

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Do you want to know a secret? The role of valence and delay in early information preference

Permalink

<https://escholarship.org/uc/item/7xw3m02x>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

Embrey, Jake

Liew, Shi Xian

Navarro, Danielle

et al.

Publication Date

2020

Peer reviewed

Do you want to know a secret? The role of valence and delay in early information preference

Jake R. Embrey (j.embrey@unsw.edu.au), Shi Xian Liew (s.liew@unsw.edu.au),
Danielle Navarro (d.navarro@unsw.edu.au), & Ben R. Newell (ben.newell@unsw.edu.au)

School of Psychology, UNSW, Sydney, NSW 2052, Australia

Abstract

People tend to place value on information even when it does not affect the outcome of a decision. Two competing accounts offer explanations for such non-instrumental information seeking. One account foregrounds the role of anticipation and the other focusses on uncertainty aversion. Both accounts make similar predictions for short cue-outcome delays and when outcomes are positively valenced, but they differ in their explanation of information preference at long delays with negative outcomes. We present a series of experiments involving both primary and secondary reinforcers that pit these accounts against each other. The results indicate a consistent preference for non-instrumental information even at long cue-outcome delays and no evidence for information avoidance with negative outcomes. This pattern appears to provide more support for the uncertainty-aversion account than one based on anticipation.

Keywords: information seeking, uncertainty; anticipation; temporal discounting; valence; delay

Introduction

Humans are often described as *informativores* (Gigerenzer & Garcia-Retamero, 2017; Hertwig & Engel, 2016; Miller, 1983): we scroll smart phones consuming meaningless media, read trivial internet comment sections, and refuse to leave even the worst movies so we can find out how they end. This tendency to want-to-know spills over into situations when the information we seek has no “instrumental” value, such as navigating to a new website and looking up the price of an item just purchased. Such behaviours challenge standard decision making accounts (Hirshleifer & Riley, 1979; Raiffa & Schlaifer, 1961) that are premised on the notion that information is only valuable for its ability to inform future decisions (i.e., its instrumentality).

Recent theories of non-instrumental information seeking can be broadly classified into two accounts. The anticipation account posits that “non-instrumental” information affords the subjective value of anticipating a reward or its related prediction error (Iigaya, Story, Kurth-Nelson, Dolan, & Dayan, 2016; Zhu, Xiang, & Ludvig, 2017). Conversely, the uncertainty aversion account describes information seeking in terms of the resolution of temporal uncertainty (Bennett et al., 2016).

These two accounts have successfully predicted a number of empirically observed effects in information seeking. In cases where a predetermined reward is delayed, but information about it is available at no cost, people show strong preferences to ‘find out’ about the nature of the reward prior to its arrival (Bennett et al., 2016, Iigaya et al., 2016, Zhu et al., 2017). This preference typically increases

with the length of delay between choice and outcome arrival (Iigaya et al., 2016) and decreases as the cost of obtaining information rises (Bennett et al., 2016).

The anticipatory and uncertainty aversion accounts, however, qualitatively diverge in their predictions when outcome valence is negative and at extended delay lengths.

The Anticipation Account

Anticipation-based models of information seeking assume that preference for advance information increases as the period for outcome anticipation builds. Importantly, they also assume that information preference is attenuated by a temporal discounting mechanism that increasingly reduces information preference with longer delays (Chapman, 1996). Consequently, preference for information is predicted to follow a unimodal trajectory: people should be indifferent between knowledge and ignorance when delays are short (e.g., 1-10 seconds) because there is very little time for anticipation to build, information preference should increase and peak for moderate delays (e.g., 20 - 30 seconds), before trending towards indifference at long delays where the preferential boost from anticipation is balanced with temporal discounting (Figure 1).

Empirical support for this account is mixed. In their own data, Iigaya et al. (2016) find partial support for the model using “erotic” images as a primary reward: at short delays participants were indifferent, but at longer delays participants had strong information preference. They found no evidence of any preference attenuation for the longest delays in their experiment (40 seconds), although this delay may simply not be long enough for the counteracting effect of discounting to take hold.

Anticipatory models of information-seeking also consider the reward valence in generating predictions. Under an anticipatory account, future negative outcomes should elicit “dread” and therefore people should avoid acquiring non-instrumental information; in contrast, if the effect is driven by uncertainty avoidance, people should still prefer to find out even if the outcomes are unpleasant (Figure 1).

Zhu et al. (2017) used “erotic” images for positive outcomes and “mutilation” images as aversive outcomes, and included three conditions (positive, negative, and mixed) at a fixed delay of 20 seconds. With positive outcomes, participants showed information preference; preferences attenuated when positive and negative trials were intermixed, and choice indifference was observed in the negative condition. Contrary to model predictions, information aversion was not observed in the negative condition. Moreover, the results were obtained from nine choice trials, with large individual differences observed.

Further evidence against this class of models is also seen in earlier literature showing *preference* for non-instrumental information about upcoming electric shocks (Averill & Rosenn, 1972; Lanzetta & Driscoll, 1966, 1968).

The Uncertainty Aversion Account

According to the uncertainty hypothesis (Bennett et al., 2016; Kreps & Porteus, 1978), people find uncertainty inherently aversive (Berlyne, 1960). Consequently, information is valued for its ability to resolve uncertainty regardless of valence (appetitive, neutral, or aversive). Uncertainty aversion models predict that preference for knowledge increases with delay as long periods of uncertainty are more aversive than short periods. These models however do not incorporate temporal discounting mechanisms¹ and predict a monotonic increase in information preference as delay increases (Figure 1). Generally, these models are consistent with the data from both Iigaya et al. (2016) and Bennett et al., (2016) which used primary (erotic images) and secondary (money) rewards respectively, and is also consistent with the older literature (Lanzetta & Driscoll, 1966, 1968).

In this paper, we compare the anticipation and uncertainty accounts by examining whether a longer delay can yield observable effects of temporal discounting on information preference across different outcome valences. In addition, we conduct these investigations using secondary (e.g., monetary reward) and primary (e.g., food) reinforcers. Our focus here is on behavioural data patterns; we leave formal model analyses for future publications.

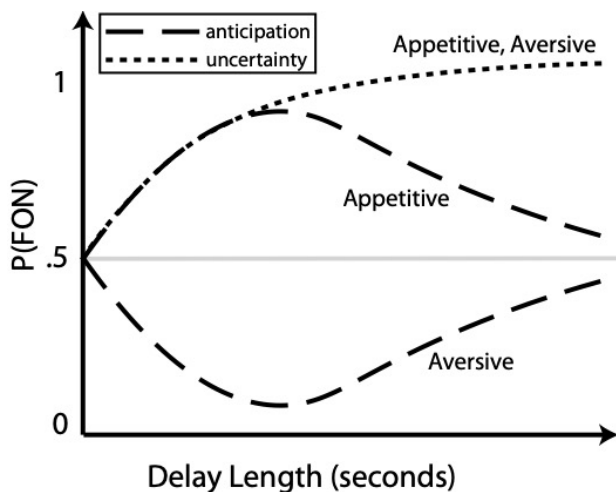


Figure 1. Qualitative predictions of the probability of choosing the informative “Find Out Now” option for the anticipation-based hypothesis (discontinuous lines with long segments) and uncertainty-based hypothesis (discontinuous lines with short segments) with different reward valences. Note that the uncertainty-based hypothesis does not distinguish between reward valence.

¹ Bennett et al. (2016) briefly consider a variant of their model with temporal discounting, but their mechanism functions across the presentation duration of the cue, as opposed to the duration between cue and outcome which we focus on in this paper.

Experiment 1

Experiment 1 employs a “secrets” task following Iigaya et al. (2016) and Zhu et al. (2017), but used points/money as the outcome (a secondary reinforcer). In the task participants could learn whether they were going to win points, lose points, or receive 0 points. Delay length between choice and outcome was varied within-subjects.

Method

Participants Forty undergraduate students from the UNSW psychology cohort ($M_{age} = 20.13$ years, 25 females, 15 males) participated in exchange for course credit, and were paid an amount depending on the points awarded ($M = \$3.50$ AUD).

Materials The experiment was implemented in jsPsych (de Leeuw, 2015) within Google Chrome on a desktop.

Design On each trial participants were presented with two options, labelled “Find Out Now” and “Keep It Secret”, and information on the delay length. We used seven delay lengths (1, 2, 5, 10, 20, 40, & 80 s), randomly intermixed throughout the task. Each participant completed two blocks of trials in a random order. In the win block, the outcome was 0 points or ~100 points (uniformly distributed between 90-110 points); in the loss block, participants either received 0 points or lost ~100 points. There was an equal chance of receiving either of the two outcomes (win/loss or zero points), regardless of choice. If the participant chose “Keep It Secret” a non-informative cue was immediately shown, and the outcome revealed after the appropriate delay. In the win block, if they chose “Find Out Now” they would be immediately shown a smiley face if the outcome was positive, and a sad face if the subsequent outcome was neutral (0 points), with the outcome again shown after the appropriate delay. In the loss block, the sad face signalled the later arrival of a loss and the smiley face signalled the zero-points outcome. This design is illustrated in Figure 2.

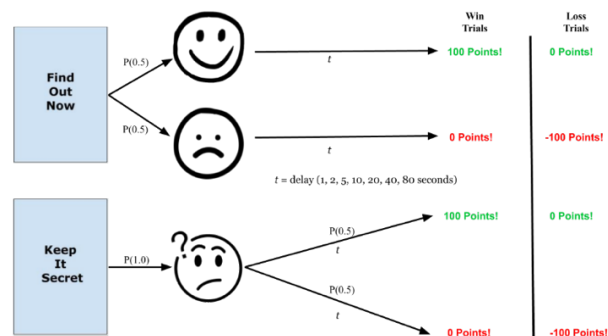


Figure 2. The win and loss conditions of the ‘secrets’ task. Choosing ‘Find Out Now’ provides the participant with early information. Choosing ‘Keep It Secret’ provides no information about the delayed outcomes.

Procedure All participants started the experiment with 3000 points to prevent the possibility of obtaining a

negative balance. There were 60 choice trials in each block (120 in total): 10 trials at the 1, 2, 5, 10, and 20 second delays, and five trials for 40 and 80 second delays. Participants were informed about the nature of the task and that their decisions did not determine the outcome; rather, the options allowed them to receive information regarding the outcome before the delay or keep it a secret. Cues were displayed immediately after their initial decision. After completion, participants were debriefed on the experiment's aims and reimbursed according to the amount of money they earned: 1000 points = \$1AUD.

Results

The results are displayed in Figure 3. Visual inspection of the results suggests an overall preference for the "Find Out Now" (FON) option that is constant across delay and unaffected by valence. Across all trials, this preference was significantly greater than chance ($M = .63$ $t(39) = 4.05$, 95% CI [.57, .69], $p < .001$; although there were individual differences – see Figure 3, right panel).

We investigated the effect of delay and valence using generalised linear mixed models (GLMM). The initial baseline model contained only random intercepts for individual participants to predict choice preference and was compared against models with random slopes (as a function of delay) and delay length as a fixed effect. This analysis found no evidence for an effect of delay or valence (all $p > .05$).

We also found a significant positive correlation between average choice preferences in win trials and loss trials ($r(38) = .74$, $p < .001$), suggesting that the mechanisms driving preference for information are related, regardless of the valence of the expected outcome.

Discussion

Experiment 1 replicates aspects of the Iigaya et al. (2016) data and earlier electric shock experiments (Lanzetta & Driscoll, 1966; 1968), that is, information preference is observed at longer delays for both positive and negative outcomes. Contrary to Zhu et al., (2017) we find no effect of outcome valence on this preference. Moreover, we do not replicate Iigaya et al.'s (2016) finding that delay influences choice preference as information is preferred to

ignorance regardless of delay. One possible explanation for this discrepancy is that previous studies used primary rewards (erotic images) as the outcome, whereas we used points/money, a secondary reward. A secondary reinforcer may not be sufficient for generating the effects of delay discounting and outcome valence in this context because money is not inherently rewarding, and is merely a stand in for primary reinforcers that can be obtained (by spending the money!) at a later point.

Experiment 2

To address this limitation of Experiment 1, an additional experiment was conducted using primary reinforcers as the outcomes. In Experiment 2 there were two conditions, one with an appetitive outcome, and the other with an aversive outcome. These conditions were run consecutively on different groups of participants but for convenience we report them as a single experiment (with appropriate caveats for cross-condition comparisons). The appetitive condition used chocolate M&Ms as the reward, the aversive condition used a high pitch aversive tone (microphone feedback). These two experiments were intended to be respective analogues of the loss and win block conditions from Experiment 1.

Participants Experiment 2 (sound) had 51 undergraduate psychology student participants ($M_{age} = 19.49$, 38 females, 13 males) who were granted course credit for their participation. Experiment 2 (chocolate) 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&Ms. We ensured participants had not previously completed Experiment 1 or other similar studies.

Materials The setup for the sound condition was analogous to the loss condition of Experiment 1 other than points being replaced with the following outcomes: 10 seconds of




Figure 3. Left – Mean choice proportions for 'Find Out Now' across delay lengths for win and loss trials in Experiment 1. Right – Individual mean preferences for 'Find Out Now' across delay for both trial type.

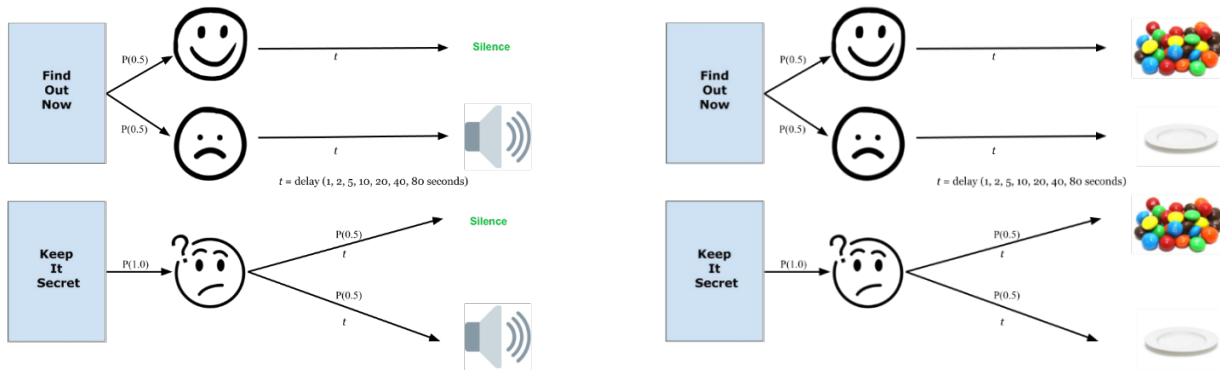


Figure 4. The designs of Experiment 2 (left: sound condition; right: chocolate condition). Regardless of choice, participants are equally likely to receive either of the two outcomes. Participants are merely choosing whether they want to know the delayed outcome prior to the delay or not.

microphone feedback (ranked 2nd most aversive from a list of 34 tested sounds; Cox, 2008) or silence (neutral). There was equal chance of receiving these outcomes regardless of choice. The chocolate condition was coded in MATLAB, displayed using PsychToolbox on a desktop computer, and was attached to an M&M dispenser. Participants either received an M&M or nothing, with equal probability. Participants were also told to consume the M&M immediately upon receipt.

Design We increased the number of trials at the “extreme” delays to 10 for each delay length (70 trials in total), otherwise the design was as in Experiment 1 (Figure 4).

Procedure Procedures were identical to Experiment 1, with the exception of not having any accumulation of points. In addition, the number of trials for the 40 and 80 second delays were increased from five to 10.

Results: Sound Condition

The results from Experiment 2 are displayed in Figure 5. We analysed the average preference for the FON option using a GLMM analysis analogous to that of Experiment 1.

Firstly, the addition of random slopes (varying across delay) significantly improved the fit of the model ($p < .001$), suggesting information preference across time varies significantly between participants. The addition of delay also significantly improved the model’s fit ($p < .01$) suggesting that delay length did have some effect in this experiment. 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 were mostly indifferent (preferences of .52 and .51 respectively), with preferences for FON increasing to .73 and .74 at the 40 and 80 second delays respectively.

We also inspected whether trial number had any influence on people’s preferences. We found a significant positive correlation between FON preference and trial number ($r = .43, p < .01$) with the preference to find out increasing with trial number, indicating that as people gained more experience with the task, their tendency to prefer early information increased.

Finally, we reiterate that there is considerable variation in individuals’ baseline preferences and the change in these

preferences across delays (Figure 5). While the majority of participants increased their preference as delay increased, there appear to be a subset of people who opted for ignorance as delay lengths increased.

Results: Chocolate Condition

Figure 6 displays the data for the chocolate condition. Similar to the sound condition, the addition of random slopes improved the model’s fit ($p < .001$). The addition of delay also significantly improved the model’s fit ($p < .01$), once again suggesting that delay length affected people’s information preferences. The pattern of results was roughly the same as in the sound condition: at 1 and 2 second delays, mean preference was .45 and .43 respectively, with these preferences increasing to .70 and .72 at 40 and 80 seconds respectively. In this condition, however, we did not find any correlation between trial number and choice preference ($r = .01, p = .73$).

As in the sound condition, we note that participants exhibit a wide range of individual preferences in this task (see Figure 6). Most show clear increases in their preference for early information as delays increase. However, a small subset show opposing preference trends. There are also a number of individuals who show a consistent, near-100% preference for early information regardless of delay.

Discussion: Experiment 2

Unlike the secondary reinforcement used in Experiment 1, both tasks using primary reinforcement found that delay had an effect on information preference. In this respect our 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 predicted by anticipatory models.

Our results diverge somewhat more sharply with this model when negative valence outcomes are considered (Experiment 2 sound). The theory posited by Iigaya et al. (2016) and Zhu et al. (2017) predicts information avoidance with aversive outcomes. We however find strong evidence that preferences for early information remain even when participants are expecting an aversive, primary outcome. This pattern of results is more consistent with the uncertainty aversion approach adopted by Bennett et al (2016).

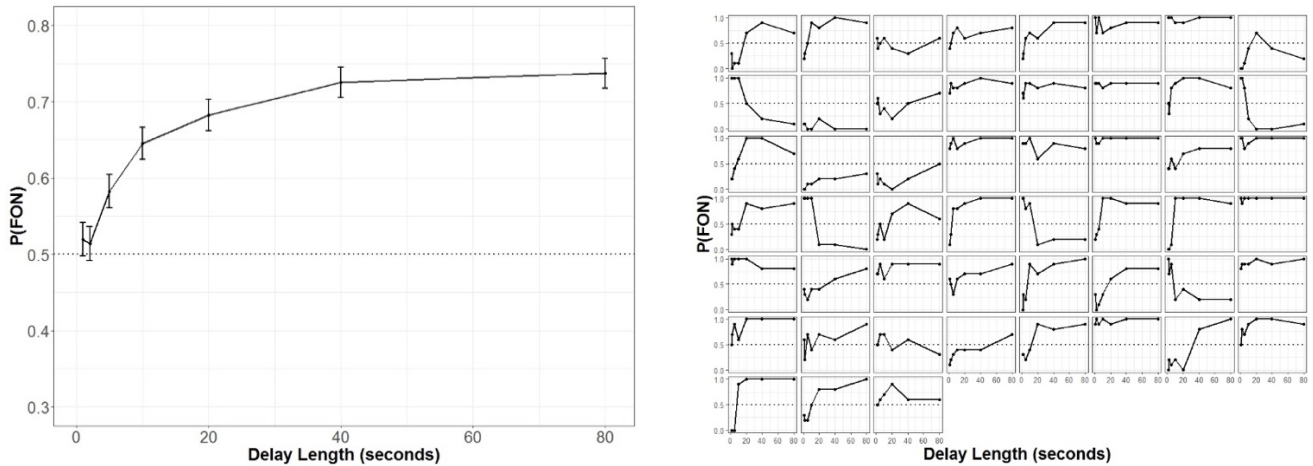


Figure 5. Sound condition of Experiment 2. Left – Mean choice preferences for Find Out Now averaged across participants for each delay length. Right – Individual choice preferences for the seven delay lengths.

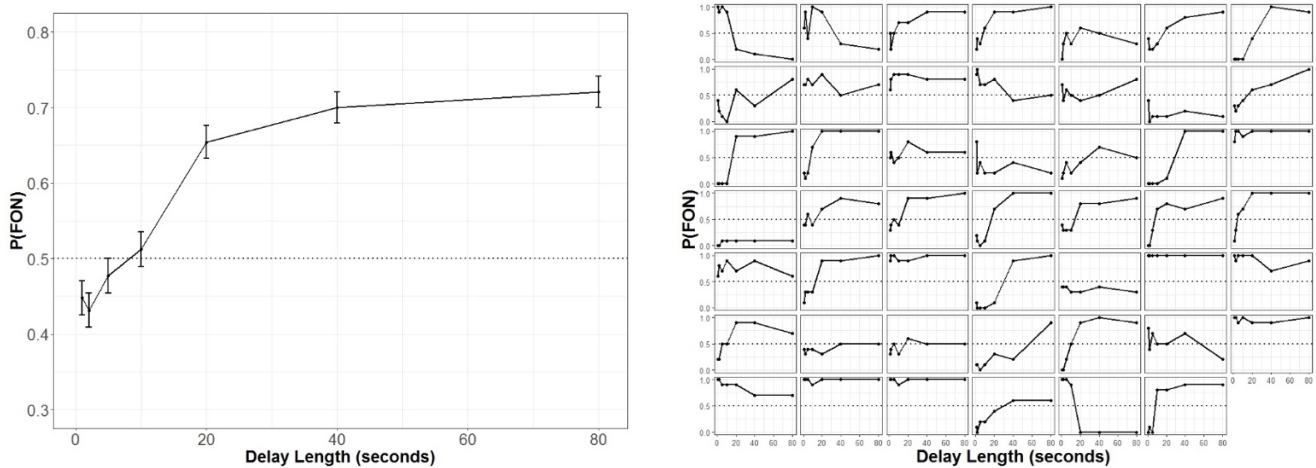


Figure 6. Chocolate condition of Experiment 2. Left – Mean choice proportions for Find Out Now averaged across participants for each delay length. Right – Individual choice proportions for Find Out Now across each delay length.

General Discussion

Since people seek information even when it cannot aid future decisions, information must itself be valuable. Separate accounts of this intrinsic value focus on uncertainty aversion or anticipation. These accounts produce conflicting predictions depending on the valence and the delay of outcomes. This paper aimed to elucidate the factors which influence people’s preferences for non-instrumental information in an experiential, decision-making setting.

At a qualitative level, the results of our two experiments seem most consistent with the uncertainty aversion account. That is, information is valued for its ability to resolve temporal uncertainty. Regardless of whether the expected outcome was positive (chocolate or monetary win) or negative (aversive sound or monetary loss), on average, participants exhibited a preference to ‘Find Out Now’ for intermediate and longer delays.

While participants presumably “anticipate” the outcome before its arrival (whether that be a positive, negative, or neutral outcome), this anticipation does not appear to be driving choice preference over and above a desire to resolve uncertainty. If it did, then participants

should avoid foreknowledge of negative outcomes, thereby attenuating the dread of a certain aversive impact. In contrast, we find a desire to know rather than be surprised by the delivery of an outcome, irrespective of its valence.

With regard to delay, we find evidence to support Igaya et al.’s (2016) findings that the length of delay influences people’s preference for early information about primary outcomes. The majority of participants increased their tendency to receive early information as the delay between choice and outcome increased. Despite trends for finding-out appearing to plateau at longer delays (i.e., 40 and 80 seconds) we did not observe an attenuation in choice preference at longer delays. While anticipatory accounts such as Igaya et al. (2016) hypothesise that reward savouring is attenuated by temporal discounting, we are not aware of any experimental demonstrations of this pattern. Future work will perhaps need to include even longer delays (in the order of minutes rather than seconds) in order to provide more robust tests of the anticipatory account.

Another potentially productive direction in which to take this research is to investigate the level of uncertainty associated with outcome delivery. Our experiments held the level of uncertainty constant at its maximum (i.e., 50/50

chance), but recent research (Charpentier et al., 2018; Iigaya et al., 2019) suggests the level of uncertainty plays a role in people's preference for resolution. In particular, Charpentier and colleagues' results suggest people's inclination for early information increases as the likelihood of a positive outcome increases. On the face of it, this finding is more aligned with anticipatory accounts than uncertainty aversion, because pure uncertainty aversion would predict maximum information seeking at maximum uncertainty. Modifications of the secrets-task would allow more specific tests of these predictions.

This series of experiments are also the first, to our knowledge, to explore the differential effects of primary and secondary rewards on information preference in the secrets-task. Primary rewards have previously been argued as necessary to exhibit delay effects (Crockett et al., 2013; Iigaya et al., 2016) because immediate consumption at the conclusion of the delay is necessary to fulfil the expected reward. Our results seem consistent with this claim: although there is a general desire for uncertainty resolution irrespective of reward-type and valence, the increasing desire with delay is only seen when the outcome has an immediate impact on participants' subjective state.

A clear challenge in advancing our understanding of how these factors of valence, delay and uncertainty drive non-instrumental information seeking is to develop quantitative computational models of observed behaviour. Several models already exist in the literature (e.g., Bennett et al., 2016; Charpentier et al., 2018; Iigaya et al., 2016; Zhu et al., 2017) but whether these models can capture the full gamut of findings across different paradigms remains to be seen. Further model development and careful comparison is an important next step for this work. For example, adjudicating between models that include a parameter for temporal discounting and those with an 'uncertainty penalty' would seem a fruitful approach. Such modelling, with the inclusion of uncertainty aversion scales (e.g., Intolerance of Uncertainty Scale; Carleton, Norton, & Asmundson, 2007) may also shed light on the robust individual differences in information preference we observe (Figures, 3, 5, 6 right panels).

There are certainly limits to the understanding of information preference that can be gained from a simple laboratory task like the one we use here (i.e., small monetary payments and relatively trivial rewards/penalties). Moreover, data from field studies and hypothetical questionnaires do appear to show robust tendencies for the active avoidance of information (Hertwig & Engel, 2016) – a pattern we do not see in our data. For example, only 55% of those tested for HIV return for their test results (Hightow et al., 2003) and only 3% to 25% of those at risk of carrying the Huntington's disease gene opt for the test (Creighton et al., 2003). Additionally, Gigerenzer and Garcia-Retamero (2017) found that over 85% of people do not wish to know the date of death for themselves or their partners (should such knowledge be available).

Predominantly, examples of these kinds of deliberate ignorance arise when potential future outcomes are *extremely* aversive (e.g., incurable illness or death) and thus go beyond the ethical constraints of experimental environments. Presumably this information avoidance is

driven by feared dread of the potential outcome, and in these instances the aversive quality of uncertainty weighs less heavily than the dread elicited by the possibility of knowing extremely negative information. Additionally, there are instances where ignorance itself provides utility (e.g., not knowing the final score of a football game you're watching); behaviours which are also difficult to reconcile with pure uncertainty aversion accounts.

While these examples of deliberate ignorance are largely found in case studies and questionnaires, a comprehensive understanding of information preference will need to address, and reconcile, these real-world observations with the findings obtained in the lab.

Acknowledgements

We would like to thank Gabrielle Kaplan for conducting the first experiment and analysing preliminary results. We would also like to thank Peter Lovibond and Jessica Lee for the M&M machine and corresponding code, and Jared Hotaling for making our experiment compatible with MATLAB. The project was supported by the Australian Research Council.

References

- Averill, J. R., & Rosenn, M. (1972). Vigilant and nonvigilant coping strategies and psychophysiological stress reactions during the anticipation of electric shock. *Journal of Personality and Social Psychology*, 23(1), 128-141.
- 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), e1005020.
- Berlyne, D. E. (1960). *Conflict, arousal, and curiosity*. New York: McGraw Hill Book Company.
- Carleton, R. N., Norton, M. P. J., & Asmundson, G. J. (2007). Fearing the unknown: A short version of the Intolerance of Uncertainty Scale. *Journal of anxiety disorders*, 21(1), 105-117.
- Chapman, G. B. (1996). Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 771-791.
- 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.
- Creighton, S., Almqvist, E. W., MacGregor, D., Fernandez, B., Hogg, H., Beis, J., & MacKenzie, J. (2003). Predictive, pre-natal and diagnostic genetic testing for Huntington's disease: the experience in Canada from 1987 to 2000. *Clinical genetics*, 63(6), 462-475.
- 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.
- Gigerenzer, G., & Garcia-Retamero, R. (2017). Cassandra's regret: The psychology of not wanting to know. *Psychological Review*, 124(2), 179-196.

- Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359-372.
- Hightow, L. B., Miller, W. C., Leone, P. A., Wohl, D., Smurzynski, M., & Kaplan, A. H. (2003). Failure to return for HIV posttest counseling in an STD clinic population. *AIDS Education and Prevention*, 15(3), 282-290.
- Hirshleifer, J., & Riley, J. G. (1979). The analytics of uncertainty and information-an expository survey. *Journal of economic literature*, 17(4), 1375-1421.
- Iigaya, K., Hauser, T. U., Kurth-Nelson, Z., O'Doherty, J. P., Dayan, P., & Dolan, R. J. (2019). The value of what's to come: neural mechanisms coupling prediction error and reward anticipation. *bioRxiv*, 588699.
- 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.
- Lanzetta, J. T., & Driscoll, J. M. (1966). Preference for information about an uncertain but unavoidable outcome. *Journal of Personality and Social Psychology*, 3(1), 96-102.
- Lanzetta, J. T., & Driscoll, J. M. (1968). Effects of uncertainty and importance on information search in decision making. *Journal of Personality and Social Psychology*, 10(4), 479-486.
- Miller, G. A. (1983). Informavores. In F. Machlup & U. Mansfield (Eds.), *The study of information: Interdisciplinary messages* (pp. 111-113). New York, NY: Wiley-Interscience.
- Raiffa, H., & Schlaifer, R. (1961). Applied statistical decision theory.
- Zhu, J. Q., Xiang, W., & Ludvig, E. A. (2017, April). *Information seeking as chasing anticipated prediction errors*. In T. T. E. J. D. G. Gunzelmann A. Howes (Ed.), *Proceedings of the 39th annual conference of the cognitive science society* (pp. 3658-3663). Austin, TX: Cognitive Science Society.