDP). The MDP is a description of the environment as a set of states that transition to each other probabilistically (possibly dependent on an agent's actions), and a set of rewards that may be available in each state. Most RL algorithms predicate action choices on "values:" the expected sum of future rewards if a specific action were to be taken at a specific state.

In model-based RL, action values are computed by mentally simulating the consequences of said actions within a representational map of the world that includes the states, the transition probabilities between them, and the rewards in each state – a "world model," hence model-based RL. The decision-making agent can learn a model of the environment from experience, by observing and keeping track of sequences of states, and the rewards available in each state (for instance, implemented through learning from state prediction errors). To decide what action to take, the agent can then search through different trajectories of the mental map ("planning" or mental simulation in a "forward model") and choose a path that yields the best estimated long-run outcome. A benefit of making decisions based on a stored mental map is that the map can easily be updated when changes in the environment occur, meaning that actions that are computed using the map will also flexibly change as needed. So if, for example, the agent learns that a particular route home is unexpectedly blocked, it can immediately re-plan accordingly. However, searching through all possible action trajectories before making a choice is computationally costly, and in some situations, prohibitively so. Such planning and re-planning also takes time, which may not be available for certain decisions that are time-sensitive. This model-based method of action selection maps directly to Tolman's action-outcome model of learning: the world model includes the contingency relationships between actions and outcomes in the state transitions, as well as the subjective worth of the different outcomes. Planning within the model therefore realizes the Dickinsonian goal-directed combination of contingency and outcome information in order to make a decision.

In contrast, in model-free RL, decisions are made without a world model. Instead, values for different actions are learned directly, through trial and error interaction with the environment every time an action is performed. Decisions then rely on those values that embody a cached summary of past experience. In particular, on each trial, the agent observes the current state, performs an action based on the estimated values of different actions at that state, and observes the results: the outcome it receives, and the state it transitions to. The agent then computes a reward prediction error: the difference between the outcomes obtained and the amount of expected reward (the original estimated value of the action taken). Importantly, because values reflect expected future rewards (not just the immediate reward), to estimate the outcomes obtained, the agent sums up the immediate reward and the value of the new state (i.e., the rewards expected from that state onward). The prediction error is then used to determine the magnitude and direction of an update of the value of the action taken: if the prediction error is positive (i.e., the action yielded better outcomes than expected), the action value is increased, and vice versa if the prediction error is negative.

Through this process of updating values trial-by-trial, prediction errors are minimized and the scalar value of each action eventually approximates the true future expected rewards. Therefore, action selection based on model-free RL is optimal once learning converges. However, because the values reflect the accumulation of previous rewards and transitions, this method of learning is less flexible in adapting to abrupt change. Known changes in the environment (for instance, hearing about the road block) cannot be incorporated into learned values without directly experiencing prediction errors. This is not just an algorithmic quirk: because cached values embody long-term expected rewards, that is, expected sums of rewards in future states, it is not possible to cleave out the portion of the value that should be changed given a change in one of the state contingencies or predicted outcomes. The information that would allow such dissection – the breakdown of the different components of the total value – is simply not stored in this type of algorithm. Model-free action selection therefore formalizes Thorndike's stimulus-response model: responding is driven by the learned value of a response at a specific state, without explicating the expected outcome.

According to the computational theory, the main benefit of model-free decision making is that action selection is computationally easy: one can simply choose the action with the highest cached value. However, a large amount of trial-and-error experience is required for values to become good estimates of future rewards. Additionally, as mentioned, the cached values are inflexible and trial-and-error learning is needed to update them. Model-based decision making is therefore advantageous in its adaptability in response to changes in outcome values or state transitions. However, it is either computationally costly due to the search over trajectories of possible action sequences, or inaccurate if not all sequences are mentally simulated. As such, neither system should dominate in *all* situations, and if the brain uses the systems optimally, each is expected to have a role to play in decision making.

As an aside, there are additional decision making algorithms that the brain can – and likely does – employ. For instance, one can select actions by sampling distinct episodic memories (in machine learning this would be akin to sampling a past trajectory and repeating it, with the sampling potentially prioritizing trajectories that led to more rewards). This algorithm would be most appropriate when very few experiences with a task exist, and both model-based and model-free RL cannot reliably estimate action values.

3 | UNCERTAINTY-BASED ARBITRATION BETWEEN THE TWO SYSTEMS

If the brain has two mechanisms to estimate the value of different actions, they may not always agree. This then creates a new decision making problem: which system to use? Experiments testing for the behavioral hallmarks of goal-directed versus habitual behavior reveal conditions under which each strategy is employed: early in learning, devaluation results confirm that behavior is goal directed, that is, consistent with a model-based RL strategy for computing action values. However, once behavior is repeated more extensively ("overtrained"), it turns habitual (see control groups in Figure 1B,C). This is consistent with the idea that model-free learning is slower and less efficient than model-based learning, and so the model-free system should be used only after considerable experience with the situation.

There are also conditions under which behavior seems to remain model-based (that is, sensitive to outcome devaluation) even through extensive training. In particular, when choosing between several different actions that lead to different outcomes (e.g., two levers, one leading to delivery of cherry-flavored food pellets and the other to banana-flavored pellets), behavior seems to remain goal-directed for longer. This may be because the choice between different (fairly equivalent) outcomes halves the experience with each action (so there are fewer updates to the model-free value of that action) or because the choice between two actions that are equally favorable in terms of model-free values utilizes the more accurate forward-modeling model-based system in an attempt to disambiguate the options.

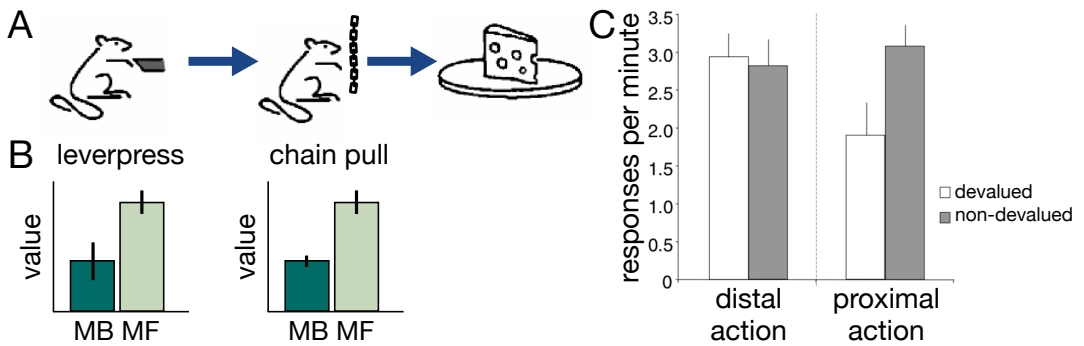


FIGURE 2 Uncertainty-based arbitration between model-free and model-based control. **A.** The sequence of actions required from the rats in Balleine et al. (2005). **B.** Each system estimates the value of pressing the lever, as well as its own confidence in that value (error bar). Here, after devaluation, the model-free system has more confidence in its value of pressing the lever, while the model-based system is more accurate in evaluating the utility of chain pulling. **C.** As a result, after devaluation rats press the lever (action distal from reward) but refrain from subsequently pulling the chain (action proximal to reward). Figure adapted from Balleine et al. (2005).

One thing that these results show—and that other findings confirm—is that animals can alternate between using each of the decision-making systems within a given local scenario. We are not habitual or goal directed, we can be either at any point in time, and we can switch between decision-making modes even within a chain of actions. This was most starkly demonstrated by Balleine and colleagues in an experiment where rats were required to perform a sequence of actions consisting of a lever-press followed by a chain pull in order to obtain a reward outcome (Figure 2A). Animals were trained on the sequence when hungry, but then tested (in extinction) when sated, as a means to devalue the food reward. Results showed that the action distal from the reward (the lever press) was not sensitive to the motivational shift, suggesting habitual control, whereas the action more proximal to the outcome was sensitive to the lower value of the reward, supporting goal-directed control (Figure 2C). That is, the rats continued to press the lever, but substantially reduced their tendency to continue to pull on the chain. One explanation for this is that proximity to the outcome allows more accurate forward-planning (and thus promotes model-based control) whereas such planning is too costly for choices that are farther away from the end-point of a sequence of actions, hence the reliance on model-free values in that case. This is similar to what expert chess players report – in the middle of a game, they rarely plan many steps ahead. Instead, they rely on well-honed knowledge of what board positions are more favorable in the long run (e.g., it is in general advantageous to have control of the center of the board). However, close to the end of the game, all relevant sequences of actions by the opponent are considered through mental planning.

Indeed, these conditions for goal-directed versus habitual control of animal behavior were what led Nathaniel Daw and colleagues to first suggest that goal-directed behavior can be mapped on to model-based RL, whereas habitual behavior is a hallmark of model-free RL. Moreover, to arbitrate between the two decision-making systems, they suggested that each system computes and reports not only values for different actions, but also its confidence in those values (i.e., standard error bars or confidence intervals on the estimated values). Arbitration between the systems can then utilize their relative confidence: at each point in time, the brain should trust the system that is most sure of its value estimate (Figure 2B).

One can interpret the behavioral findings in this light: the model-free system calculates its confidence as proportional to the amount of experience with this particular action in this situation (state). Like in experimental science

where larger sample sizes decrease the error in estimating the mean of a population, more experience with an action in a certain situation decreases the error in estimating the future rewards that can be expected if that action is chosen. The model-based system, in contrast, loses precision the farther into the future it has to simulate, given that the model it plans with is not fully accurate. The more iterations of using the model to figure out the future rewards contingent on an action (i.e., value), the more errors will accumulate in the estimate. Additionally, if the planning tree has many branches (that is, there are many possible future courses of action to evaluate and average over), computational noise or working memory limitations (or both) can inject noise into the value estimate. As such, the model-based system might put less confidence in estimates that required more computation, due to either depth or width of the simulated options.

This could explain why the action more distal from the reward outcome was deemed habitual, or model free, in Balleine and colleagues' action sequence study, even while the action proximal to reward was goal directed/model based. The more the animal needs to simulate forward, the more uncertain the model-based system, whereas the extensive training process in that experiment was sufficient to reduce the uncertainty of the model-free system for both actions (Figure 2B). According to the theory, arbitration between the two systems is done on a decision-by-decision basis: each system provides value estimates for each of the possible actions, complete with error bars on these estimates, and for each action, an arbitration system uses the value from the most confident system (or, alternatively, averages the two values inversely weighted by precision) to make the final action choice.

We note here that we have used the term "uncertainty" not in the sense of "risk," which is sometimes called uncertainty in the field of neuroeconomics, but rather in the sense of "ambiguity." Risk is a property of the stochasticity of the world – tossing a fair coin for a \$1 reward if heads, and \$1 loss if tails involves high risk as we are maximally uncertain about the outcome of the toss. However, this risk can be rationally and accurately averaged over to produce a mean expectation of \$0. The uncertainty we are referring to is, instead, akin to ambiguity about the fairness of the coin, that is, not knowing the probability of reward itself. Reinforcement learning naturally averages over risk as it learns expected values (in the statistical notion of "expectation"). However, it is susceptible to ambiguity – the whole learning process is aimed at reducing ambiguity and estimating the expectation more and more accurately. As such, the hypothesis above, regarding each system estimating its own uncertainty, is different from ideas about RL tracking risk and not only expected values.

4 | TWO DISTINCT BUT INTERACTING NEURAL SYSTEMS

Supporting the behavioral results described above, neuroscientific studies have shown that the brain harbors two largely parallel decision making circuits. Habitual control using model-free RL is associated with a decision-making loop from cortical afferents, through the dorsolateral striatum (putamen in primates), the ventrolateral thalamus, and back to cortex. Learning in cortico-striatal synapses in this loop is modulated by dopaminergic signals that convey prediction errors, as predicted by the theory. A parallel loop including the prelimbic cortex (dmPFC in primates), the dorsomedial striatum (caudate in primates), ventromedial thalamus and back to cortex is thought to mediate goal-directed control and model-based computations. Learning in the latter loop seems to be dopamine-independent (though recent evidence suggests that this may be an oversimplified view), but relies on inputs from the hippocampus. Lesioning or inactivating either of these two loops will revert behavioral control to the remaining, intact system (Figure 1).

Although we have been discussing model-based and model-free RL as if they are exclusive, the chess analogy highlights that the two systems might better work in tandem. For instance, one can use model-based planning to

search several steps forward in a simulated tree of future possibilities, computing action values based on cached model-free values of the end-points of the search rather than simulating the game until its end. Similarly, early pruning of some branches of the search tree can be determined based on cached values. This may result in failure to exploit sequences of actions that initially lead to negative outcomes but eventually to winning the game (e.g., sacrificing the queen for a later check mate). However, pruning can allow the decision maker to focus cognitive resources on simulating action sequences that are more likely to be advantageous in the long run. Of course, as the model-free system improves its estimates, pruning will be even more optimal.

In sum, although the distinction between two competing may seem simplistic, evidence for more mixed modes of decision making does not necessarily mean that the two systems do not exist as separate algorithms. It does suggest that competition and collaboration can take place at the level of individual actions and sweeps of forward planning in the model. Indeed, the two cortico-basal-ganglia loops are actually interconnected. This connection may underlie both the competition and cooperation between the two systems – the locus of arbitration between model-free and model-based decisions is not yet known. Additionally, the neurotransmitter norepinephrine (NE) has been associated with signaling environmental change (known as “unexpected uncertainty”) – which should trigger a shift from habitual, model-free behavior back to deliberative, goal-directed model-based reasoning. High levels of NE may therefore modulate the arbitration between the two systems, in addition to increasing learning rates for both the world model and the cached values.

5 | A PARADIGM FOR STUDYING MODEL-BASED VERSUS MODEL-FREE DECISION MAKING IN HUMANS

To test for model-based or model-free behavior in animals, the studies above used a single change in the environment (most often: a change in the value of the outcome), and a subsequent test. However, the effects of planning in a model versus prediction-error based stimulus-response learning can be observed even without such abrupt changes in the environment. In particular, Nathaniel Daw has taken advantage of the specific algorithmic assumptions of model-based versus model-free learning to develop a task that can quantify, on a continuous scale, the extent to which human behavior conforms with either mechanism.

This paradigm is known as the “two-step task” (Figure 3). On each trial, an initial choice between two options leads probabilistically to either of two second-stage states (“planets” in Figure 3, from an adaptation of the original task for use with children). At the second step, another choice between two options (“aliens”) is made. Rewards at the second step are probabilistic and change slowly and independently for each of the four aliens. As a result, to maximize reward, participants must constantly learn which alien is best in each planet, and which planet is best overall. The transition probabilities from first to second step are fixed (and sometimes pre-trained): each first-level option leads to one planet 70% of the time (“common” transition), and to the other planet 30% of the time (“rare” transition).

Teasing apart model-based and model-free choice strategies occurs by observing how rewards received after a rare transition impact first-step choices on subsequent trials. In particular, after a reward is obtained in the second step, model-free RL suggests that the first-step choice for that trial should be repeated in the subsequent trial, as this action was accompanied by a positive prediction error at the end of the two-step chain of actions. In contrast, model-based RL updates the probability of reward for that alien, and by taking into account the transition structure of the task, knows that the best way to reach this now-more-valuable planet is by choosing the opposite first-step action. Therefore, model-free RL predicts that choice in the first step should, in general, be consistent with feedback on the previous trial: repeat actions that were rewarded, avoid those that were not. Model-based RL, in contrast,

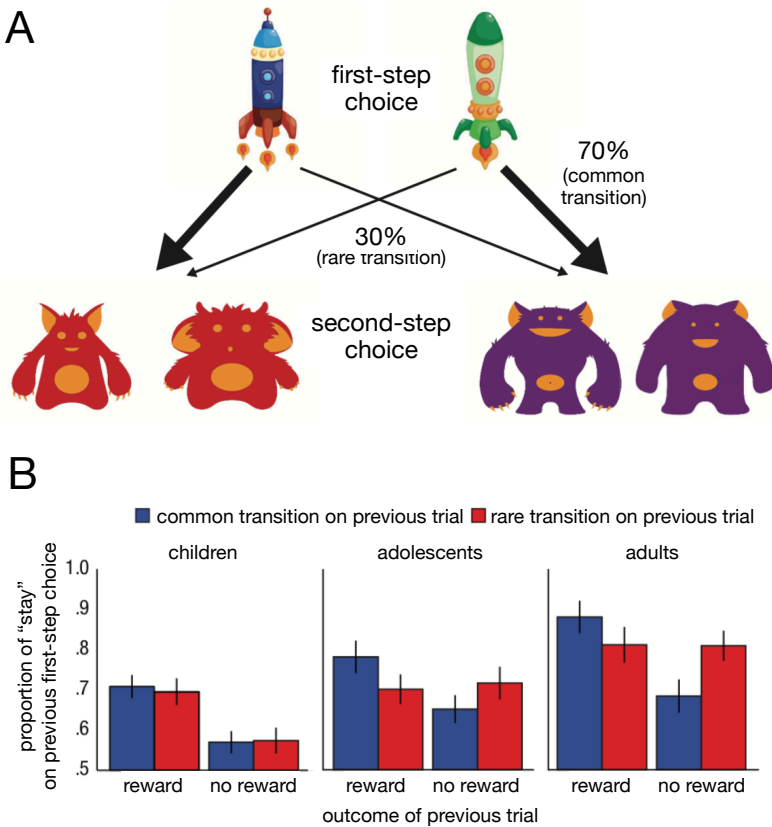


FIGURE 3 The "two-step task." A. In a version of the task used to study how model-based and model-free control changes over development, in the first stage, participants are presented with two options (spaceships) and asked to pick one. Each spaceship leads to one "planet" with 70% chance and to the other with 30% chance. In the second stage, participants meet two "aliens" on the planet they reached, and can choose which alien to ask for space treasure to give. Aliens are associated with different (and constantly changing) probabilities of having space-treasure reward to give. **B.** The interaction between transition type and the effects of reward on subsequent choice that is characteristic of model-based learning emerges during adolescence and continues in adulthood. Figures adapted from Decker et al. (2016).

takes into account whether the outcome followed a rare transition or a common one to determine whether to repeat or avoid the previous action.

The 2-step task has been leveraged to test the effects of a host of manipulations on the trade-off between model-based and model-free decision-making. For instance, Ross Otto and colleagues showed that a cold-pressor stress manipulation decreases model-based decisions in the task, with subjects with better working memory showing more resilience in the face of stress. In a separate study, they showed that working memory load due to a dual task decreased model-based decision making, and that this competition between the two systems could be controlled on a trial-by-trial basis by imposing working memory constraints that restrict the availability of cognitive resources. These results can be explained by assuming that limiting computational resources (due to stress, or a concurrent task) renders the model-based system more uncertain in its (less rigorous and more noisy) calculations, and therefore control

reverts to the model-free system. Developmentally, Catherine Hartley and colleagues have shown that young children show dominance of model-free learning, with model-based strategies emerging in adolescence and increasing through adulthood (Figure 3B). Rodent variants of the task have also been developed, and are being used to dissect the neural mechanisms of both methods of decision-making, capitalizing on the ability to measure both model-based and model-free influences throughout many sessions of data collection.

6 | CONCLUSION

Returning to the image of two competing decision-making systems sitting on our shoulders, it is now clear why neither is the devil. Values estimated by either model-based or model-free RL are equally rational—though not necessarily equal. Indeed, because these systems use different algorithms to compute action values, they are differently susceptible to situational elements that can make decisions suboptimal or seemingly irrational. This may be why the brain employs both algorithms: when possible, it uses forward planning in a sophisticated, prospective, model-based system. However, the high cognitive demands of such decision making requires exclusive use of precious resources, preventing other concurrent tasks. Luckily, our good-old habitual, model-free RL system is ever ready to jump in, and to allow us to maintain the illusion that we can multi-task.

Acknowledgements

This work was funded in part by the National Institute on Drug Abuse R01DA042065 and by the John Templeton Foundation. We are grateful to Ross Otto, Mehdi Keramati, Nathaniel Daw and Geoffrey Schoenbaum for valuable feedback on the manuscript.

Further Reading

- Balleine, B. W.; Garner, C.; Gonzalez, F. & Dickinson, A. (1995). Motivational control of heterogeneous instrumental chains *Journal of Experimental Psychology: Animal Behavior Processes*, 21, 203-217
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204-1215.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12), 1704-1711.
- 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(6), 848-858.
- Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior*, 22(1), 1-18.
- Dickinson, A., Balleine, B., Watt, A., Gonzalez, F., & Boakes, R. A. (1995). Motivational control after extended instrumental training. *Animal Learning & Behavior*, 23(2), 197-206.
- Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral cortex*, 13(4), 400-408.
- Lengyel, M., & Dayan, P. (2008). Hippocampal contributions to control: the third way. In *Advances in neural*

information processing systems (pp. 889-896).

Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013a). The curse of planning: dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychological science*, 24(5), 751-761.

Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013b). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences*, 110(52), 20941-20946.

Yin, H. H., Ostlund, S. B., Knowlton, B. J., & Balleine, B. W. (2005). The role of the dorsomedial striatum in instrumental conditioning. *European Journal of Neuroscience*, 22(2), 513-523.