

# Anxiety, Avoidance, and Sequential Evaluation

Samuel Zorowitz<sup>1</sup>, Ida Momennejad<sup>2</sup>, and Nathaniel D. Daw<sup>1,3</sup>

<sup>1</sup>Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey, USA

<sup>2</sup>Department of Biomedical Engineering, Columbia University, New York, New York, USA

<sup>3</sup>Department of Psychology, Princeton University, Princeton, New Jersey, USA

an open access journal



**Keywords:** anxiety, avoidance, fear generalization, decision theory

## ABSTRACT

Anxiety disorders are characterized by a range of aberrations in the processing of and response to threat, but there is little clarity what core pathogenesis might underlie these symptoms. Here we propose that a particular set of unrealistically pessimistic assumptions can distort an agent's behavior and underlie a host of seemingly disparate anxiety symptoms. We formalize this hypothesis in a decision-theoretic analysis of maladaptive avoidance and a reinforcement learning model, which shows how a localized bias in beliefs can formally explain a range of phenomena related to anxiety. The core observation, implicit in standard decision-theoretic accounts of sequential evaluation, is that the potential for avoidance should be protective: If danger can be avoided later, it poses less threat now. We show how a violation of this assumption—via a pessimistic, false belief that later avoidance will be unsuccessful—leads to a characteristic, excessive propagation of fear and avoidance to situations far antecedent of threat. This single deviation can explain a range of features of anxious behavior, including exaggerated threat appraisals, fear generalization, and persistent avoidance. Simulations of the model reproduce laboratory demonstrations of abnormal decision-making in anxiety, including in situations of approach–avoid conflict and planning to avoid losses. The model also ties together a number of other seemingly disjoint phenomena in anxious disorders. For instance, learning under the pessimistic bias captures a hypothesis about the role of anxiety in the later development of depression. The bias itself offers a new formalization of classic insights from the psychiatric literature about the central role of maladaptive beliefs about control and self-efficacy in anxiety. This perspective also extends previous computational accounts of beliefs about control in mood disorders, which neglected the sequential aspects of choice.

## INTRODUCTION

Though anxiety disorders differ in their particular symptomology, and in the content and situations that elicit symptoms, they all are similarly characterized by aberrations in the processing of and response to threat (American Psychiatric Association, 2013). In particular, at least three symptoms manifest across many of the anxiety disorders. First, anxiety is associated with exaggerated threat appraisal, or a bias toward evaluating threat as disproportionately greater in likelihood and severity than is warranted (Clark & Beck, 2011). Second, anxiety is also associated with fear generalization, wherein the primary threat becomes associated with increasingly distal locations, events, and thoughts (Dymond, Dunsmoor, Vervliet, Roche, & Hermans, 2015).

Citation: Zorowitz, S., Momennejad, I., & Daw, N. D. (2020). Anxiety, avoidance, and sequential evaluation. *Computational Psychiatry*, 4, 1–17. [https://doi.org/10.1162/cpsy\\_a\\_00026](https://doi.org/10.1162/cpsy_a_00026)

DOI: [https://doi.org/10.1162/cpsy\\_a\\_00026](https://doi.org/10.1162/cpsy_a_00026)

Received: 4 August 2019  
Accepted: 29 December 2019

Competing Interests: The authors declare no conflict of interest.

Corresponding Author:  
Samuel Zorowitz  
[szorow1@gmail.com](mailto:szorow1@gmail.com)

Copyright: © 2020  
Computational Psychiatry, Inc. and the  
Massachusetts Institute of Technology  
Published under a Creative Commons  
Attribution 4.0 International  
(CC BY 4.0) license

Finally, anxiety is associated with persistent avoidance behavior, which often occurs well in advance of the materialization of actual threat (Arnaudova, Kindt, Fanselow, & Beckers, 2017). (Here we distinguish between *avoidance* and *escape*, where the former describes actions taken to prevent the onset of threats, whereas the latter describes defensive responses to proximal threat). Excessive avoidance behaviors are an especially harmful aspect of anxiety disorders, both because they interfere with daily life and because they indirectly maintain anxiety by preventing learning from the nonoccurrence of perceived threats. Though laboratory studies of decision-making and learning in anxious populations have corroborated these clinical observations (Aylward et al., 2019; Harlé, Guo, Zhang, Paulus, & Angela, 2017; Norbury, Robbins, & Seymour, 2018), none offer an explanation as to the root of these symptoms.

These symptoms are particularly puzzling from a decision-theoretic perspective (Huys, Daw, & Dayan, 2015). In many circumstances, distant threat should not impinge upon decision-making in the present. Indeed, we argue that fear and avoidance of situations far in the future violate the basic logic of evaluation over sequential trajectories of action. This is because avoidance is by nature protective: The ability to successfully avoid danger in the future means an agent need not also do so now. For instance, cars endanger pedestrians but can be reliably avoided by following traffic signals; given that, staying indoors offers little or no additional protection from accidents. This is an instance of a fundamental property of evaluation in sequential decision-making: The value of present action turns fundamentally on assumptions about subsequent events, which importantly include the agent's own subsequent choices. Typically, it is appropriate to assume that an agent will continue to make good (i.e., reward-maximizing/harm-minimizing) choices down the line, and that good choices at the current stage should therefore anticipate this.

This line of reasoning hints that a fundamental aberration in anxiety disorders may relate to this assumption, which otherwise should preclude the spread of threat to antecedent situations and subsequent excessive avoidance. Indeed, anxiety disorders are associated with pessimistic beliefs about the future (Clark & Beck, 2011). Clinically and subclinically anxious individuals judge future threat as more likely than do nonanxious individuals (Butler & Mathews, 1983, 1987; MacLeod & Byrne, 1996). Importantly, the development and maintenance of clinical anxiety is strongly tied to diminished perceived control (Bandura & Adams, 1977; Barlow, 2002; Gallagher, Bentley, & Barlow, 2014) such that anxious individuals are more likely to endorse the belief that they are unlikely or unable to mitigate future threat. Indeed, lack of belief in one's ability to successfully navigate future danger is associated with anxiety (Davey, Jubb, & Cameron, 1996; Dugas, Freeston, & Ladouceur, 1997), and an increased belief in perceived control over threat is correlated with symptom reduction across the family of anxiety disorders (Gallagher, Naragon-Gainey, & Brown, 2014).

Here we develop this idea—that symptoms of anxiety may arise from misbeliefs about future avoidance—into a formal model of evaluation under pessimistic assumptions about future choices. We show that a single, localized deviation from normative evaluation can explain a surprising range of features of anxious behavior, including exaggerated threat appraisal, fear generalization, and persistent avoidance. This account also offers a new formalization of classic insights from the psychiatric literature about the central role of beliefs about control and self-efficacy in anxiety (Bandura & Adams, 1977; Barlow, 2002). Specifically, we show through simulation that a model with a misbelief about the reliability of future self-action gives rise to a number of characteristic symptoms and laboratory results concerning anxiety. Our perspective also extends previous computational accounts of beliefs about control in mood disorders (e.g., Huys & Dayan, 2009), which neglected the sequential aspects of choice.

## MODEL DESCRIPTION

We model anxious decision-making in the context of Markov decision processes (MDPs). A standard normative assumption is that agents attempt to optimize the expected cumulative discounted reward:


$$Q^\pi(s, a) = r(s, a) + \gamma \sum_{s'} p(s' | s, a) \sum_{a'} \pi(a' | s') Q^\pi(s', a'). \quad (1)$$

For any particular state-action  $(s, a)$ , this is necessarily defined relative to a policy  $\pi(a' | s')$  specifying the assumed distribution of *future* choices. (It is also defined relative to a discount rate  $\gamma$ , controlling the present value of future outcomes). The return can be optimized self-consistently under the assumption that the agent makes the return-maximizing choice at each step in the future, leading to the familiar expression for the optimal values,

$$Q^*(s, a) = r(s, a) + \gamma \sum_{s'} p(s' | s, a) \max_{a'} Q^*(s', a'). \quad (2)$$

The “max” operator in Equation 2 yields a fundamental asymmetry between approach and avoidance, as illustrated in Figure 1. It formalizes the assumption that the agent makes the return-maximizing choice at each step. Through this operation, opportunities for reward (which maximize the argument) propagate recursively to earlier steps, but avoidable dangers do not, because the return-maximizing action is to avoid. (To the extent obtaining reward or avoiding harm are only imperfectly achievable, this propagation and attenuation at each step are graded according to success probability, but the basic asymmetry remains). This principle is highlighted in a toy MDP (Figure 1): a deterministic open gridworld with two terminal states, a rewarding state and a punishing one (Figure 1A). The optimal state values  $V^* = \max_a Q^*(s, a)$  (Figure 1B) reflect a “mountain” of opportunity propagating recursively from the reward. Conversely, because harm is avoidable in this environment, its negative value is contained: All states (even those adjacent to threat) represent the positive opportunity for reward.

Although the return-maximizing assumption self-consistently defines optimal behavior, an agent need not be restricted to it (Symmonds, Bossaerts, & Dolan, 2010) and might in principle anticipate encountering danger under different (e.g., pathological) assumptions about the future. For example, an agent may expect to fail to take the correct protective actions in later states [i.e., to use a suboptimal  $\pi(a | s)$ ] or may believe that the world’s future dynamics do not guarantee reliable avoidance even so [i.e., under stochastic or adversarial transition dynamics



**Figure 1. An example open field environment.** A) A simple deterministic gridworld with two terminal states: one rewarding (blue) and one aversive (red). B–D) States colored by their value under different levels of pessimism, with arrows showing an optimal trajectory. In B, for an optimistic agent ( $w = 1$ ), all states (other than the harmful state) take on positive value. In C, for a pessimistic agent ( $w = 0.5$ ), negative value spreads from the source to antecedent states. In D, with increasing pessimism ( $w = 0$ ), the extent of the spread grows worse, and the return-optimizing trajectory becomes more distorted and avoidant. (Parameters:  $\gamma = 0.95$ ).

$P(s' | s, a)$ ]. Consider an agent who has such pessimistic expectations about dangerous events at future steps. Note that assumptions of this sort, even if incorrect, can serve adaptive purposes. In general, pessimistic assumptions can help to ensure robustness and safety under uncertain or even adversarial scenarios (Garcia & Fernandez, 2015). Related work in reinforcement learning shows how computing returns with pessimistic predictions can help to quantify variability in outcomes (i.e., to learn different points in the distribution of possible returns; Bellemare, Dabney, & Munos, 2017), which is one way to explicitly and tunably take account of risk tolerance. Indeed, it is common in machine learning theory to optimize outcomes under worst-case assumptions.

Here we propose unrealistically pessimistic assumptions as a root cause of many anxious symptoms. Such pessimism can be encoded either in the policy,  $\pi(a | s)$ , or transition probabilities,  $p(s' | s, a)$ . These, respectively, correspond to misbeliefs about one's own avoidance actions or the environment's responses to them, a point to which we return in discussion. Here, for concreteness, we focus on distortions in the policy. In particular, we adopt the  $\beta$ -pessimistic value function from Gaskett (2003) to define state-action value in expectation over a mixture of the best and worst action:

$$Q^w(s, a) = r(s, a) + \gamma \sum_{s'} p(s' | s, a) \left( w \max_{a'} Q^w(s', a') + (1 - w) \min_{a'} Q^w(s', a') \right). \quad (3)$$

The weight parameter  $w$  controls the degree of pessimism. An optimistic agent ( $w = 1$ ) expects in the future to act fully in accordance with its preferences, whereas a pessimistic agent ( $w = 0$ ) expects to act contrary to its preferences. Importantly, this belief is false under the model; that is, we assume, at each step, that the agent actually chooses by maximizing the values given by Equation 3. But these values are always computed under the assumption that the agent will then not maximize at all later steps. (Another way to say this is that we assume choices at each step are optimal, but only under the assumption that later choices will not be). Figures 1C and 1D illustrate the consequences of different levels of pessimism for valuation in an example gridworld. The value of threat propagates, with increasing distance across the state space, due to an increasing expectation that the agent may fail to avoid threat in the future.

This simple simulation reflects a localized violation of the core decision theoretic assumption of future return-optimizing action. Therefore the model's behavior already echoes several core symptoms of anxiety disorders. Namely, the pessimistic agents in Figures 1C and 1D exhibit exaggerated threat appraisals (otherwise neutral states unrealistically signal danger), generalization of fear (threat value spreads across the gridworld), and persistent avoidance (early on, the agent takes paths that maintain increasing distances from threat). Importantly, as we elaborate in the following pages, this deviation from the usual assumptions is supported by prominent clinical theories of anxiety.

## SIMULATIONS

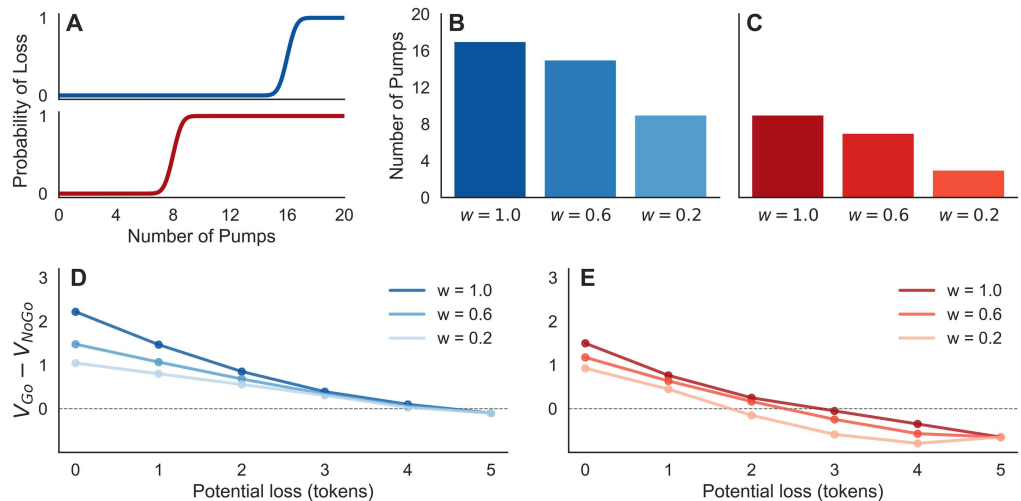
In what follows, we demonstrate through simulation how our simple model can account for anxious behavior in laboratory-based studies of *sequential learning* and decision-making. (Because in our model, anxiety arises through biased sequential evaluation, we will not address one-step bandit tasks, where others have reported learning deficits associated with anxiety, e.g., Aylward et al., 2019; Harlé et al., 2017). We also show that our model is consistent with clinical theory describing the transition from clinical anxiety to depression. Unless otherwise noted, state and action values under varying degrees of pessimism were solved for using

the value iteration dynamic programming method (Sutton & Barto, 2018). All simulations were implemented in the Python programming language, and the code is publicly available at <https://github.com/ndawlab/seqanx>.

### Approach–Avoidance Conflict

One behavioral finding characteristic of anxiety disorders is unbalanced processing of approach–avoidance conflict (Aupperle & Martin, 2010). Anxious individuals are more likely to forgo potential gains to avoid potential danger. Many of the disruptions anxiety causes to everyday functioning (e.g., avoiding social obligations for fear of possible social humiliation) can be understood in these terms. As such, many have sought to probe and measure this behavior in the laboratory. For instance, in the balloon analog risk task (BART; Lejuez et al., 2002), participants attempt to earn money by pumping virtual balloons. With each pump, the balloon inflates and money is earned, but so too increases the chance that the balloon will pop and the accumulated earnings will be lost. At any point in a trial, a participant may cash out, banking the money earned and ending the trial. Anxiety is correlated with fewer pumps of the balloon and earlier cash-outs in the BART (Maner et al., 2007; Ramírez, Ortega, & Del Paso, 2015).

As shown in Figure 2, our model easily accommodates this result. Whereas optimistic agents pump until the marginal gain of a pump no longer offsets the chance of the balloon bursting, optimal choice under increasingly pessimistic (i.e., anxious) assumptions cashes out progressively earlier—similar to empirical findings (Maner et al., 2007; Ramírez et al., 2015). This is because it anticipates and avoids future errors in choice, which would otherwise result in the balloon popping. Our model can analogously explain other manifestations of biased



**Figure 2. The balloon analog risk task (BART; Lejuez et al., 2002).** A) The risk of balloon burst (point loss) increases with each pump and does so earlier for the high-risk (red) than low-risk (blue) balloons. (For full rules of the task, see Methods). The optimal policy (number of pumps) under increasingly pessimistic valuation is presented for B) low-risk and C) high-risk balloons. The optimistic agent ( $w = 1$ ) prefers a policy reflecting the true environmental risk. The moderate ( $w = 0.6$ ) and strongly ( $w = 0.2$ ) pessimistic agents cash out earlier, as is observed in anxious individuals. D, E) In the sleeping predator task, the risk of loss is constant, but the cost of loss still increases as more rewards are gathered. The value of reward pursuit under increasingly pessimistic valuation is presented for scenarios with D) low risk and E) high risk of predator awakening. The relative value of approach (vs. avoid) decreases with loss amount and threat level, and more so under pessimistic assumptions. (Parameters:  $\gamma = 1.0$ ).

approach–avoid conflict in anxiety, such as in the predator avoidance task (Fung, Qi, Hassabis, Daw, & Mobbs, 2019).

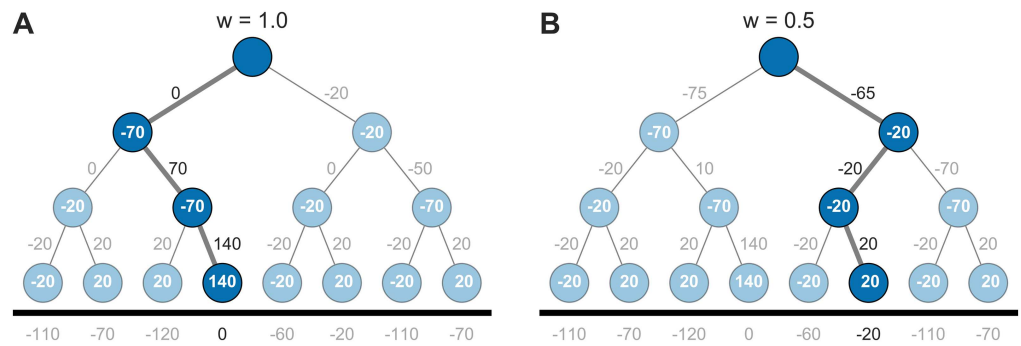
A unique prediction of the model (because we assume optimal choice under the assumption of *future* suboptimality) is that bias should arise only when beliefs about *future* avoidance are involved, rather than direct conflict between immediate impulses. Recent data (Fung et al., 2019) using a predator avoidance task, analogous to the BART, support this view. In this task, increasing trait anxiety predicted earlier escape (analogous to cash-out in the BART) for slow predators (for whom future decisions to escape were a relevant consideration) but not for fast ones (who would attack immediately, mooted consideration of future steps).

Relatedly, the model can also capture findings of increased behavioral inhibition, measured as prolonged response times, in anxious individuals under threat (Bach, 2015). In the behavioral inhibition task, participants seek tokens adjacent to a virtual “sleeping predator,” which are all lost if the predator awakes. Though in this variant, the risk of predation is constant throughout a trial (rather than increasing, as in the BART), the potential loss from capture still increases with each token collected. Bach found that participants are slower to collect tokens as this potential loss increases and that this slowing is enhanced by subclinical anxiety. We can capture this effect in our model by noting that the relative value of approach compared to avoidance is reduced as potential loss and the risk that the predator awakes increase (Figures 2D and 2E). These effects are amplified under more pessimistic (i.e., in our model, more anxious) assumptions about future actions. Thus, as before, anxious pessimism in our model produces greater and earlier choice of avoidance, analogous to earlier cash-out in the BART. To extend this effect to reaction times, we adopt the further, standard assumption that actions (here approach) are slower when their relative values compared to alternatives are lower (Oud et al., 2016). In this case, the model captures behavioral inhibition (slower responses as threat increases) and its enhancement by anxiety, as measured by Bach (2015). (Note that we need not assume that the coupling of reaction times to action value spread is due to difficulty in decision formation per se, which Bach argues against: It may, for instance, reflect Pavlovian initiation biases; Niv, Daw, Joel, & Dayan, 2007).

### ***Aversive Pruning***

Another laboratory phenomenon associated with anxiety is *aversive pruning* in planning (Huys et al., 2012; Lally et al., 2017). This refers to the idea that when evaluating future action trajectories in a sequential task like chess, people are resource limited, cannot evaluate all possible options, and must selectively consider certain paths and neglect others. One proposal for how people accomplish this is aversive pruning (Huys et al., 2012), wherein choice sequences involving large losses are discarded from further evaluation. An example of aversive pruning is shown in Figure 3A. Although the optimal choice in the decision tree is to weather an initial large loss (e.g.,  $-70$ ) to reap the large gain that follows, people tend to disfavor this path, suggesting they prune it and consequently neglect the later gain. The degree of such pruning correlates, depending on the study, with subclinical depressive (Huys et al., 2012) or anxiety (Lally et al., 2017) symptoms.

Our model predicts this result (Figure 3B) as specifically linked to our model of anxious pessimism, though for a somewhat different reason than in Huys’s original modeling. In our model, pessimistic (anxious) agents neglect large gains deeper in the tree, not because they fail to consider them (here we assume full evaluation of the Bellman equation), but because with increasing anxiety, they increasingly expect the potential of choosing incorrectly afterward, thus failing to recoup the loss (and, mathematically, probabilistically pruning the better



**Figure 3.** The decision tree environment (Huys et al., 2012; Lally et al., 2017). A) An optimistic agent ( $w = 1$ ) prefers the optimal loss-minimizing policy through the initial large loss (left branch). B) A pessimistic agent ( $w = 0.5$ ) comes to prefer the branch without the large loss so as to avoid being unable to recoup the large initial loss. One-step rewards (or costs) are presented in each state; the net value  $Q$  on each path is shown numerically. (Parameters:  $\gamma = 1.0$ ).

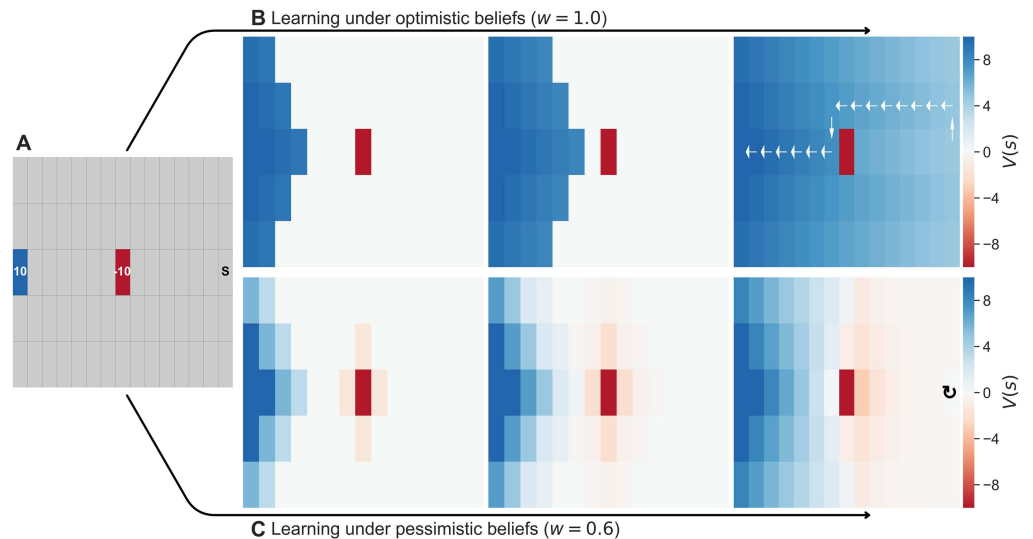
branches). Future research could use slight variants in the decision trees to tease apart these different interpretations, for example, by comparing decision trees that differ only in what follows the large initial loss. Under a model of aversive pruning, such a change should not impact the proportion of agents selecting the left branch; in contrast, our model predicts that choice should parametrically increase with the extent of the amelioration.

#### The Anxiety–Depression Transition

So far, we have considered only the asymptotic preferences implied by our pessimistic value function, which we computed directly through value iteration. But we can also consider the process of learning under this value function (e.g., by variants of Q-learning, Sutton & Barto, 2018, or DYNA, Russek, Momennejad, Botvinick, Gershman, & Daw, 2017; Sutton, 1991, using the  $\beta$ -pessimistic return). The dynamics of such learning may speak to the progression of symptoms.

Of note in this respect, anxiety and depression are highly comorbid, with almost half of individuals with a lifetime depression diagnosis also diagnosed with an anxiety disorder (Kessler et al., 2015). One notable proposal is that this association often (though by no means exclusively) arises longitudinally, in particular, that clinical anxiety precedes certain types of depression (Alloy, Kelly, Mineka, & Clements, 1990; Jacobson & Newman, 2014). The idea, in brief, is that uncertainty in one's ability in the face of future threat results in anxiety and avoidance behaviors. Persistent avoidance, in turn, begets forgone reward, leading ultimately to a belief that reward is unobtainable and subsequently to depression. This informal story can be captured by simulations of learning in our model (Figure 4) in environments like that of Figure 1. Over the course of learning, the penumbra of negative value under pessimistic assumptions spreads gradually throughout the environment. This can in turn lead the agent to expect no reward and, also echoing the anergic symptoms of depression, forgo action altogether. This last point in particular dovetails nicely with theoretical accounts of the anergic aspects of depression (Huys et al., 2015), which point out that low experienced reward rates should in decision-theoretic accounts lead to reduced response vigor (Niv et al., 2007), leading to a potentially self-reinforcing downward spiral.

In addition to suggesting one explanation for the comorbidity of anxiety and depression, our model hints at a reason for the longevity and recurrence of anxiety disorders even with treatment.



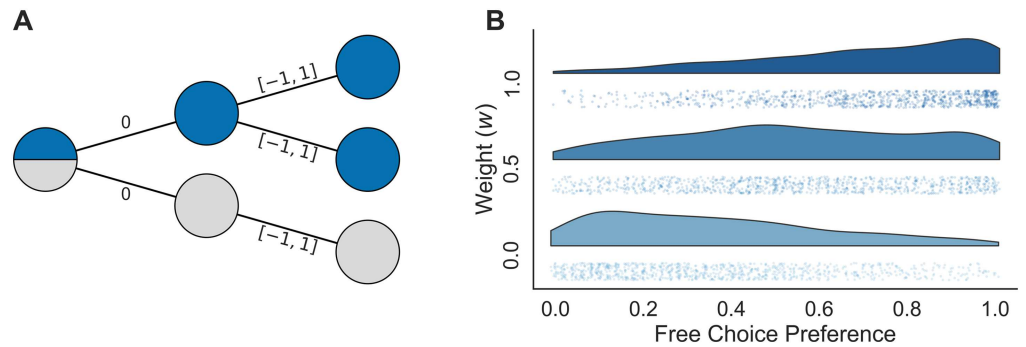
**Figure 4. An example of the anxiety–depression transition.** A) A simple deterministic gridworld with two terminal states: one rewarding (blue) and one aversive (red). B, C) The development of value expectancies over three steps of learning, for two levels of pessimism. States are colored by their values under different levels of pessimism, with arrows showing an optimal trajectory. In B, for an optimistic agent ( $w = 1$ ), all states (other than the harmful state) take on positive value with learning. In C, for a pessimistic agent ( $w = 0.6$ ), negative value spreads from the source to antecedent states. As a result of avoidance, the agent learns reward is unobtainable and develops anergic symptoms (i.e., forgoes action). (Parameters:  $\gamma = 0.95$ ).

Because pessimistic expectations allow for threat value to spread to states and actions far antecedent of the primary danger (e.g., Figure 1D), it would accordingly also take a great many steps of iterative learning to correct all these exaggerated appraisals of threat. Frustratingly, these biased estimates of value may still remain even after a misbelief in the efficacy of future action is corrected for in a course of therapy. This phenomenon (similar to failures of model-free reinforcement learning algorithms to adjust to reward revaluation without extensive relearning; Daw, Niv, & Dayan, 2005) may offer at least a partial answer to a classic puzzle in pathological avoidance, that is, why it is so resistant to extinction (Moutoussis, Shahar, Hauser, & Dolan, 2018), and to the unfortunately high rates of anxiety recurrence following treatment (Pittig, Treanor, LeBeau, & Craske, 2018).

#### Free Choice Premium

Finally, the model also offers a novel prediction tying anxious beliefs to a classic, but hitherto separate, phenomenon known as the *free choice premium*. This refers to the finding that, all else being equal, people tend to treat choice as itself valuable, that is, choices that lead to more choice opportunities in the future are preferred to those that lead to fewer future choice opportunities. A free choice premium has been observed in multiple behavioral experiments (Leotti, Iyengar, & Ochsner, 2010; Ly, Wang, Bhanji, & Delgado, 2019). A variant of a free choice premium paradigm from two previous studies (Leotti & Delgado, 2011, 2014) is presented in Figure 5A. In the task, participants repeatedly choose between a free choice option, allowing for an additional future choice, and a fixed choice option. Importantly, both choices lead to identical, stochastic outcomes (e.g., 50–50 chance of [1, -1]). Empirical studies have found that human subjects (from a general, healthy population) prefer the free choice option despite it conferring no additional benefits relative to the fixed choice option.





**Figure 5. The free choice premium task (Leotti & Delgado, 2011, 2014) with equal chance of outcome  $R \in [-1, 1]$ .** A) Nonanxious participants exhibit a preference for the free choice option (blue) despite it conferring no benefit over the fixed choice option (gray). B) Pessimistic agents show an attenuated free choice bias. The fractional preference for the free choice option over the simulated experiment is shown for three populations of subjects with levels of pessimism  $w$  ( $y$ -axis; each dot represents an individual simulated agent, and the smoothed density of free choice bias for each pessimism level is also shown). (Parameters: Q-learning with  $\gamma = 1.0$ , and inverse temperature,  $\beta$ , increased from 1 to 15 over 100 episodes).

On one account (Ly et al., 2019), this free choice preference directly and specifically reflects the assumption about sequential choice whose violation we argue is core to anxiety, that is, that the agent will make reward-maximizing choices in the future. Under such an optimistic assumption (and given noisy and imperfect knowledge about the value of each option, due to learning), additional options are valuable in the sense that free choice can be expected to exploit the best among them. Namely, the maximum over several noisy values is, in general, larger than a single option from the same distribution. Our proposal makes the novel prediction that if, as we hypothesize, anxiety reflects a violation of this optimistic assumption, then anxious individuals will exhibit a diminished or reversed free choice bias, as shown in simulation in Figure 5B. Future empirical research will be required to test this prediction.

## DISCUSSION

Central to anxiety disorders are symptoms including exaggerated threat appraisal, threat generalization, and excessive avoidance (Arnaudova et al., 2017; Clark & Beck, 2011; Dymond et al., 2015). We have presented a simple computational account suggesting how a single underlying pessimistic misbelief can give rise to these aberrations in learning and choice. We use a reinforcement learning approach in which undue pessimism results in maladaptive policy. Specifically, we show how a failure to believe in the reliability of one's future actions can effectively backpropagate negative value across states of the environment. This process results in a range of inferences and behaviors resembling those observed in clinical anxiety. Though it is by no means a complete account of anxiety, our account ties together a surprisingly wide range of symptoms of anxiety disorders.

We are not the first to propose a formal theory of control in psychiatry using MDPs. Huys and Dayan (2009) also provided a computational account of learned helplessness through simple models of one-step action–outcome contingencies. Our accounts differ particularly in our exclusive focus on control in the sequential setting, which Huys and Dayan did not address. Indeed, we propose that ultimately the key to anxiety is precisely the way in which evaluation in sequential tasks is necessarily reliant on expectations about future choice and events. Similarly,

research and modeling by Bishop and colleagues (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015; Gagne, Dayan, & Bishop, 2018) also taking a decision-theoretic approach has stressed the importance of uncertainty as a core feature of anxiety. Specifically, they have described uncertainty as inherently aversive in anxiety and have presented models of how uncertainty may be increased in anxiety (e.g., aberrant processing of environmental volatility). The present work is compatible with deficits in processing uncertainty and might instead be viewed as an attempt to unpack why uncertainty is aversive—because, in our view, it is resolved (i.e., marginalized) under pessimistic distributional assumptions. As for control, we extend this view to focus on how uncertainty is resolved in the sequential setting and also to zero in on particular instances of uncertainty (about future actions and some other options discussed next) and misbeliefs about them that, we argue, are particularly consequential.

For concreteness, we formalized pessimistic assumptions in terms of only one of several variants of a more general family of models, but we do not mean this restriction as a substantive claim. In particular, we focused on the agent's beliefs about its own future actions: expecting failure in handling or avoiding future threat. However, this is just one of several different pessimistic misbeliefs that could satisfy the basic logic of our model and produce similar symptoms. These other beliefs need not be mutually exclusive, though they might reflect different cognitive routes to symptoms that would, in turn, imply different psychotherapeutic strategies.

For instance, one variant of the model is suggested by the observation that Equation 1 is also computed in expectation over the anticipated future environmental dynamics  $p(s' | s, a)$ . Thus pessimism can alternatively be encoded in this distribution; for example, a false belief that the world's response to one's choices is unpredictable or adversarial. Because the Bellman equation for the return averages over this distribution in addition to the choice policy at each step, and because an unpredictable environment also reduces the efficacy of avoidance, either formulation can produce ultimately similar results in our simulations here. A third variant of our model arises from uncertainty about the current state  $s$  of the environment. Although we have taken it as fully observed, if the world state is only partly known, then this distribution too must be averaged out in evaluating each action (Kaelbling, Littman, & Cassandra, 1998), and here also a pessimistic skew will propagate the expectation of danger and result in exaggerated avoidance (Paulus & Yu, 2012). In summary, pessimistic resolution of several different varieties of uncertainty (e.g., about future action, environmental dynamics, or environmental state) could each produce similar symptoms for analogous reasons. However, from the perspective of cognitive theories of anxiety, these represent quite different maladaptive beliefs: a key difference that may be relevant in guiding treatment (especially cognitive psychotherapies aimed at ameliorating the false beliefs) of a host of anxiety disorders.

Particularly due to the way it encompasses several such variants, our account formalizes a long-standing range of theory on the role of control in anxiety. Central to many prominent cognitive theories of anxiety in the psychiatric literature is a perceived lack of control. For example, self-efficacy theory (Bandura & Adams, 1977) and the triple vulnerabilities model (Barlow, 2002) both posit that a reduced belief in the ability to effectively respond to future threat is involved in the genesis and maintenance of clinical anxiety. In contrast, and focused less on the self, the learned helplessness theories of anxiety (Alloy et al., 1990) claim that clinical anxiety results from an uncertain belief in the controllability of the environment, such that future threat cannot be effectively mitigated or avoided. As we note earlier, the present model and analysis (though simulated here in terms of self-efficacy) can accommodate either variant and show how they relate to one another as members of a more general family of accounts.

The possibility of multiple anxious phenotypes, each characterized by unique but not mutually exclusive beliefs, suggests the need for behavioral assays designed to isolate and interrogate such biases. One such task is the free choice premium paradigm described earlier, which captures pessimism (or optimism) about one's own choices. An analogous task might measure pessimistic expectations about environmental state transition probabilities. For instance, this could be accomplished with a variant of a sequential decision-making task that requires subjects to learn the transition structure of a multistep decision tree (Gläscher, Daw, Dayan, & O'Doherty, 2010; Lee, Shimojo, & O'Doherty, 2014) and make choices to gather rewards in it. Pessimistic expectations about environmental state transitions would bias choices in this type of task. Individually, these tasks could test our hypothesis that either sort of misbelief is associated with symptoms of anxiety; they could also be compared to one another (and to more detailed self-report assessments of beliefs about control or self-efficacy) to investigate potential heterogeneity across patients in the antecedent of anxiety.

Importantly, although we have considered the most pathological cases of pessimism, these may reflect the exaggerative extremes of an otherwise adaptive evaluation strategy. Traditionally, the goal of reinforcement learning algorithms is to find a reward-maximizing policy with respect to the expectation (average) of returns. However, depending on one's risk attitude and uncertainty about the environment (e.g., if there is potential for catastrophic loss), it may instead be preferable to learn a policy with respect to an alternative and more pessimistic objective function, similar to the one considered here. Accordingly, returns and agent behaviors similar to ours arise in previous research on learning risk-sensitive and robust policies (Bellmare et al., 2017; Chow, Tamar, Mannor, & Pavone, 2015; Morimura, Sugiyama, Kashima, Hachiya, & Tanaka, 2012).

We have centered our discussion at Marr's (1982) computational level: on beliefs and their consequences in terms of action values. We have so far remained agnostic as to how, algorithmically or mechanistically, these misbeliefs are implemented in the brain (Friston, Stephan, Montague, & Dolan, 2014). Importantly, the brain is believed to contain multiple distinct mechanisms for evaluating actions (e.g., model-based and model-free learning; Daw et al., 2005; Huys et al., 2015), and pessimistic beliefs might play out either differentially or similarly through each of these mechanisms. One promising possibility is that these symptoms mainly reflect aberrations in model-based planning (Huys et al., 2015), that is, explicitly evaluating actions by mentally simulating possible trajectories. Recent work has suggested that this process may be accomplished by mentally "replaying" individual potential trajectories (Matar & Daw, 2018; Momennejad, Otto, Daw, & Norman, 2018). In this setting, the biases we suggest would amount to overcontemplating, or overweighting, certain pessimistic trajectories (Hunter, Meer, Gillan, Hsu, & Daw, 2019). Such a bias might be detectable using neuroimaging, as a change in which types of events that tend to be replayed (Ambrose, Pfeiffer, & Foster, 2016; Momennejad et al., 2018). Such a biased replay process, in turn, may also correspond to worry and rumination. Indeed, in line with the present results, chronic worry is associated with reduced perceived control, diminished belief in self-efficacy in response to threat, and exaggerated threat appraisal (Berenbaum, 2010). This suggests that clinical anxiety may in part result from planning processes gone awry.

It is important to note that the present model may not describe all anxiety disorders with equal accuracy. Indeed, our analysis of pessimistic sequential evaluation is, by definition, a model of prospective cognition. Thus the present results are more likely to accurately describe the anxiety disorders that primarily involve aberrations in future-oriented cognitive processes, such as generalized and social anxiety disorders. Naturally, the present model can

account neither for the compulsive behaviors of obsessive-compulsive disorder (OCD) nor the memory disturbances of posttraumatic stress disorder (PTSD). That said, recent clinical studies have suggested that diminished perceived control is a vulnerability factor common to all anxiety disorders, including OCD and PTSD (Gallagher, Bentley, & Barlow, 2014; Gallagher, Naragon-Gainey, & Brown, 2014). Importantly, aspects of other psychiatric disorders that involve future-oriented misbeliefs, worry, and avoidance behaviors (e.g., eating disorders; Konstantellou, Campbell, Eisler, Simic, & Treasure, 2011) may similarly be well described by the current account. Indeed, and much more speculatively, the model may also have implications for bipolar disorders. The onset of manic symptoms is associated with overoptimistic perceived control (Alloy, Abramson, Walshaw, & Neeren, 2006), and our same model and basic reasoning (though now envisioning excessively optimistic rather than excessively pessimistic beliefs) may help to explain how this bias may translate into risk-seeking behavior and dysregulated goal pursuit. As such, the present results are transdiagnostic and not limited to one particular diagnosis.

Finally, the relationship between evaluation, planning, and neural replay discussed herein suggests potential future work that might help to bring this model into contact with the more memory-related aspects of PTSD. For instance, the same replay processes that can be used to evaluate actions can also, in theory, update predictive representations, cognitive maps, or models of the environment (Russek et al., 2017), such as the successor representation (Dayan, 1993; Momennejad et al., 2017). If so, similar biases in replay could result in not only aberrant avoidance behavior but also progressive, aberrant remodeling of world models or cognitive maps, an observation that may connect to the rich and complex set of issues on memory involvement in PTSD.

## METHODS

We model anxious decision-making in the context of MDPs. Tasks were modeled as deterministic, infinite-horizon, discrete-time environments. Some (detailed later) were modeled with discounted returns  $\gamma < 1$ . All simulations were implemented in the Python programming language, and the code is publicly available at <https://github.com/ndawlab/seqanx>.

For all but the free choice premium task, we defined state-action values,  $Q(s, a)$ , in accordance with our modified, pessimistic Bellman Equation 3, reproduced for convenience:

$$Q^w(s, a) = r(s, a) + \gamma \sum_{s'} p(s' | s, a) \left( w \max_{a'} Q^w(s', a') + (1 - w) \min_{a'} Q^w(s', a') \right).$$

Here  $Q$ -values were solved for directly through value iteration (Sutton & Barto, 2018). By contrast,  $Q$ -values in the free choice task were computed using the  $\beta$ -pessimism temporal-difference learning algorithm (Gaskett, 2003), where the reward prediction error is defined as

$$\delta^w = r(s, a) + \gamma \left( w \max_{a'} Q^w(s', a') + (1 - w) \min_{a'} Q^w(s', a') \right) - Q^w(s, a)$$

and the update rule is defined as:

$$Q^w(s, a) \leftarrow Q^w(s, a) + \eta \cdot \delta^w,$$

where  $\eta$  is a learning rate. The parameterizations of the MDP environments and learning algorithms are next specified in turn.

#### **Toy MDP/Anxiety-to-Depression Transition**

Both the Toy MDP and anxiety–depression transition simulations were performed using simple gridworlds. Both environments involved only two nonzero states, one rewarding ( $r = 10$ ) and one aversive ( $r = -10$ ). In both environments, we solved for the discounted, asymptomatic  $Q$ -values using value iteration with  $\gamma = 0.95$ . In the toy gridworld, we performed simulations under pessimistic assumptions,  $w \in [1.0, 0.5, 0.0]$ . In the transition gridworld, we performed simulations under pessimistic assumptions,  $w \in [1.0, 0.6]$ . To highlight the effects of learning, we took “snapshots” of  $Q$ -values prior to asymptote, three and five steps into value iteration.

#### **Balloon Analog Risk Task/Predator Avoidance Task**

The balloon analog risk task (Lejuez et al., 2002) has participants inflate a virtual balloon for points. Earnings rise with each pump, but so too does the risk of the balloon popping and subsequent point loss. Unbeknownst to participants, the number of pumps before balloon pop is predefined and drawn randomly from some distribution (e.g., uniform, normal, exponential), where the mean controls the risk (i.e., average number of pumps before point loss).

Here we modeled the BART as an undiscounted MDP with 20 states, where transitioning to each successive state yielded  $r = 1$ . The only available actions were to transition to the next state (e.g.,  $S_1 \rightarrow S_2, S_2 \rightarrow S_3$ ) or end the episode (i.e., cash out). With each act to move to the next state, there was some probability of transitioning to a bad terminal state (i.e., balloon pop) with reward equal to the negative equivalent of accumulated gain thus far. The probability of this bad transition was modeled using normal density function, with parameters  $\mathcal{N}(16, 0.5)$  for low risk and  $\mathcal{N}(8, 0.5)$  for high risk. The asymptomatic  $Q$ -values were solved for using value iteration for both the low-risk and high-risk conditions under pessimistic assumptions,  $w \in [1.0, 0.6, 0.2]$ .

The predator avoidance task (Fung et al., 2019) can be analogously modeled. There a virtual predator approaches participants over discrete time steps while participants “forage.” For every time step the participant remains (i.e., does not flee), points are accumulated. However, if the participant is caught, all earnings are lost, and an additional penalty is received. Thus the predator avoidance task bears a striking resemblance to the BART; foraging is equivalent to virtual pumps, and fleeing is equivalent to cashing out.

#### **Behavioral Inhibition Task**

In the behavioral inhibition task (or sleeping predator task; Bach, 2015), participants collect virtual tokens while evading capture from a virtual predator. Unlike the BART, the risk in the behavioral inhibition task (i.e., the predator “waking up”) is constant. However, the cost of capture increases as sequential tokens are collected.

True to the original, we model the task as an undiscounted MDP with six states, where transitioning to each successive state yielded  $r = 1$ . Identical to the BART, the only available actions were to transition to the next state or to end the episode (i.e., avoid the predator). With each act to move to the next state, there was a constant probability of transitioning to a bad terminal state (i.e., capture by the predator), with reward equal to the negative equivalent of accumulated gain thus far. The probability of bad transition was defined as  $p = 0.10$  for low risk and  $p = 0.15$  for high risk, based on the objective risk probabilities in the empirical

experiment (Bach, 2015). The asymptomatic  $Q$ -values were solved for using value iteration for both the low-risk and high-risk conditions under pessimistic assumptions,  $w \in [1.0, 0.6, 0.2]$ .

### **Aversive Pruning**

In the aversive pruning task (Huys et al., 2012; Lally et al., 2017), participants learn to navigate a six-state graphworld, where each state is directly connected to only two other states. Each state is associated with some reward (cost), and participants must plan trajectories through the state-space so as to maximize reward (minimize cost).

We model the task as an undiscounted MDP where the original 6-node network has been “unravelling” into a 15-state decision tree. This is equivalent to having the participant start in one state and plan to make three actions. The rewards associated with transitioning to each state were taken directly from Huys et al., 2012, and Lally et al., 2017. The asymptomatic  $Q$ -values were solved for using value iteration for both the low-risk and high-risk conditions under pessimistic assumptions,  $w \in [1.0, 0.5]$ .

### **Free Choice Task**

In the free choice task (Leotti & Delgado, 2011), participants complete a series of two-stage trials. In the first stage, they select between a free choice option, allowing them to make an additional choice in the second stage, and a fixed choice, permitting no choice in the second stage. In the second stage, participants select between one of two bandits (free choice) or are randomly assigned a bandit (fixed choice). Importantly, all bandits pay out under an identical reward distribution.

We model the task as an undiscounted MDP with a six-state decision tree structure. In the free choice branch, agents are able to select between two terminal bandits; in the fixed choice branch, agents can choose only one terminal bandit. All bandits pay out identically, in this case, randomly in the set,  $r \in [-1, 1]$ . Learned  $Q$ -values were computed using a pessimistic temporal difference learning algorithm, with learning rate  $\eta = 0.4$ . Each simulated agent learned the values of each action over 100 trials, with an increasing, logarithmically spaced inverse temperature in the range of  $\beta \in [0, 15]$ . (Inverse temperature was gradually increased over learning to facilitate exploration of choice options). The resulting fraction of free choices made over the last 50 trials was stored for 1,000 simulated agents, run separately for pessimistic assumption,  $w \in [1.0, 0.5, 0.0]$ .

## **AUTHOR CONTRIBUTIONS**

Samuel Zorowitz: Conceptualization: Lead; Formal analysis: Lead; Software: Lead; Visualization: Lead; Writing - original draft: Equal; Writing - review & editing: Equal. Ida Momennejad: Conceptualization: Supporting; Formal analysis: Supporting; Writing - original draft: Supporting; Writing - review & editing: Supporting. Nathaniel D. Daw: Conceptualization: Supporting; Formal analysis: Supporting; Supervision: Lead; Writing - original draft: Equal; Writing - review & editing: Equal.

## **FUNDING INFORMATION**

Samuel Zorowitz, National Science Foundation (<https://doi.org/10.13039/100000001>). Nathaniel D. Daw, National Institute on Drug Abuse (<https://doi.org/10.13039/100000026>), Award ID: 1R01DA038891.

## ACKNOWLEDGMENTS

The authors are grateful to Zeb Kurth-Nelson and Will Dabney for helpful discussion.

## REFERENCES

- Alloy, L. B., Abramson, L. Y., Walshaw, P. D., & Neeren, A. M. (2006). Cognitive vulnerability to unipolar and bipolar mood disorders. *Journal of Social and Clinical Psychology, 25*(7), 726–754. <https://doi.org/10.1521/jscp.2006.25.7.726>
- Alloy, L. B., Kelly, K. A., Mineka, S., & Clements, C. M. (1990). Comorbidity in anxiety and depressive disorders: A helplessness/hopelessness perspective. In J. D. Maser & C. R. Cloninger (Eds.), *Comorbidity of anxiety and mood disorders*. Washington, DC: American Psychiatric Press.
- Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron, 91*(5), 1124–1136. <https://doi.org/10.1016/j.neuron.2016.07.047>
- American Psychiatric Association. (2013). *Diagnostic and statistical manual of mental disorders (DSM-5)*. Washington, DC: Author.
- Arnudova, I., Kindt, M., Fanselow, M., & Beckers, T. (2017). Pathways towards the proliferation of avoidance in anxiety and implications for treatment. *Behaviour Research and Therapy, 96*, 3–13. <https://doi.org/10.1016/j.brat.2017.04.004>
- Aupperle, R., & Martin, P. P. (2010). Neural systems underlying approach and avoidance in anxiety disorders. *Dialogues in Clinical Neuroscience, 12*(4), 517–531.
- Aylward, J., Valton, V., Ahn, W.-Y., Bond, R. L., Dayan, P., Roiser, J. P., & Robinson, O. J. (2019). Altered learning under uncertainty in unmedicated mood and anxiety disorders. *Nature Human Behaviour, 1*. <https://doi.org/10.1038/s41562-019-0628-0>
- Bach, D. R. (2015). Anxiety-like behavioural inhibition is normative under environmental threat-reward correlations. *PLoS Computational Biology, 11*(12), e1004646. <https://doi.org/10.1371/journal.pcbi.1004646>
- Bandura, A., & Adams, N. E. (1977). Analysis of self-efficacy theory of behavioral change. *Cognitive Therapy and Research, 1*(4), 287–310. <https://doi.org/10.1007/BF01663995>
- Barlow, D. H. (2002). *Anxiety and its disorders: The nature and treatment of anxiety and panic*. New York, NY: Guilford Press.
- Bellemare, M. G., Dabney, W., & Munos, R. (2017). A distributional perspective on reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning* (Vol. 70, pp. 449–458). N.P.: JMLR.
- Berenbaum, H. (2010). An initiation-termination two-phase model of worrying. *Clinical Psychology Review, 30*(8), 962–975. <https://doi.org/10.1016/j.cpr.2010.06.011>
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience, 18*(4), 590–596. <https://doi.org/10.1038/nn.3961>
- Butler, G., & Mathews, A. (1983). Cognitive processes in anxiety. *Advances in Behaviour Research and Therapy, 5*(1), 51–62. [https://doi.org/10.1016/0146-6402\(83\)90015-2](https://doi.org/10.1016/0146-6402(83)90015-2)
- Butler, G., & Mathews, A. (1987). Anticipatory anxiety and risk perception. *Cognitive Therapy and Research, 11*(5), 551–565. <https://doi.org/10.1007/BF01183858>
- Chow, Y., Tamar, A., Mannor, S., & Pavone, M. (2015). Risk-sensitive and robust decision-making: A CVAR optimization approach. In *Advances in neural information processing systems* (pp. 1522–1530). Montreal, Quebec: NeurIPS.
- Clark, D. A., & Beck, A. T. (2011). *Cognitive therapy of anxiety disorders: Science and practice*. New York, NY: Guilford Press.
- Davey, G. C., Jubb, M., & Cameron, C. (1996). Catastrophic worrying as a function of changes in problem-solving confidence. *Cognitive Therapy and Research, 20*(4), 333–344. <https://doi.org/10.1007/BF02228037>
- 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. <https://doi.org/10.1038/nn1560>
- Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural Computation, 5*(4), 613–624. <https://doi.org/10.1162/neco.1993.5.4.613>
- Dugas, M. J., Freeston, M. H., & Ladouceur, R. (1997). Intolerance of uncertainty and problem orientation in worry. *Cognitive Therapy and Research, 21*(6), 593–606. <https://doi.org/10.1023/A:1021890322153>
- Dymond, S., Dunsmoor, J. E., Vervliet, B., Roche, B., & Hermans, D. (2015). Fear generalization in humans: Systematic review and implications for anxiety disorder research. *Behavioral Therapy, 46*(5), 561–582. <https://doi.org/10.1016/j.beth.2014.10.001>
- Friston, K. J., Stephan, K. E., Montague, R., & Dolan, R. J. (2014). Computational psychiatry: The brain as a phantastic organ. *The Lancet Psychiatry, 1*(2), 148–158. [https://doi.org/10.1016/S2215-0366\(14\)70275-5](https://doi.org/10.1016/S2215-0366(14)70275-5)
- Fung, B. J., Qi, S., Hassabis, D., Daw, N., & Mobbs, D. (2019). Slow escape decisions are swayed by trait anxiety. *Nature Human Behaviour, 1*, 702–708. <https://doi.org/10.1038/s41562-019-0595-5>
- Gagne, C., Dayan, P., & Bishop, S. J. (2018). When planning to survive goes wrong: Predicting the future and replaying the past in anxiety and PTSD. *Current Opinion in Behavioral Sciences, 24*, 89–95. <https://doi.org/10.1016/j.cobeha.2018.03.013>
- Gallagher, M. W., Bentley, K. H., & Barlow, D. H. (2014). Perceived control and vulnerability to anxiety disorders: A meta-analytic review. *Cognitive Therapy and Research, 38*(6), 571–584. <https://doi.org/10.1007/s10608-014-9624-x>
- Gallagher, M. W., Naragon-Gainey, K., & Brown, T. A. (2014). Perceived control is a transdiagnostic predictor of cognitive-behavior therapy outcome for anxiety disorders. *Cognitive Therapy and Research, 38*(1), 10–22. <https://doi.org/10.1007/s10608-013-9587-3>
- Garcia, J., & Fernandez, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning, 16*, 1437–1480.
- Gaskett, C. (2003). Reinforcement learning under circumstances beyond its control. In *Proceedings of the International Conference on Computational Intelligence, Robotics and Autonomous Systems*. Washington, DC: IEEE.

- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66(6), 4585–595. <https://doi.org/10.1016/j.neuron.2010.04.016>
- Harlé, K. M., Guo, D., Zhang, S., Paulus, M. P., & Angela, J. Y. (2017). Anhedonia and anxiety underlying depressive symptomatology have distinct effects on reward-based decision-making. *PLoS One*, 12(10), e0186473. <https://doi.org/10.1371/journal.pone.0186473>
- Hunter, L. E., Meer, E. A., Gillan, C. M., Hsu, M., & Daw, N. D. (2019). Excessive deliberation in social anxiety. *bioRxiv*, 522433. <https://doi.org/10.1101/522433>
- Huys, Q. J., Daw, N. D., & Dayan, P. (2015). Depression: A decision-theoretic analysis. *Annual Review of Neuroscience*, 38, 1–23. <https://doi.org/10.1146/annurev-neuro-071714-033928>
- Huys, Q. J., & Dayan, P. (2009). A Bayesian formulation of behavioral control. *Cognition*, 113(3), 314–328. <https://doi.org/10.1016/j.cognition.2009.01.008>
- Huys, Q. J., Eshel, N., O'Nions, 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(3), e1002410. <https://doi.org/10.1371/journal.pcbi.1002410>
- Jacobson, N. C., & Newman, M. G. (2014). Avoidance mediates the relationship between anxiety and depression over a decade later. *Journal of Anxiety Disorders*, 28(5), 437–445. <https://doi.org/10.1016/j.janxdis.2014.03.007>
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
- Kessler, R. C., Sampson, N. A., Berglund, P., Gruber, M. J., Al-Hamzawi, A., Andrade, L., . . . Wilcox, M. A. (2015). Anxious and non-anxious major depressive disorder in the World Health Organization World Mental Health Surveys. *Epidemiology and Psychiatric Sciences*, 24(3), 210–226. <https://doi.org/10.1017/S2045796015000189>
- Konstantellou, A., Campbell, M., Eisler, I., Simic, M., & Treasure, J. (2011). Testing a cognitive model of generalized anxiety disorder in the eating disorders. *Journal of Anxiety Disorders*, 25(7), 864–869. <https://doi.org/10.1016/j.janxdis.2011.04.005>
- Lally, N., Huys, Q. J., Eshel, N., Faulkner, P., Dayan, P., & Roiser, J. P. (2017). The neural basis of aversive Pavlovian guidance during planning. *Journal of Neuroscience*, 37(42), 10215–10229. <https://doi.org/10.1523/JNEUROSCI.0085-17.2017>
- Lee, S. W., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, 81(3), 687–699. <https://doi.org/10.1016/j.neuron.2013.11.028>
- Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., . . . Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: The balloon analogue risk task (BART). *Journal of Experimental Psychology: Applied*, 8(2), 75–84. <https://doi.org/10.1037/1076-898x.8.2.75>
- Leotti, L. A., & Delgado, M. R. (2011). The inherent reward of choice. *Psychological Science*, 22(10), 1310–1318. <https://doi.org/10.1177/0956797611417005>
- Leotti, L. A., & Delgado, M. R. (2014). The value of exercising control over monetary gains and losses. *Psychological Science*, 25(2), 596–604. <https://doi.org/10.1177/0956797613514589>
- Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: The origins and value of the need for control. *Trends in Cognitive Sciences*, 14(10), 457–463. <https://doi.org/10.1016/j.tics.2010.08.001>
- Ly, V., Wang, K. S., Bhanji, J., & Delgado, M. R. (2019). A reward-based framework of perceived control. *Frontiers in Neuroscience*, 13, Article 65. <https://doi.org/10.3389/fnins.2019.00065>
- MacLeod, A. K., & Byrne, A. (1996). Anxiety, depression, and the anticipation of future positive and negative experiences. *Journal of Abnormal Psychology*, 105(2), 286–289. <https://doi.org/10.1037/0021-843X.105.2.286>
- Maner, J. K., Richey, J. A., Cromer, K., Mallott, M., Lejuez, C. W., Joiner, T. E., & Schmidt, N. B. (2007). Dispositional anxiety and risk-avoidant decision-making. *Personality and Individual Differences*, 42(4), 665–675. <https://doi.org/10.1016/j.paid.2006.08.016>
- Marr, D. C. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. Cambridge, MA: MIT Press.
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609–1617. <https://doi.org/10.1038/s41593-018-0232-z>
- Momennejad, I., Otto, A. R., Daw, N. D., & Norman, K. A. (2018). Offline replay supports planning in human reinforcement learning. *eLife*, 7, e32548. <https://doi.org/10.7554/eLife.32548>
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., & Gershman, S. J. (2017). The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9), 680–692. <https://doi.org/10.1038/s41562-017-0180-8>
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H., & Tanaka, T. (2012). *Parametric return density estimation for reinforcement learning*. arXiv preprint arXiv:1203.3497.
- Moutoussis, M., Shahar, N., Hauser, T. U., & Dolan, R. J. (2018). Computation in psychotherapy, or how computational psychiatry can aid learning-based psychological therapies. *Computational Psychiatry*, 2, 50–73. [https://doi.org/10.1162/CPSY\\_a\\_00014](https://doi.org/10.1162/CPSY_a_00014)
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: Opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507–520. <https://doi.org/10.1007/s00213-006-0502-4>
- Norbury, A., Robbins, T. W., & Seymour, B. (2018). Value generalization in human avoidance learning. *eLife*, 7, e34779. <https://doi.org/10.7554/eLife.34779.001>
- Oud, B., Krajbich, I., Miller, K., Cheong, J. H., Botvinick, M., & Fehr, E. (2016). Irrational time allocation in decision-making. *Proceedings of the Royal Society B: Biological Sciences*, 283(1822), 20151439. <https://doi.org/10.1098/rspb.2015.1439>
- Paulus, M. P., & Yu, A. J. (2012). Emotion and decision-making: Affect-driven belief systems in anxiety and depression. *Trends in Cognitive Sciences*, 16(9), 476–483. <https://doi.org/10.1016/j.tics.2012.07.009>
- Pittig, A., Treanor, M., LeBeau, R. T., & Craske, M. G. (2018). The role of associative fear and avoidance learning in anxiety disorders: Gaps and directions for future research. *Neuroscience & Biobehavioral Reviews*, 88, 117–140. <https://doi.org/10.1016/j.neubiorev.2018.03.015>



- Ramírez, E., Ortega, A. R., & Del Paso, G. A. R. (2015). Anxiety, attention, and decision making: The moderating role of heart rate variability. *International Journal of Psychophysiology*, *98*(3), 490–496. <https://doi.org/10.1016/j.ijpsycho.2015.10.007>
- Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., & Daw, N. D. (2017). Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLOS Computational Biology*, *13*(9), e1005768. <https://doi.org/10.1371/journal.pcbi.1005768>
- Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, *2*(4), 160–163. <https://doi.org/10.1145/122344.122377>
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Symmonds, M., Bossaerts, P., & Dolan, R. J. (2010). A behavioral and neural evaluation of prospective decision-making under risk. *Journal of Neuroscience*, *30*(43), 14380–14389. <https://doi.org/10.1523/JNEUROSCI.1459-10.2010>