Ines Aitsahalia, Joanna Kuo, and James Packman

Too Chicken to Cross the Road: Reinforcement Learning Models of Anxiety

Introduction

Anxiety disorders are characterized by persistent maladaptive beliefs and actions, despite contrary evidence. This disconnect between the environment and the emotional response can cause intense distress (APA, 2013). For example, imagine an anxious agent attempting to cross the street, but is afraid of getting hit by a car. This is a simplistic but instructive example. Naturally, pedestrians should have a healthy fear of cars to motivate them to look both ways before crossing the street. At a reportedly perilous intersection, even a healthy agent would feel some anxiety, and perhaps might even cross the street elsewhere. However, a pathologically anxious agent would, in contrast, not even dare to set foot in the street. They might only cross at a select few intersections that are known to be 100% safe. They might generalize this fear to a fear of crossing all streets, including ones that are not frequented by cars. Clearly, this avoidance would constitute a high degree of dysfunction and distress for the agent, distinguishing them from non-pathologically anxious pedestrians. Our current understanding of anxiety does not fully explain how such maladaptive fears and behaviors, such as avoidance, are maintained. Furthermore, current psychotherapeutic approaches leave much to be desired in that, even after treatment, avoidant behavior often resurfaces (Moutoussis et al., 2017). To better help people with anxiety disorders, we must incorporate knowledge from computational models of anxiety.

To that end, this paper asks how reinforcement learning models of anxiety can explain and inform the behavior and brain activity characteristic of anxiety. We begin by presenting two complementary models: a Q-learning model built on sequential evaluation, as well as the Latent Cause model. These models present two different but connected accounts of the maladaptive beliefs underlying anxiety. To supplement this knowledge, we then explore a model of the modulatory role of dopamine in anxiety.

Q-Learning

What kinds of computational models can explain anxious behavior? Decision-theoretic models describe how problems in **sequential evaluation** belie anxiety symptoms (Zorowitz et al., 2019). Zorowitz et al. (2019) present a fundamental tension underlying anxiety: that fears about the future drastically and disproportionately impact current risk assessment. They argue that pessimistic beliefs about the future can distort the entire decision-making process, leading to the avoidance behaviors seen in anxiety. They model anxious learning through a Markov Decision Process, with the standard assumption that an agent seeks to maximize total reward. Their model, shown below, includes *w*, a weight parameter for pessimism. This pessimism, when it becomes unrealistically inflated, can cause distorted learning (Zorowitz et al., 2019). Here, *w* can modulate both policy decisions and transition probabilities, meaning that this could explain a slew of anxious and avoidance behaviors that will be discussed in the behavior section.

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

The work of Zorowitz et al. draws on the decision-theoretic research of Huys et al. (2012). In their paper on depression, Huys et al. describe the ways that pessimism can be self-perpetuating, and that a completely healthy learning model can lead to drastic results based on a distorted initial condition (2012). This is also supported by Bach, who found that behavioral inhibition is largely linearly dependent on the magnitude and probability of a threat, which is mediated by an individual's anxiety (2015).

What *are* these negative beliefs about the future? Specifically, anxious agents assume that the future will not be optimal because the world is noisy and volatile and/or because the agent doubts their own efficacy. These negative beliefs about the world and the self, in addition to being hallmarks of depression, result in the notion that later avoidance of negative states will be futile (Zorowitz et al., 2019). Logically, then, the agent must engage in avoidance *now*. Such actions demonstrate on-policy learning vis-à-vis SARSA. Q-learning assumes that the agent will execute the optimal policy for an optimal future. But, if an agent espouses low self-efficacy beliefs and/or negative beliefs about the state of the world, then future avoidance of negative outcomes is futile, and the optimal future action is infeasible (Zorowitz et al., 2019). Consequently, the agent will use SARSA to account for the uncontrollable futute, resulting in exaggerated behaviors to avoid negative outcomes. Consequently, anxious agents evince fear generalization, exaggerated fear appraisal, and maladaptive avoidance and safety behaviors.

I. Fear Generalization

Anxious agents may be particularly likely to encounter distress and dysfunction when fear is generalized beyond an objectively negative state. Let us return to our anxious pedestrian agent. To any reasonable agent, being hit by a car may be considered a negative state. However, an anxious agent believes that the future will not be optimal (Zorowitz et al., 2019). Consequently, they may fear that 1) drivers on the road might be driving too fast or too recklessly and/or 2) the agent will not be able to avoid being hit, either due to their own clumsiness or to some other factor. For instance, the agent might fear that they might trip and fall or fail to see a car coming quickly round a bend. These fears reflect negative beliefs about the uncontrollable world or the inefficacious self (Zorowitz et al., 2019). Either one of these states coming to pass could result in a car hitting the agent (an objectively negative outcome). As a result, states antecedent to the negative state (i.e., prior to crossing an intersection, or even walking on the sidewalk) become intolerable. With an increasingly anxious agent, states further and further antecedent to the objectively negative state become negative (Zorowitz et al., 2019). Negative value "bleeding over" into antecedent states means that the agent cannot stomach standing near the street at all. Fear generalization thus leads to avoidance.

II. Exaggerated Fear Appraisal

Fear generalization also belies exaggerated fear appraisal. Because anxious agents do not assume that they will be able to execute an optimal future policy (a la Q-learning), they assume that they can *only* avoid future calamity by worrying sufficiently in the present. This behavioral pattern is more suited to truly terrifying future situations. For example, we should all be worried about climate change because, in the years to come, it will be far too late to reverse its negative effects. This example illustrates the point that worrying about negative future states far in advance may be exaggerated in everyday, non-existentially threatening cases. The Balloon Analog Risk Task (BART; Lejuez et al., 2002) elegantly illustrates how negative beliefs about the future and/or the self encourage an exaggerated appraisal of the negative outcome, even in a trivial setting. In the BART game, participants are given a virtual balloon which they may inflate with a pump. For each pump, the participant earns points or money; however, if the balloon pops, all the reward that would have been obtained is lost. In some versions, participants encounter balloons at high or low risk of popping. Anxious agents stop inflating their balloons much sooner than non-anxious (i.e., “optimistic”) agents, regardless of the risk level of the balloon. While all agents pump less in the high-risk condition, the effect is more pronounced for anxious agents (Lejuez et al., 2002). Although it is important for an agent to perceive high versus low risk, the anxious agents exaggerate their appraisal of the high-risk condition, resulting in them earning far fewer points. This result implies that anxious agents who “cash out” sooner thus miss out on potential positive outcomes.

III. Avoidance

The anxious agent’s avoidance may also cause them to miss out on potential positive outcomes. Suppose the agent had to cross the street to purchase some Nutella. To an anxious agent, if reaching the Nutella means crossing the artificially expanded negative space (i.e., the street and the fear-inducing areas surrounding it), the agent may simply forego Nutella. The trajectory towards a positive outcome becomes more distorted (Zorowitz et al., 2019) as the agent searches high and low to find something that approximates Nutella without crossing the street. Exploring the area more broadly, perhaps to find a special tub of Nutella, is out of the question. More generally speaking, this generalization of fear results in avoidance (Zorowitz et al., 2019), thereby precluding exploration and the enjoyment of positive outcomes. This phenomenon is sometimes called “aversive pruning,” where an anxious agent refuses to incur a large loss in the short-term out of fear of not being able to recoup that loss in the longer-term (Huys et al., 2012; Lally et al., 2017). In our example, the anxious agent will not cross the street for fear that they might be hit by a car.

How well does this sequential evaluation model explain different types of anxiety disorders? The anxious pedestrian agent example seems to revolve around a phobia, a persistent, maladaptive fear of a particular object/situation (in this case, cars in the street; APA, 2013). However, it is worth noting that one can also understand Obsessive-Compulsive Disorder (OCD) in terms of sequential evaluation. OCD is characterized by persistent, intrusive thoughts (“obsessions”) that cause distress (APA, 2013). The gent then tries to reduce that distress via typically repeated, ritualistic behaviors (“compulsions”) which themselves tend to compound longer-term distress (APA, 2013). For example, an Obsessive-Compulsive (OC) agent might have obsessions about someone breaking into her house. Attempting to reduce distress, the agent checks and re-checks all the locks in her house. In a particularly dangerous area, doing so may indeed be adaptive. What distinguishes the OC agent, however, is *repeatedly* checking the locks over the course of an hour or more, causing immense dysfunction and distress. To budget more time for this ritual, the agent may begin checking the locks hours before she needs to leave. By the logic of sequential evaluation, the agent's low self-efficacy beliefs necessitate action in the present because future action to avoid catastrophe will be futile (Zorowitz et al., 2019). As such, the OC agent must engage in compulsive behaviors (i.e., checking and re-checking the locks) in the present state, far antecedent to the negative outcome over which they obsess.

Of course, this Q-learning-based, sequential evaluation model is not the only computational attempt to explain anxious behavior. We now turn to the Latent Cause model for a complementary perspective on anxious behaviors.

Latent Cause Theory

I. Bayesian Inference and Distorted Priors

Another model of anxious decision making focuses on differences in the assignment of events to representative states. This model, called the Latent Cause model, introduced by Niv and Gershman (2012), explains anxious behavior in classical conditioning as Bayesian inference with an infinite capacity prior. They posit that animals, including humans, group observations of the world into clusters and explain them with, sometimes hidden, causes. These latent causes are unobservable, but explain stimuli. Different animals may assign the same stimulus to different latent causes, depending on context *how* animals make the decision of which cause to assign an observation to could explain why people with anxiety consistently assume threat in otherwise safe environments. Furthermore, a patient who confuses latent causes may benefit less from therapy: If a patient attributes anxious feelings to a certain stimulus (i.e., one latent cause), but then attributes the positive feelings of therapy only to the therapist or to the therapy session (i.e., some new latent cause, the patient fails to update their learning about the anxiety-inducing stimulus. This process has been called "overaccommodation" (Moutoussis et al., 2017).

Pisupati proposes that agents can be modeled as having two extremes of possible prior beliefs: the world is overly stochastic, or the world is overly deterministic.[1] With the overly stochastic distortion, an agent is more likely to believe that surprising information and events are common, and is less likely to assign new evidence to a new latent cause, rather updating the old latent cause to include more variability. However, a belief in the determinism of the world would lead to an over-assignment of information to new latent causes, which could explain the reemergence of anxiety symptoms despite treatment, as explained by Moutoussis et al. (2017). If the world is random (or stochastic), it's possible that new or contradictory stimuli are simply the result of chance. If the world is fixed (or deterministic), new information must be brought about by a new cause since everything corresponds to a specific, recurrent, identifiable cause.

Applying this work to instrumental conditioning, it becomes clear how an overly deterministic distortion in prior beliefs can lead to the safety behaviors observed in patients with anxiety and obsessive disorders. An OC agent's compulsions may have little to no effect on the undesirable outcome; rather, they reduce the belief that the undesirable outcome will occur and thereby reduce fear. If the world is controllable, then an agent simply has to perform certain

[1] Pisupati's current research, unpublished as of 8 December, 2020.

actions to achieve the desired outcome, and the seemingly irrational safety behaviors are completely logical.

II. State Transitions

However, there is another component to anxiety and state transition beyond simply distorted prior beliefs. Take as an example an anxious agent seeing a glass near the edge of a trable. This agent may fear it will fall (negative beliefs about the future), will prefer the object be farther from the edge (generalization of negative value to antecedent states), and may then engage in a safety behavior; namely, moving the object farther from the edge. But, is the agent learning that antecedent states predict negative outcomes? Does an object being near the edge *necessarily* mean it will fall? Clearly not; however, Obsessive Compulsive (OC) agents doubt their prior experience when predicting future states (Fradkin et al., 2020). It is possible that OC agents have a confirmation bias dealing towards experiences that support their obsessions, and doubt experiences that contradict them. In Fradkin et al.'s Bayesian learning experiments (2020), OC participants, due to their uncertainty about state transitions, fail to properly weight previously learned contingencies. Consequently, OC participants are not only surprised by expectable outcomes, but engage in excessive "exploratory behavior" (Fradkin et al., 2020). "Exploration" may seem to be a counterintuitive term because anxiety disorders tend to involve a lack of exploration out of fear. However, with OCD, since agents fail to learn from past experience, they must look elsewhere for information about state transitions. Specifically, the agent will expend significant effort obtaining more information from the current state while discounting priors. For example, an OC agent will repeatedly check the stove to make sure it is off, discounting the fact that they have never forgotten to do so. This compulsion may be described as "exploratory" in this way: Uncertainty about state transitions may result in anxious agents feeling compelled to engage in safety behaviors that involve obtaining more information pertaining to their obsessions (e.g., checking the stove). Obsessions in OCD stem from a kind of prospective evaluation problem: the agent will continually replay dangerous future scenarios (Lohr et al., 2007), worrying about them while discounting their prior unlikelihood (Fradkin et al., 2020). This problem is also an example of exaggerated threat appraisal. The agent then believes that compulsive behaviors can mitigate their worries, resulting in the repetitive, excessive performance of safety rituals (Lohr et al., 2007). These behaviors become reinforced because the obsessed-over scenario never came to pass (Lohr et al., 2007). For example, our anxious pedestrian agent may replay the scenario of getting hit by a car repeatedly, and may look both ways over and over again before crossing the street, discounting both the information gained from looking both ways before as well as the information gained from safely crossing the street on previous occasions.

In this section, we've shown how skewed beliefs about the world as well as uncertainty about state transitions can inform anxious and compulsive behavior. The Latent Cause model does not only apply to anxiety. An overly stochastic worldview can explain the depressive symptoms, such as lack of motivation, that are also frequently seen in patients with anxiety. Believing that one has no control over the environment can lead to despair and frustration. Should an agent begin with a deterministic view, they are likely to engage in safety or avoidance behaviors in order to feel control. However, if they continually experience stimuli that are contradictory to their beliefs, it is possible that their view of the world (their prior in the Bayesian equation) will swing to the opposite extreme, and be one of extreme stochasticity. This change could explain the high comorbidity of anxiety and depression. In addition, a belief in the stochasticity of the world can also inform anxious exploration. An anxious agent who believes the world is noisy and volatile (Zorowitz et al., 2019) may thus obsess over obtaining more information.

Following Marr's three levels of analysis, we have already discussed the computational and the algorithmic components. Thus, we turn to the implementational level as we explore the neurophysiological mechanisms that underlie anxiety and how computational models of these mechanisms can better provide medications for patients.

Opponent Actor Learning

I. The Neurobiology of Anxiety Disorders

Much research has shown that the limbic system, which includes the amygdala, hippocampus, basal ganglia, and medial prefrontal cortex is the major system that controls our emotions (Martin et al. 2009). During times of stress or anxiety, the limbic system is activated in order to gather information about the environment and to respond to it appropriately. In anxiety disorders, however, this system becomes overly activated, causing agents to misinterpret stimuli and therefore respond to them inappropriately.

II. OpAl and the Modulatory Role of Dopamine in Anxiety Disorders

One neurotransmitter that has been heavily suggested to play a modulatory role in anxiety disorders is dopamine. Dopamine is a neurotransmitter that regulates motor control, emotions, and decision making. The different types of dopamine receptors can be characterized into two classes — D1 and D2. Each is present in different parts of the brain respectively and have opposing functions, but work together to modulate behavior. The two main pathways that dopamine neurons innervate are the Go and No-Go pathways, located in the basal ganglia. They are named as such due to the effects each pathway elicits. D1 dopamine neurons trigger the Go pathways, exciting adenylatecyclase activity and increasing cyclic adenosine monophosphate (cAMP), which leads to the positive reinforcement of an action (Go), whereas D2 dopamine neurons inhibit cAMP activity, inducing long term depression (No-Go) (Collins et. al, 2014).

To get a better understanding of how D1 and D2 neurons work together, we will look at the algorithmic model called OpAl (Opponent Actor Learning) and analyze how parameters in the model could explain the behaviors seen in anxiety disorders. Developed by Anne G. E. Collins and Michael J. Frank, OpAl is built on the classic actor-critic model, but instead of only having one actor, there is a dual opponent actor system that discriminates between negative and positive values. The purpose of this is to "model the influence of dopamine on incentive choice — the tendency to differentially weigh costs and benefits — after learning has occurred" (Collins et al., 2014). This is relevant to anxiety disorders because it can explain why agents display avoidance behaviors even after treatment.

The formulae for the two actor weights are as follows:

$$G_a(t+1) = G_a(t) + [\alpha_G G_a(t)] \times \delta(t)$$
$$N_a(t+1) = N_a(t) + [\alpha_N N_a(t)] \times -\delta(t)$$

where $\delta(t)$ is the previously defined critic prediction error and α_G and α_N are the learning rates for the Go and No-Go weights. As seen from the equation, G weights are positively correlated with the true expected value whereas N weights are negatively correlated with the true expected value. The overall weight for a specific action, $Act_{a(t)}$, then is

$$Act_{a(t)} = \beta_G G_a(t) - \beta_N N_a(t)$$

where it is dependent on β_G and β_N which are parameters that modulate the extent to which the G and N weights are represented in a given trial. This then gives an overall probability for a specific action occurring as the weighted difference between G_a and N_a compared to all the other actions.

$$p(a) = \frac{e^{Act_a(t)}}{\sum_i e^{Act_{a_i}(t)}}$$

The reason why there are two different pairs of parameters in the model is to capture the asymmetries in both learning and choice incentive effects. The learning rates represent the ability to distinguish between good and bad options, whereas β_G and β_N weighs these options during time of choice.

Simulating experiments where mice press a lever to receive an optogenetic laser stimulation, by modulating the G and N activity levels, and comparing it to experimental results, the researchers found that OpAl was able to accurately account for the learning and incentive effects of striatal dopamine. With regards to learning they were able to depict approach learning and avoidance learning and found that D1 MSN stimulation developed stronger G weights for the action and therefore was more likely to be chosen, whereas D2 MSN stimulation led to stronger N weights and a decrease in likelihood for an action to be chosen. Maintaining α_G and α_N levels, they were also able to show how β_G and β_N which represent choice incentive effects, affect the likelihood of an action occurring. By increasing β_G , where $\beta_G > \beta_n$, they found that it corresponded to an increase in motivation.

From this data, we can conclude that dopamine plays a modulatory role not only during learning but also at the time of choice. The implications of this model with regards to anxiety disorders, suggests that there could be differences in learning and motivation in agents with anxiety, where it is a combination of avoidance learning and the inability to weigh options correctly, whether it is due to exaggerated fears or distorted priors, during times of choice that ultimately results in the agent making sub-optimal decisions. This can also explain why anxiety behaviors resurface in agents even after therapy.

III. Pharmacological Treatments for Anxiety Disorders

As of today, the courses of treatment for anxiety disorders are mostly SSRIs, SNRIs and benzodiazepines (McGowen, 2019). The former two inhibit the reuptake of serotonin and norepinephrine in order to increase the mood of the individual, whereas the latter, increases the effects of GABA receptors. However, due to a wide range of anxiety disorders and their severity, only 60% of individuals see a decrease in their anxiety while taking medication (McGowen, 2019). And many of them who do oftentimes relapse. In addition, there are many side effects from these medications including weight gain, sexual dysfunction, and sleep disturbance (Ferguson, 2001). From our study of the dopaminergic system, we know that there must be a fine balance between the Go and No-Go pathways in order to make the optimal decision. In addition, the weights of these pathways not only have effects in decision making but also learning. Therefore, to treat anxiety disorders could not just be developing medications that target mood, but also at mechanisms that enhance learning and incentive choice.

Conclusion

This paper examined how computational models can explain the distinctive behavior and brain activity characteristic of anxiety disorders. We considered three different mathematical models in this review. Exploring sequential evaluation and Latent Cause Theory revealed two complementary computational models for the cognition underlying anxiety disorders. These models illuminate behavioral patterns across anxiety-related disorders, including OCD. Finally, the neurobiology of anxiety disorders and the modulatory role of dopamine in producing anxiety-like behaviors allowed us to consider current pharmacological treatments for anxiety disorders. One hopes that computational psychiatry can inform future treatments for anxiety disorders. Common models, behaviors, and neurobiological mechanisms implicated in a range of anxiety disorders might lend credence to a transdiagnostic approach to treatment. For example, it is hypothesized that the key difference between Post-Traumatic Stress Disorder (PTSD), Panic Disorder (PD), specific phobias, and Obsessive-Compulsive Disorder (OCD) are 1) the degree to which the agent feels that they can predict the onset of stress, and 2) the degree to which the agent feels they can control the offset of stress (Lohr et al., 2006). For instance, OCD's obsessions are unpredictable in onset, but the agent uses safety behaviors in an attempt to reduce distress because they believe doing so allows them to control the offset of stress (Lohr et al., 2006). Training clinicians to understand these underlying patterns explained by these models and mechanisms will, one hopes, improve treatment in the future.

Acknowledgements

We would like to thank Dr. Sashank Pisupati and Professor Niv for their guidance and help in refining our topic, surveying the literature, and understanding the complexities of computationally modeling the brain. In addition, we would like to thank Jean and Max for their feedback on our first draft as well as to the entire class of NEU 338 for providing comments and suggestions during our presentations and participating in the BART task.

References

American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders*. 5th ed. Washington, DC: American Psychiatric Association; 2013.

Bach, DR. (2015). Anxiety-like behavioural inhibition is normative under environmental threat-reward correlations. *PLoS Comput Biol 11*(12): e1004646. https://doi.org/10.1371/journal.pcbi.1004646.

Collins, Anne G. E., and Michael J. Frank. "Opponent Actor Learning (OpAL): Modeling Interactive Effects of Striatal Dopamine on Reinforcement Learning and Choice Incentive." *Psychological Review*, vol. 121, no. 3, 2014, pp. 337–366., doi:10.1037/a0037015.

Ferguson, James M.. "SSRI Antidepressant Medications: Adverse Effects and Tolerability." Primary care companion to the Journal of clinical psychiatry vol. 3,1 (2001): 22-27. doi:10.4088/pcc.v03n0105.

Fradkin I, Ludwig C, Eldar E, Huppert JD (2020) Doubting what you already know: Uncertainty regarding state transitions is associated with obsessive compulsive symptoms. PLoS Comput Biol 16(2): e1007634. https://doi.org/10.1371/journal.pcbi.1007634.

Lohr, Jeffrey M., Olatunji, Bunmi O., Sawchuk, Craig N. (2007). A functional analysis of danger and safety signals in anxiety disorders.*Clinical Psychology Review, 27*(1) https://doi.org/10.1016/j.cpr.2006.07.005.

Gershman, S.J., Niv, Y. (2012). Exploring a latent cause theory of classical conditioning. *Learn Behav 40*, 255–268. https://doi-org.ezproxy.princeton.edu/10.3758/s13420-012-0080-8.

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.

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.

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.

Martin, Elizabeth I., et al. "The Neurobiology of Anxiety Disorders: Brain Imaging, Genetics, and Psychoneuroendocrinology." Psychiatric Clinics of North America, vol. 32, no. 3, 2009, pp.

549–575., doi:10.1016/j.psc.2009.05.004.

Mcgowan, O.o. “Pharmacogenetics of Anxiety Disorders.” Neuroscience Letters, vol. 726, 2020, p. 134443., doi:10.1016/j.neulet.2019.134443.

Mora, Miguel Pérez De La, et al. (2010). “Role of Dopamine Receptor Mechanisms in the Amygdaloid Modulation of Fear and Anxiety: Structural and Functional Analysis.” Progress in Neurobiology, *90*(2), 198–216., doi:10.1016/j.pneurobio.2009.10.010.

Moutoussis, M., Shahar, N., Hauser, T. U., & Dolan, R. J. (2017). 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.

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.