1

2 Comparing Anticipation and Uncertainty-Penalty Accounts of Non-Instrumental
3 Information Seeking

4 Shi Xian Liew, Jake R. Embrey, Danielle J. Navarro, and Ben R. Newell

5 School of Psychology, University of New South Wales, NSW, Australia

6

Author Note

7 Correspondence concerning this article should be addressed to: Shi Xian Liew, UNSW,
8 NSW Australia E-mail: shixianliew@gmail.com

Abstract

Proposed psychological mechanisms generating non-instrumental information seeking in humans can be broadly categorised into two competing accounts: the maximisation of anticipating rewards versus an aversion to uncertainty. We compare three separate formalisations of these theories on their ability to track the dependency of information seeking behaviour on increasing levels of cue-outcome delay as well as their sensitivity to outcome valence. Across three experiments using a variety of different stimuli, we observe a flat to monotonically increasing pattern of delay dependency and minimal evidence of sensitivity to outcome valence—patterns which are better predicted, qualitatively and quantitatively, by an uncertainty aversion information model.

Keywords: information seeking; temporal discounting; reinforcement learning; anticipation; computational modeling;

1 Introduction

The human tendency to seek out information is well documented (Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014; Kidd & Hayden, 2015; Navarro, Newell, & Schulze, 2016). This is perhaps unsurprising under situations where information can be used to guide future decisions. However, this tendency to want-to-know also manifests when information is apparently non-instrumental, that is, when it is irrelevant for future actions (Charpentier, Bromberg-Martin, & Sharot, 2018; Iigaya, Story, Kurth-Nelson, Dolan, & Dayan, 2016; Sharot & Sunstein, 2020; Vasconcelos, Monteiro, & Kacelnik, 2015; Zhu, Xiang, & Ludvig, 2017). Such behaviour is observed even when the information incurs a cost (Bennett, Bode, Brydevall, Warren, & Murawski, 2016; Bennett, Sutcliffe, Tan, Smillie, & Bode, 2020; Cabrero, Zhu, & Ludvig, 2019; Eliaz & Schotter, 2010; Pierson & Goodman, 2014). To illustrate the basic idea, imagine awaiting the draw of a raffle ticket from a lottery: the ticket is drawn, you can see it's not the colour of the ticket you bought (so you have not won) but the temptation to know the winning number remains strong.

Several theories explaining non-instrumental information seeking behaviour adopt a reinforcement learning approach (e.g., Beierholm & Dayan, 2010; Bromberg-Martin & Hikosaka, 2011). This class of models offers two important predictions: first, they typically involve a temporal discounting mechanism that assumes preference for information decreases as the delay between cue and outcome increases (Loewenstein, 1987). This is generally founded on the idea that rewards themselves lose value the longer we have to wait to obtain them, consequently resulting in a decrease in the value of information on those rewards. Second, they suggest that people will avoid non-instrumental information about potentially aversive future outcomes (Iigaya et al., 2016; Kobayashi, Ravaoli, Baranès, Woodford, & Gottlieb, 2019; Zhu et al., 2017). Within their general framework of information seeking, Sharot and Sunstein (2020) suggest that the knowledge of aversive outcomes can induce negative affect resulting in the avoidance of such information. Discussion involving general (usually instrumental) information avoidant behaviour is not

new (e.g., see Golman, Hagmann, & Loewenstein, 2017; Hertwig & Engel, 2020), however the reinforcement learning models we discuss here are among the few to explicitly consider the avoidance of *non-instrumental* information.

A different class of models focuses on temporal resolution of uncertainty (e.g., Bennett et al., 2016; Epstein & Zin, 1989; Kreps & Porteus, 1979). Specifically, they describe information seeking behaviour as the preference for resolving uncertainty at an earlier (rather than later) point in time. These models are agnostic to the valence of the outcome—they predict a preference for uncertainty reduction irrespective of whether a future outcome is potentially positive or negative.

In the current work, we examine these competing accounts of non-instrumental information seeking using an experimental task in which we manipulate the delays between cues and outcomes as well as the valence of outcomes. We then examine the capacities of three models of information seeking (two from the reinforcement learning tradition, and one based on uncertainty resolution or aversion) in accounting for our behavioural data. While all three models offer reasonable predictions of behavioural patterns, the uncertainty aversion account ultimately emerges as the preferred model.

We use an experimental paradigm known as the *secrets* task. In this task, observers are presented with a risky event and are given the choice to receive advance information in the form of completely predictive cues (Find Out Now: FON) or to receive an ambiguous cue (Keep It Secret: KIS; Figure 1). After the cue is revealed, the observer is made to wait for a period of time before the outcome and the corresponding reward is presented. Recent studies adopting this task have found evidence for an increase in information seeking as the delay between cues and outcomes increases (Iigaya et al., 2016), as well as a reduction in information seeking for negative compared to positively valenced outcomes (Zhu et al., 2017). However, no studies to date using the secrets task have found clear evidence of temporal discounting, nor information-avoidant behaviour. Even at cue-outcome delays of 40 seconds there is no diminished preference for information compared with shorter delays,

and when outcomes are highly aversive images of mutilations, preference for advanced knowledge remains at approximately 50%. Such findings are potentially problematic for accounts which assume temporal discounting and avoidance.

In a related study, using a modified secrets task in which probabilities of outcomes were varied but delays were held constant, Charpentier et al. (2018) found that when the probability of winning increased so too did information seeking, with the opposite pattern observed for losses. This study thus provides further evidence indicating an effect of outcome valence on the strength of information seeking but again there was no evidence for information avoidance. Even when faced with a high probability of a loss, average choice proportions still showed a majority for information seeking.

One impediment to comparing many of the existing studies is that specifics of the tasks used vary, thus making it more difficult to draw general conclusions regarding the impact of different variables. For example, Zhu et al. (2017) varied valence in the secrets task but did not manipulate delay; Iigaya et al. (2016) and Iigaya et al. (2020) did the opposite, and Charpentier et al. (2018) used a subtly (but importantly)¹ different task and used monetary outcomes in comparison to the primary reinforcers used by Iigaya et al. (2016), Iigaya et al. (2020) and Zhu et al. (2017). Primary reinforcers refer to rewards that are available for immediate consumption or indulgence and have been argued to be important in eliciting delay-dependent choice effects (Crockett et al., 2013; Iigaya et al., 2016) and contrasts with secondary reinforcers, which are rewards provided as a proxy for primary reinforcers (e.g., money, a secondary reinforcer, being used as a medium to buy food, a primary reinforcer). However, the impact of varying reinforcer type in the same task is yet to be examined, thus leaving open questions about the importance of the reinforcer for eliciting non-instrumental information-seeking.

Here, we take a comprehensive approach by manipulating delays—including delays that are twice as long as those in previous work—the valence of outcomes and the type of

¹In Charpentier et al. (2018) the presentation of cues themselves were probabilistic, and gamble outcomes of the noninformative cue were also hidden.

outcomes. Specifically, we use monetary outcomes (gains and losses) in Experiment 1, and then primary reinforcers of chocolate (gain) and aversive noise (loss) in Experiments 2 and 3. We chose these outcome types over images because although the images used in previous work can be considered primary reinforcers, it is difficult to measure the extent to which participants actively engaged with the stimuli at the time of image presentation. In contrast, our primary reinforcers directly stimulated sensory modalities outside the visual aesthetics provided by images.

As noted above, of the three models we apply to our data, two are based on the anticipation of future outcomes: the Reward Prediction Error-Anticipation (RPE-A) model (Iigaya et al., 2016) and the Anticipated Prediction Error (APE) model (Zhu et al., 2017), and one, based on uncertainty aversion is known as the Uncertainty Penalty (UP) model (Bennett et al., 2016). Although all three models allow for varying levels of cue-outcome delay dependency, they do so using different explanatory mechanisms. Our overall goal is to see which of these models offers the best account of information-seeking behaviour in our experiments.

2 Experiment 1

In their original study, Iigaya et al. (2016) did not find strong temporal discounting effects on information seeking behaviour at their longest delay duration (40 s), speculating that such effects may only manifest at even longer delay durations. We designed our first experiment with this objective, manipulating delay lengths up to twice as long as Iigaya et al. (2016) (i.e., 80 s) in a secrets task using secondary reinforcers where participants could separately win or lose points.

2.1 Method

2.1.1 Participants

Following the sample sizes of previous studies (e.g., Charpentier et al., 2018), we recruited 40 UNSW psychology undergraduate students ($M_{age} = 20.13$ years, 25 females, 15 males) in exchange for course credit, as well as a cash amount depending on the points awarded during the task ($M = 3.50$ AUD).

2.1.2 Materials

We implemented the experiment in jsPsych (De Leeuw, 2015) within the Google Chrome browser on a desktop computer.

2.1.3 Design and Procedure

Each participant completed two blocks of trials in a randomised order. All participants started the experiment with 3000 points to prevent the possibility of obtaining a negative balance. In the “win block”, the positive outcome of each trial was receiving a random sample from a uniform distribution between 90-110 points, and the negative outcome was receiving 0 points. In the “loss block”, the positive outcome was receiving 0 points and the negative outcome was losing a random sample from a uniform distribution between 90-110 points. Participants were informed that the probability of obtaining a positive outcome or a negative outcome were equal and both were independent from the participants’ choice of option.

On each trial, participants were presented with the cue-outcome delay duration as well as the option to either *Find Out Now* (FON) or *Keep It Secret* (KIS). If participants chose the FON option, they immediately received either a smiley face (cueing a positive outcome) or a sad face (cueing a negative outcome) with equal probability. If participants chose the KIS option, they were presented with a confused face (a noninformative cue)

Figure 1: General experimental design of the secrets task. t indicates the delay between presentation of the cue and the reward. Specific rewards varied across experiments. Experiment 1 (win block): participants could win 100 points or otherwise lose nothing. Experiment 1 (loss block) participants could win nothing or otherwise lose 100 points. Experiment 2 (chocolate): participants could win an M&M chocolate or lose nothing. Experiment 2 (sound): participants could win nothing or otherwise receive an aversive noise. Experiment 3: participants could win an M&M chocolate or otherwise receive an aversive noise.

instead. Cues remained on screen for as long as the delay duration required, followed by
 delivery of the outcome. There were seven delay lengths of 1, 2, 5, 10, 20, 40 and 80
 seconds randomly interleaved throughout the session for each participant, with 10 trials for
 each delay length from 1 to 20 seconds, and 5 trials for delay lengths of 40 and 80 seconds,
 resulting in 60 trials in total per participant. The generic structure of the experiment is
 presented in Figure 1. After completion, participants were debriefed on the experiment's
 aims and reimbursed according to the amount of money they earned: 1000 points = \$1
 AUD.

Figure 2: Mean FON proportions across increasing delay (log-scaled) for each experiment condition. Error bars indicate 1 standard error of the mean.

Table 1: GLMM statistics for Experiment 1

Linear Model	parameters	BIC	χ^2	df	<i>p</i>
Baseline	2	5646.9	-	-	-
Random Slopes	4	5577.1	86.73	2	< .001
Random Slopes with Delay	5	5583.2	2.38	1	.12
Random Slopes with Delay and Valence	6	5591.6	0.02	1	.89

2.2 Results and Discussion

The mean probability of choosing FON across increasing delay durations is presented in Figure 2a for the win blocks, and Figure 2b for the loss blocks. Visual inspection of the results suggests an overall preference for the FON option that is constant across delay and unaffected by valence. Averaged across all trials, this preference was significantly greater than chance for the win ($M = 0.63, t(39) = 3.67, 95\% \text{ CI}[0.56, 0.69], p < .001$) and loss blocks ($M = 0.63, t(39) = 3.72, 95\% \text{ CI}[0.56, 0.69], p < .001$).

We used generalised linear mixed models (GLMM) to investigate the effect of delay and valence on information seeking behaviour. We compared the baseline model using only random intercepts for individual participants to models with random slopes (as a function of delay) as well as delay length and valence as a fixed effects. We found no evidence for an effect of delay or valence (Table 1).²

Contrary to Zhu et al. (2017) we found no effect of outcome valence on this preference. In addition, unlike Iigaya et al. (2016) we found no delay-dependent effect on

²While it falls outside the immediate scope of this study, we also analysed the correlation between FON choices and trial number. Details of this analysis can be found in the supplementary materials.

information seeking behaviour. These patterns suggest that primary reinforcers are important for eliciting delay-dependent changes in the preference for non-instrumental information.

3 Experiment 2

In light of the absence of delay-dependent effects with the secondary reinforcers of Experiment 1, Experiment 2 used primary reinforcers; namely M&M chocolates as the positively-valenced reward, and a high pitched aversive tone as the negative stimulus (microphone feedback, ranked 2nd-most aversive stimulus out of 34 other stimuli; Cox, 2008). These two conditions were designed to be analogous to the win and loss block conditions from Experiment 1, but in this experiment were run on different groups of participants (i.e., valence of reward was between-subjects rather than within).

3.1 Methods

3.1.1 Participants

The sound condition in Experiment 2 had 51 undergraduate psychology student participants ($M_{age} = 19.49$, 38 females, 13 males) who were granted course credit for their participation. The chocolate condition in Experiment 2 had 49 participants ($M_{age} = 21.82$, 20 females, 29 males). Participants were a mix of undergraduate students who received course credit for participation and paid participants who received \$15 AUD. Participants were required to fast for two hours prior to the experiment and self-reported that they liked M&M chocolates. We also ensured participants had not previously completed Experiment 1 or other similar studies.

3.1.2 Design and Procedure

The chocolate condition was similar to the win block in Experiment 1. Participants who received a positive outcome were presented with a single M&M chocolate delivered via an automated dispenser, otherwise a negative outcome delivered nothing. In the sound condition (comparable to the loss block in Experiment 1), a negative outcome presented an aversive microphone feedback noise that played for 10 seconds via headphones. A positive outcome in the sound condition did not result in any specific event, however participants had to wait for 10 seconds before they were able to proceed to the next trial. As with Experiment 1, the probability of positive and negative outcomes were equal.

In order to increase the robustness of estimates, we raised the number of trials at the 40 and 80 second delays to 10 for each delay length resulting in 70 trials in total across the experiment. The design was otherwise identical to Experiment 1.

3.2 Results and Discussion

Mean information preferences across increasing levels of delay are presented in Figure 2c for the chocolate condition and Figure 2d the sound condition. On average, across all levels of delay, the chocolate condition did not yield significant information seeking behaviour ($M = 0.56, t(47) = 1.89, 95\% \text{ CI} [.50, .63], p = .06$). The sound condition indicated significant preference for advance information ($M = 0.63, t(50) = 4.07, 95\% \text{ CI} [.57, .69], p < .001$).

We analysed the average preference for the FON option using a GLMM analysis analogous to that of Experiment 1 (see Table 2). In the chocolate condition, we found that the addition of random slopes significantly improved the fit of the model, suggesting information preference across time varied significantly between participants. The addition of delay also significantly improved the model's fit, indicating that delay length had an effect on information preference. On average we observed that the preference for the FON option increased from short to long delays. Specifically, at 1 and 2 second delays, people

Table 2: GLMM statistics for Experiment 2

Condition	Linear Model	parameters	BIC	χ^2	df	p
Chocolate	Baseline	2	3959.5	-	-	-
	Random Slopes	4	3269.1	694.43	2	< .001
	Random Slopes with Delay	5	3261.4	9.67	1	.002
Sound	Baseline	2	4056.3	-	-	-
	Random Slopes	4	3609.2	463.53	2	< .001
	Random Slopes with Delay	5	3607.7	9.61	1	.002

were slightly information avoidant (preferences of .45 and .43 respectively), with preferences for FON increasing to .70 and .72 at the 40 and 80 second delays respectively.

Similar to the chocolate condition, the addition of random slopes and delay improved the model's fit (both $p < .01$) in the sound condition. The pattern of results was also similar to the chocolate condition: at 1 and 2 second delays, mean preference was .52 and .51 respectively, with these preferences increasing to .73 and .74 at 40 and 80 seconds respectively.

Unlike the secondary reinforcers used in Experiment 1, both tasks using primary reinforcers found that delay has an effect on information preference. These results are consistent with those of Iigaya et al. (2016). Despite doubling the length of the longest delay (40 seconds to 80 seconds) we did not, however, see the non-monotonic pattern - decreasing preference for information as delay increases - speculated by Iigaya et al. (2016). In addition, we did not observe any clear information avoidant behaviour in any condition; a finding which is at odds with the predicted pattern of avoidance, or at least lower preference for non-instrumental information about negative outcomes (Charpentier et al., 2018; Sharot & Sunstein, 2020; Zhu et al., 2017).

4 Experiment 3

Experiment 2 was successful in eliciting delay-dependent information behaviour, however the absence of information avoidance or a reduction in information seeking when facing negative primary reinforcers is puzzling given the predictions of some accounts (Iigaya et al., 2016; Sharot & Sunstein, 2020). In Experiment 3 we examined the possibility that information seeking depends not only on the absolute valence of the outcome, but also its valence relative to other potential outcomes. Specifically, it is possible that the negative outcomes in Experiments 1 and 2 were not *relatively negative* enough compared to the other neutral outcomes (i.e., winning 0 points, or not receiving the aversive sound). By replacing the neutral outcomes with a positively valenced reward, we increase the relative negativity of the negative outcomes. If information avoidant behaviour is contingent on both absolute and relative valences of outcomes, we predicted that this manipulation would lead to a pattern of information avoidance. To investigate this possibility, in Experiment 3 we removed the neutral outcome and used primary reinforcers for all possible outcomes. Specifically, M&M chocolates were rewarded as the positive outcome and the aversive microphone sound was delivered as the negative outcome. This mixed paradigm is similar to the mixed condition in Zhu et al. (2017) which used erotic and aversive images for the positive and negative rewards, respectively (but held delay constant at 20 seconds for all trials).

4.1 Method

4.1.1 Participants

Fifty-one undergraduate students from the UNSW psychology cohort ($M_{age} = 19.27$ years, 30 females, 21 males) were recruited via the SONA platform in exchange for course credit.

Table 3: GLMM statistics for Experiment 3

Linear Model	parameters	BIC	χ^2	df	p
Baseline	2	3859.8	-	-	-
Random Slopes	4	3575.6	300.44	2	< .001
Random Slopes with Delay	5	3583.5	0.29	1	.59

4.1.2 Design and Procedure

The design and procedure were similar to Experiment 2, with a single block consisting of 70 trials (10 per delay) where the positive outcome was an M&M chocolate and the negative outcome was the aversive microphone feedback noise.

4.2 Results and Discussion

Overall, a significant preference for advance information was observed ($M = 0.61, t(48) = 3.27, 95\% \text{ CI} [.55, .68], p = .002$), but somewhat surprisingly the average choice proportion of FON did not vary considerably across the different delay lengths, indicating no clear effect of delay-dependence (Figure 2e).

Applying the same GLMM analysis to this data indicated that preference for information varied significantly between participants, where the addition of random slopes significantly improved the fit of a baseline model (Table 3). However, the addition of delay did not significantly improve the model's fit.

Despite using primary reinforcers, and offering a choice between a positive and a negative outcome on each trial, we saw no evidence of information avoidance. If anything, overall preference for information appeared to be highest (on average) of all three experiments. Contrary to the suggestions of previous studies (e.g., Iigaya et al., 2016; Sharot & Sunstein, 2020) these findings suggest that the avoidance of information is not dependent on the relative valence of negative outcomes compared to positive outcomes.

In addition, the results of Experiment 3 yielded no evidence of delay dependent

information preference. Not only was there no attenuation in information preference at longer delays as predicted by earlier literature Iigaya et al. (2016), there was no observation of any increase in shorter delays. While a general effect of delay dependence was not the central prediction in this experiment, its absence here is a point worth raising for future investigation, which we further address in the General Discussion.

5 Models of Information Seeking

We consider three formal models that provide different accounts of non-instrumental information seeking. Two of the models give an explanation based on anticipatory utility, in which the decision maker is presumed to derive additional value from the anticipation of the arrival of the rewarding outcome. According to such an account, people seek (or avoid) non-instrumental information about a future outcome because it allows us to savor (or dread) the event to come. The third model relies on an entirely different principle, namely that uncertainty is inherently aversive, and people seek information in order to remove uncertainty about an outcome, quite irrespective of the utility of the outcome itself. In this section we describe all three models. Notation schemes vary across papers, and to minimise confusion we apply a consistent notation scheme to describe events in the experiment, as outlined in Figure 3. If an agent stores some internal representation of the value of state s it is denoted $V(s)$. The reward experienced when the trial progresses from some state s to a later state s' is denoted $R(s, s')$. In cases where this reward does not depend on the initial state s , this is denoted $R(\cdot, s')$. Unless stated otherwise, the rewards and values are subject to temporal discounting and the notation refers to the temporally discounted reward/value at the time the agent makes the choice. For reinforcement learning models, prediction errors—between expected and experienced rewards associated with a state change—are denoted $E(s, s')$. Additionally, when a value, reward or prediction error includes an anticipatory component, the relevant functions are denoted using $\tilde{V}$, $\tilde{R}$ or $\tilde{E}$.

Figure 3: Notation used to describe the events within a single trial in the secrets task. When the agent makes a decision the trial enters an action state s_a , which is either s_f if the agent chooses to find out now, or s_k if they decide to keep it secret. This is then followed by a cue state s_c , in which the agent is presented a cue that may signal a reward (state s_+), signal no reward (state s_-) or be uninformative (state $s_?$). This is then followed by wait states that span some delay interval d , then finally an outcome state s_o is reached. Irrespective of which decision the agent made, this outcome state will be a win state s_w with probability q , or a loss state s_l with probability $1 - q$.

To maintain consistency with previous work, the free parameters in each model are expressed in the notation from the original papers.

5.1 Reward Prediction Error with Anticipation

The first model we consider is the reinforcement learning model introduced by Iigaya et al. (2016), which they referred to as a *reward prediction error model with anticipation* (RPE-A). The central assumption of this model is that observers savour the subjective utility from the anticipation of future events (Loewenstein, 1987; Story et al., 2013), and allows the savouring itself to serve as a reinforcer.

The formal structure of the model is as follows. On any given trial the agent assigns a predicted reward value $V(s)$ to the two choice states s_f and s_k , and makes decisions in accordance with a softmax rule

$$P(s_f) = \frac{1}{1 + \exp(\beta(V(s_k) - V(s_f)))} \quad (1)$$

where β is a free parameter in the model that governs the agent's willingness to choose a lower value option. After making the choice s_a that eventually leads to outcome s_o and yields an accumulated, temporally discounted reward $R(s_a, s_o)$ that is experienced by the agent, the value assigned to that action state s_a is updated according to the following rule:

$$V(s_a) \leftarrow V(s_a) + \alpha E(s_a, s_o) \quad (2)$$

where the reward prediction error is the difference

$$E(s_a, s_o) = R(s_a, s_o) - V(s_a) \quad (3)$$

This expression describes a standard reward prediction error learning rule with α denoting a learning rate parameter: the agent had assigned value $V(s_a)$ to this state, and adjusts this slightly to more closely match the experienced reward $R(s_a, s_o)$.

Where the RPE-A model departs from more traditional reinforcement learning accounts, the total reward received has two distinct components: there is a primary reward $R(\cdot, s_o)$ corresponding to the temporally-discounted reward from the outcome itself, but there is also an anticipatory reward $\tilde{R}(s_a, s_o)$ experienced during the wait states that last for the delay length d , because the agent savors (or dreads) the arrival of the outcome, and this also is subject to temporal discounting. Formally, the present value (at the time of

choice) of total reward to be experienced by the agent is given by the weighted sum

$$R(s_a, s_o) = R(\cdot, s_o) + \eta \tilde{R}(s_a, s_o) \quad (4)$$

where we use $\tilde{R}$ to indicate that the experienced reward is anticipatory and

$$\eta = \eta_0 + c |E(s_a, s_c)| \quad (5)$$

In this expression the “baseline anticipation” parameter η_0 and the “error boosting” parameter c are both free parameters of the model, and $E(s_a, s_c)$ denotes the reward prediction error experienced when the cue state is revealed. For example, if the agent knows that a keep it secret choice s_k is always followed by the ambiguous $s_?$ then there will be no prediction error when the cue is revealed, because $V(s_k) = V(s_?)$. However, when the agent decides to find out now some prediction error will be encountered because $V(s_f) < V(s_+)$ and $V(s_f) > V(s_-)$. As such the magnitude of the prediction error will be greater when the agent chooses to find out now, $|E(s_f, s_c)| > |E(s_k, s_c)|$, simply by virtue of the fact that information is revealed at the cue stage when a find out now decision is made. This “boosting” mechanism is central to the RPE-A model. It asserts that the savouring is more intense when it follows a surprising cue than when it follows a predictable cue.

To complete the model, we require expressions for the reward prediction error $E(s_a, s_c)$ that follows the cue, as well as primary reward value $R(\cdot, s_o)$ and the anticipatory reward value $\tilde{R}(s_a, s_o)$. As noted earlier these two quantities need to be assessed at the time of the choice, so temporal discounting mechanism is applied. Thus, if $R_{\dagger}(\cdot, s_o)$ denotes the primary reward value of the outcome state s_o at the time it is delivered, the relevant discounted reward after a delay of length d is given by

$$R(\cdot, s_o) = R_{\dagger}(\cdot, s_o) e^{-\gamma d} \quad (6)$$

and as noted by Iigaya et al. (2016), the present value at time of choice for the accumulated anticipation can be calculated analytically by integrating the instantaneous time-discounted anticipatory reward across the duration of the delay d , thereby marginalising the wait states. If the momentary strength of anticipatory savouring rises over time with rate ν , the integral yields

$$\tilde{R}(s_a, s_o) = (\eta_0 + c |E(s_a, s_c)|) \frac{R_{\uparrow}(\cdot, s_o)}{\nu - \gamma} (e^{-\gamma d} - e^{-\nu d}) \quad (7)$$

(see original paper for details). Finally, regarding the cue prediction error itself Iigaya et al. (2016) derived an analytic expression for $E(s_a, s_c)$ which holds whenever the agent has learned the true choice-cue association (i.e., the probability q of seeing a particular cue following the choice),

$$E(s_a, s_c) = \frac{(1 - q)(\eta_0 \tilde{R}(s_a, s_o) + R(\cdot, s_o))}{1 - (1 - q)c \tilde{R}(s_a, s_o)} \quad (8)$$

Note that the expressions for the anticipatory reward $\tilde{R}(s_a, s_o)$ and the cue prediction error $E(s_a, s_c)$ depend on one another, but as noted by Iigaya et al. (2016) there are generally stable solutions to this system of equations: our implementation used fixed point iteration to find the stable values.

The predictions of the RPE-A model is a consequence of three distinct mechanisms: the amount of momentary anticipation rises over time, the value of future rewards (both primary and anticipatory) is discounted as a function of delay, and that the amount of anticipation experience is larger when it follows a surprising information signal (the cue). The qualitative effect of these mechanisms is illustrated in the left panel of Figure 4.

Figure 4: Qualitative predictions of RPE-A, APE, and UP.

5.2 Anticipated Prediction Error

The second model we consider is a reinforcement learning model that assumes anticipation is the primary driver of information seeking behaviour but – unlike the previous model – it does not incorporate any notion of savoring a future reward. Rather, the theoretical claim is that the agent makes decisions not merely on the perceived value associated with an action, but also on the (attention weighted) anticipated learning signal that they will receive on the basis of that choice. Formally, the model assumes that the decision weight $\tilde{V}(s_a)$ associated with the choice action s_a is given by

$$\tilde{V}(s_a) = V(s_a) + \tilde{E}(s_a, s_c) \quad (9)$$

where $V(s_a)$ denotes the temporally discounted reward value associated with state s , and $\tilde{E}(s_a, s_c)$ denotes the anticipated value of the prediction error that will occur when the cue is revealed.

According to the APE model, the subjective values associated with each state in the task are learned by a standard temporal difference learning mechanism.

$$V(s_a) \leftarrow V(s_a) + \alpha E(s_a, s_o) \quad (10)$$

375 where α is the learning rate parameter and $E(s_a, s_o)$ is the reward prediction error. For the
 376 APE model in this task – where the outcomes are independent of the learner actions and
 377 no savouring is built into the subjective reward – the prediction error is

$$E(s_a, s_o) = R(\cdot, s_o) - V(s_a) \quad (11)$$

378 where $R(\cdot, s_o)$ is the temporally-discounted value of the reward to be received at the end of
 379 the delay period. Like the RPE-A model the discounting function is presumed to be
 380 exponential in nature, but is parameterized differently

$$R(\cdot, s_o) = \gamma^d R_{\dagger}(\cdot, s_o) \quad (12)$$

381 where d is the delay, γ is a discount parameter, and $R_{\dagger}(\cdot, s_o)$ is the primary reward value of
 382 the outcome s_o .

383 Because the expected reward is independent of the agent’s decision, the long run
 384 behaviour of this learning rule ensures that the reward value $V(s_f)$ of the “find out now”
 385 state s_f is the same as the value $V(s_k)$ of the “keep it secret” state. The reward signal
 386 itself provides no reason for the agent to prefer one choice over the other. Any preference
 387 that appears in the model is entirely due to the anticipatory component of the decision
 388 weight, $\tilde{E}(s_a, s_c)$.

389 As noted earlier, the APE model does not incorporate any inherent notion of
 390 savoring, and instead adopts the perspective that the agent constructs a myopic model of
 391 the task environment that allows them to seek out future learning signals. Specifically,
 392 what the learner does is anticipate every possible prediction error that might be
 393 encountered one step into the future. Formally, the anticipated prediction error is given by

$$\tilde{E}(s_a, s_c) = \sum_{s_c} P(s_c | s_a) w(s_c) (\gamma^d V(s_c) - V(s_a)) \quad (13)$$

In this expression, $P(s_c|s_a)$ denotes the probability that cue state s_c will follow the action state s_a . For instance the the probability that the positive cue state s_+ follows the find out now state s_f is equal to the reward probability q , and the probability that the neutral cue state s_0 follows the keep it secret state s_k is 1. The $\gamma^d V(s_c) - V(s_a)$ term denotes the (temporally discounted) reward prediction error that would be encountered if state s_c does in fact follow state s_a .³ Finally, the value of $w(s_c)$ is an attention weight parameter whose value denotes the extent to which the learner’s anticipation process is “focused” on the possibility of reaching state s_c .

Because the design of the secrets task is simple, it is possible to construct analytic expressions that correspond to the difference $\tilde{V}(s_k) - \tilde{V}(s_f)$ for the two choices, as described by Zhu et al. (2017), which makes implementation of the model considerably simpler. For our simulations, we apply the following equation:

$$\tilde{V}(s_k) - \tilde{V}(s_f) = (w(s_w) - w(s_l)) \frac{\gamma^d |R(\cdot, s_w) + R(\cdot, s_l)|}{2}, \quad (14)$$

where s_w and s_l refer to win and lose states respectively.

As with the other models, the probability of choosing to find out now is determined by a softmax decision rule applied to the decision weights $\tilde{V}$

$$P(s_f) = \frac{1}{1 + \exp\left(\beta(\tilde{V}(s_k) - \tilde{V}(s_f))\right)} \quad (15)$$

An illustration of the typical profile of choice probabilities as a function of delay length d is shown in the middle panel of Figure 4b.

³More generally, if the delay between making the choice and observing the cue is anything other than 1, the γ term would be raised to an exponent that reflects that delay length.

5.3 Uncertainty Penalty Model

The *uncertainty penalty* (UP) model differs from the anticipation based models in two respects, one theoretical and the other technical. Theoretically, instead of anticipating future rewards (like RPE-A) or future reward prediction errors (like APE), the agent is presumed to have an inherent information preference: uncertainty is assumed to be aversive and agents should prefer to receive advance information that resolves uncertainty faster than if the information was withheld. More formally, the outcome uncertainty $U(s)$ associated with any state s is given by the entropy function

$$U(s) = -P(s_w|s) \log_2 P(s_w|s) - P(s_l|s) \log_2 P(s_l|s), \quad (16)$$

where $P(s_w|s)$ and $P(s_l|s)$ are the respective probabilities of (eventually) reaching the winning state s_w or the losing state s_l given that the agent is currently in state s . This function is maximised when the probabilities of winning and losing are equal, and minimised (i.e., zero) either the probability of winning or losing is equal to one.

At a technical level, the UP model differs slightly in how value functions are computed. The RPE-A and APE models use iterative update rules to learn the value of a state $V(s)$ based on prediction error signals. The UP model does not specify a specific learning rule and simply notes that the optimal value function must satisfy Bellman's equation, which can be solved using dynamic programming in the general case. For the present purposes, it suffices to note that once the choice is made the events of a secrets trial are independent of the agent. In this case, the value $V(s)$ of any state in the task is given by

$$V(s) = \sum_{s'} P(s'|s) (R_+(s, s') + V(s') \exp(-kU(s'))) \quad (17)$$

where $P(s'|s)$ denotes the probability that the next state will be s' , $R_+(s, s')$ is the primary reward received when the agent leaves state s to arrive at state s' (which is zero unless s' is the win state s_w), and $V(s')$ is the value of the next state. In this expression the free

Table 4: Features of qualitative predictions of information-seeking models in valence sensitivity and increasing levels of delay length.

Feature	RPE-A	APE	UP
Delay Effect	Positive skew	Monotonic decreasing	Monotonic increasing
Information Avoidance	Yes	Yes	No

parameter k is a weighting assigned to the agent’s uncertainty $U(s')$, but in this task it functions in a fashion that is broadly similar to the temporal discounting parameter used in the other models. As with previous models, we apply a softmax rule to calculate choice probabilities

$$P(s_f) = \frac{1}{1 + \exp(\beta(V(s_k) - V(s_f)))} \quad (18)$$

where β is a free parameter in the model.

The qualitative behaviour of the UP model is illustrated in the right panel of Figure 4. Because the UP model relies on the uncertainty in the outcome independent of whether that outcome is good or bad, it generates the same curves for both positive and negative rewards. These curves are positively sloped because the uncertainty penalty associated with the entropy is applied repeatedly during the delay: the longer one has to wait for uncertainty to be resolved, the more aversive that uncertainty becomes.

Ultimately, each model offers differing accounts of how cue-outcome delay affects information seeking behaviour, as well as how information avoidant behaviour can be produced (if at all). The key features of the qualitative behaviour of each model are summarised in Table 4.

6 Model Fitting

We fit the RPE-A, APE, and UP models to the average FON choice proportions in each experiment. For Experiments 1 and 2, we fitted the models separately to each gain (or

chocolate) and loss (or sound) condition and present them here as distinct fits for clarity of illustration. We also conducted individual-level model analysis by fitting each model to each participant’s data and measuring the proportion of each sample to which a model was best fit.

- For the RPE-A model, we simulated 100 sessions of 200 trials to obtain stable measures of average choice probability of the informative cue. Six parameters were freely estimated: the temporal discounting parameter γ , the rate of anticipation gain ν , the relative baseline weight of anticipation η_0 , the weight of prediction error boost c , the learning rate α and the softmax determinism parameter β .
- For the APE model, we varied four parameters: the temporal discounting parameter γ , the attention weight towards the winning cue $w_+ = w(s_+)$, the attention weight towards the losing cue $w_- = w(s_-)$, and the softmax determinism parameter β .
- For the UP model, two parameters were estimated: the weight placed on uncertainty penalty k , and the determinism parameter for the choice rule β .

For the RPE-A model, reward values for positively valenced outcomes are specified with a value of 1, neutral outcomes are specified as 0, and negatively valenced outcomes are specified as -1. The APE and UP models are not sensitive to outcome valence, consequently for these models we fix the value of the more positive outcome to 1 and the less positive outcome to 0.

6.1 Model Fitting Results

The optimised model parameters are presented in Tables 5-7. Generally, these results seem to accord with our intuition—given that most data sets here do not show any evidence of temporal discounting, we would expect the role of the γ parameter for each anticipation model to be minimal. Note that the RPE-A uses γ as a negative exponential term,

Table 5: Optimised parameters for RPE-A for each data set

Data set	γ	ν	η_0	c	α	β
E1: Points Win	0.337	< 0.001	15.043	1.724	0.167	0.010
E1: Points Loss	1.073	5.92×10^{16}	< 0.001	1.014	< 0.001	0.908
E2: Chocolate	< 0.001	0.437	1.216	0.873	< 0.001	0.034
E2: Sound	89.028	$> 10^{100}$	< 0.001	0.004	1.000	5.480
E3: Mixed	8.397	$> 10^{100}$	< 0.001	0.089	< 0.001	2.562

Table 6: Optimised parameters for APE for each data set

Data set	γ	w_+	w_-	β
E1: Points Win	1.000	42.880	41.025	0.551
E1: Points Loss	1.000	-18.338	-39.124	0.050
E2: Chocolate	1.000	-9.970	-10.919	0.537
E2: Sound	1.000	17.198	-74.093	0.012
E3: Mixed	1.000	2.337	-5.963	0.110

consequently, values close to 0 indicate minimal influence of γ , and increasing values indicate stronger temporal discounting. For data sets where RPE-A makes reasonable predictions, these γ values appear close to zero. For data sets that involve negatively valenced stimuli, this is not the case. APE implements γ as a multiplicative term bounded between 0 and 1, with values closer to 0 indicating greater influence of temporal discounting. Across all data sets, the model is optimised at γ values very close to 1.

Model predictions are overlaid on behavioural data in Figure 5. A visual inspection of the model predictions indicates that the UP model appears to be producing the best fits

Table 7: Optimised parameters for UP for each data set

Dataset	k	β
E1: Points Win	0.556	1.132
E1: Points Loss	0.356	1.248
E2: Chocolate	0.017	2.754
E2: Sound	0.062	2.101
E3: Mixed	0.781	0.959

Figure 5: Behavioural data (large gray circles, vertical lines represent 1 standard error of the mean) with model predictions of mean FON preferences (black points and lines).

across all data sets. While RPE-A can generate comparable fits in data sets from conditions involving only positively valenced outcomes, the model struggles to predict data from conditions involving negatively valenced outcomes.

As indicated earlier, APE assumes no growth in anticipation over time, consequently preventing it from generating patterns of increasing information preference with increasing delay length. With such data, the best predictions it can generate resemble flat lines at the approximate average level of information seeking across all delay lengths. Although this is a poor fit to behavioural data showing delay-dependent behaviour (i.e., both conditions in Experiment 2), APE is a reasonably accurate model of data indicating delay-invariant information preference (Experiments 1 and 3).

Quantitatively, a comparison of BIC values (Table 8) reveals that the UP model is preferred across all data sets. Even in conditions where only positively valenced outcomes are involved, UP ultimately shows better performance due to its ability to account for the patterns in the data using fewer free parameters.

The strong performance of UP is also observed at the individual level where it is the

Table 8: Model comparison statistics across all experiments.

Dataset	BIC			BIC Weights		
	RPE-A	APE	UP	RPE-A	APE	UP
E1: Points Win	3213.53	3206.64	3187.0	0.00	0.00	1.00
E1: Points Loss	3373.81	3202.52	3177.1	0.00	0.00	1.00
E2: Chocolate	4484.81	4636.27	4477.1	0.02	0.00	0.98
E2: Sound	4998.15	4739.89	4613.1	0.00	0.00	1.00
E3: Mixed	4803.83	4612.30	4595.0	0.00	0.00	1.00

best-fitting model for the majority of participants in each data set (Figure 6). Although APE’s predictions of averaged FON choices were not particularly accurate, the model was able to strongly predict a minority of individual data. RPE-A was not found to best fit any individual data.

7 General Discussion

The observation of non-instrumental information seeking suggests information can be valued beyond how it can guide future decisions. While the literature indicates a general consensus that non-instrumental information seeking is robust under different situations, they deviate in their descriptions of how such behaviour can vary. Specifically, contemporary models disagree on 1) how the length of cue-outcome delay can affect magnitude of information seeking, and 2) how, if at all, information avoidance behaviour might be generated. Our work sought to address these two questions at an empirical and theoretical level.

Unlike predictions offered by prior research (e.g., Iigaya et al., 2016; Spetch, Belke, Barnet, Dunn, & Pierce, 1990) we did not observe any effect of temporal discounting on information preferences at long delay lengths. None of our empirical data indicated any significant decrease in information seeking behaviour as delay length increased. Following speculation by Iigaya et al. (2016), it is possible that such effects would only be observable

Figure 6: Proportion of participants best fit to each model for each data set. Note that RPE-A data is absent from the figure because it was not the best fit to any participant in any data set.

in additionally extreme lengths of delay. However, we remain skeptical of this position for two reasons. First, our current data do not indicate any negative trends with increasing delay lengths. Any extrapolation of current patterns indicate either delay invariance or a monotonically increasing preference for information. Second, our modelling work indicates that across all experiments, each model’s best account of the data involve specifications of their temporal discount parameter to a value indicating minimal influence. Model simulations with these specific parameter values cannot produce any significant pattern of temporal discounting at any arbitrarily long delay length. In essence, our best theoretical accounts indicate no evidence for temporal discounting in the current paradigm.

In Experiment 3 we found not only no evidence of temporal discounting, but also an absence of any clear delay dependency of information seeking behaviour. This absence is surprising because there is no clear mechanism to suggest how the lack of a neutral outcome can make seeking information less dependent on cue-outcome delays. More generally, the findings of Experiment 3 indicate that the neutral reward may be necessary in order to produce a monotonically increasing delay-effect—as seen in Iigaya et al. (2016) who also used neutral rewards. One possible explanation for this pattern is that with neutral rewards, people may have been more interested in obtaining information about gaining/losing nothing at long delays than short delays (possibly due to having to “endure” longer periods of time for nothing). However, current models do not immediately allow for a straightforward account of this interaction between neutral rewards and delay length, nor do they explain how this can differ between primary and secondary reinforcers.

Our research also provided no clear evidence of non-instrumental information avoidant behaviour in the current paradigm. Contrary to predictions by Iigaya et al. (2016); Zhu et al. (2017), the presence of losses as a possible outcome indicated primarily information seeking behaviour. The strongest hints of information avoidance was observed only in the chocolate condition of Experiment 2, and only specifically at very low delays (≤ 2 seconds). Even then, the average results for that data set indicated information

indifference as opposed to information avoidance. It is possible that robust non-instrumental information avoidance can be generated under different situations, however current models do not clearly indicate how this can occur.

Our modelling work ultimately indicates that when people seek non-instrumental information, they are seeking the resolution of uncertainty rather than savouring the anticipation of future outcomes. Unlike the latter, the former assumes no temporal discounting of rewards and no dependency of information avoidance on any specific mechanism, which is clearly observed in our data.

While both RPE-A and UP models are able to accurately reproduce patterns of information seeking seen in the data from conditions with only positively valenced outcomes, it is the relative simplicity of UP in these comparisons that makes this model particularly appealing. Without assuming multiple aspects to its core mechanism of information seeking (as is seen in RPE-A where anticipation is decomposed into a baseline, boost strength, and gain rate separately), UP elegantly describes information seeking using a single parameter (assuming minimal impact of β). Even though RPE-A can produce similar behaviour with only positively valenced rewards, it does so with a structurally and conceptually more complex design.

Although it is conceptually similar to RPE-A, the APE model struggled to reproduce the data to the same extent of the other models. In particular, its inability to accumulate anticipation over increasing delay length prevented it from providing a reasonable fit to averaged data from Experiment 2.

It is possible that these complexities of RPE-A as defined by the original authors are not completely necessary—we could imagine that the mechanistic reinforcement-learning structure is reduced into a single construct (with fewer free parameters) that captures the relevant amount of anticipation associated with a given choice. This simplified anticipation-based model may end up performing just as well as, or even better than, the UP model. However, this can only happen when no negatively

valenced rewards are available in the task. As long as the utility from anticipation is directly mapped onto outcome valence, it will not be possible for the model to predict the information seeking patterns observed with negatively valenced rewards. We could then go even further and allow the single anticipation-based construct to be free from the valence of rewards—that is, the model can be agnostic to outcome valence just as the UP model is. However, doing so risks re-defining the original RPE-A model into one that is nearly indistinguishable (at least, within our paradigm) from that of UP, at which point a model comparison between the two may no longer be diagnostic.

While our results challenge prior speculative predictions, empirically they align closely with findings from previous studies (e.g., Charpentier et al., 2018; Iigaya et al., 2016; Zhu et al., 2017), none of which have found reliable attenuation of information seeking at long delays or consistent information avoidant behaviour. The diverse background of participants from these various studies and ours (from university students to general members of the public) suggest that the observed patterns of behaviour are robust and not specific to any particular sample of adults. However, we note that these patterns have only been tested and observed within controlled laboratory studies, and consequently we have no evidence to suggest that such behaviour necessarily generalizes outside a controlled environment. We have no reason to believe that our specific results depend on other characteristics of the participants, materials, or context.

The apparent simplicity of the observed data compared to current theories also raises a separate issue: under which other circumstances (beyond cue-outcome delay) can we expect non-instrumental information seeking behaviour to change? While one way to explore this further would be to investigate other empirical means of producing non-instrumental information avoidance, alternative avenues include examining how quickly people can learn the relative optimality of choosing an informative vs. noninformative cue, or whether the absolute magnitude of rewards affects information seeking behaviour. Suffice to say, the robustness of general non-instrumental information

seeking is easily apparent, but the conditions under which (and relatedly, how) such behaviour systematically differs remains an important avenue for further investigation.

8 Data availability

The datasets analysed during the current study are available in the GitHub repository, <https://github.com/shixianliew/infoseekDelayDist>.

9 Code availability

The model scripts used during the current study are available in the GitHub repository, <https://github.com/shixianliew/infoseekDelayDist>.

10 Ethics declarations

10.1 Competing interests

The authors declare no competing interests.

10.2 Ethical compliance

All experiments followed all ethical guidelines and was reviewed by University of New South Wales Human Research Advisory Panel C: Behavioural Sciences with approval number 3205. Informed consent was obtained from all participants prior to participation.

References

- Beierholm, U. R., & Dayan, P. (2010). Pavlovian-instrumental interaction in ‘observing behavior’. *PLoS computational biology*, 6(9), e1000903.
- Bennett, D., Bode, S., Brydevall, M., Warren, H., & Murawski, C. (2016). Intrinsic valuation of information in decision making under uncertainty. *PLoS computational biology*, 12(7).
- Bennett, D., Sutcliffe, K., Tan, N. P.-J., Smillie, L. D., & Bode, S. (2020). Anxious and obsessive-compulsive traits are independently associated with valuation of noninstrumental information. *Journal of Experimental Psychology: General*.
- Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature neuroscience*, 14(9), 1209–1216.
- Cabrero, J. M. R., Zhu, J.-Q., & Ludvig, E. A. (2019). Costly curiosity: People pay a price to resolve an uncertain gamble early. *Behavioural processes*, 160, 20–25.
- Charpentier, C. J., Bromberg-Martin, E. S., & Sharot, T. (2018). Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences*, 115(31), E7255–E7264.
- Cox, T. J. (2008). Scraping sounds and disgusting noises. *Applied Acoustics*, 69(12), 1195–1204.
- Crockett, M. J., Braams, B. R., Clark, L., Tobler, P. N., Robbins, T. W., & Kalenscher, T. (2013). Restricting temptations: neural mechanisms of precommitment. *Neuron*, 79(2), 391–401.
- De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47(1), 1–12.
- Eliaz, K., & Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2), 304–324.
- Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of

consumption and asset returns: a theoretical framework. *Econometrica*, 57, 937–969.

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1), 96–135.

Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, reward, and information seeking. *Journal of Neuroscience*, 34(46), 15497–15504.

Hertwig, R., & Engel, C. (2020). Deliberate ignorance: Choosing not to know..

Iigaya, K., Hauser, T. U., Kurth-Nelson, Z., O’Doherty, J. P., Dayan, P., & Dolan, R. J. (2020). The value of what’s to come: Neural mechanisms coupling prediction error and the utility of anticipation. *Science advances*, 6(25), eaba3828.

Iigaya, K., Story, G. W., Kurth-Nelson, Z., Dolan, R. J., & Dayan, P. (2016). The modulation of savouring by prediction error and its effects on choice. *Elife*, 5, e13747.

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3), 449–460.

Kobayashi, K., Ravaoli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature human behaviour*, 3(6), 587–595.

Kreps, D. M., & Porteus, E. L. (1979). Dynamic choice theory and dynamic programming. *Econometrica: Journal of the Econometric Society*, 91–100.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387), 666–684.

Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore–exploit dilemma in static and dynamic environments. *Cognitive psychology*, 85, 43–77.

Pierson, E., & Goodman, N. (2014). Uncertainty and denial: a resource-rational model of the value of information. *PloS one*, 9(11), e113342.

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*, 1–6.

Spetch, M. L., Belke, T. W., Barnet, R. C., Dunn, R., & Pierce, W. D. (1990). Subop-

- 667 timal choice in a percentage-reinforcement procedure: Effects of signal condition and
 668 terminal-link length. *Journal of the experimental analysis of behavior*, 53(2), 219–234.
- 669 Story, G. W., Vlaev, I., Seymour, B., Winston, J. S., Darzi, A., & Dolan, R. J. (2013).
 670 Dread and the disvalue of future pain. *PLoS computational biology*, 9(11).
- 671 Vasconcelos, M., Monteiro, T., & Kacelnik, A. (2015). Irrational choice and the value of
 672 information. *Scientific reports*, 5(1), 1–12.
- 673 Zhu, J.-Q., Xiang, W., & Ludvig, E. A. (2017). Information seeking as chasing anticipated
 674 prediction errors. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.),
 675 *Proceedings of the 39th annual conference of the cognitive science society* (pp. 3658–
 676 3663). Austin, TX: Cognitive Science Society.