

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The effect of uncertainty and reward probability on information seeking behaviour

Permalink

<https://escholarship.org/uc/item/8512q7qt>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN

1069-7977

Authors

Embrey, Jake Ryan
Liew, Shi Xian
Ghai, Ishaan
[et al.](#)

Publication Date

2021

Peer reviewed

The effect of uncertainty and outcome probability on non-instrumental information seeking

Jake R. Embrey (j.embrey@unsw.edu.au), Shi Xian Liew (s.liew@unsw.edu.au),
Ishaan Ghai (ishaan.ghai02@gmail.com), & Ben R. Newell (ben.newell@unsw.edu.au)

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

Abstract

People's desire to seek or avoid information is not only influenced by the possible outcomes of an event, but the probability of those particular outcomes occurring. There are competing explanations however as to how and why people's desire for non-instrumental information is affected by factors including expected value, probability of outcome, and a unique formulation of *outcome uncertainty*. Over two experiments, we find that people's preference for non-instrumental information is positively correlated with probability when the outcome is positive (i.e., winning money) and negatively correlated when the outcome is negative (i.e., losing money). Furthermore, at the aggregate level, we find the probability of an outcome to be a better predictor of information preference than the expected value of the event or its *outcome uncertainty*.

Keywords: information seeking; information avoidance; uncertainty; anticipation; probability; expected value

Introduction

Animals are reliant upon information. Whether the purpose is to source food, sex, or money, information is often the means to a desired end. Classical theories of information-seeking (Hirshleifer & Riley, 1979; Raiffa & Schlaifer, 1961) posit that humans and other animals value information for its ability to inform and guide subsequent decisions – in other words, the information's instrumentality. Examples of this type of information search could be checking reviews before going to a restaurant, reading policies before voting, or checking the weather before leaving the house. Information such as this updates our understanding of the world and guides our behaviour accordingly.

Recent research (Bennett et al., 2016; Embrey, Liew, Navarro, & Newell, 2020; Igaya, Story, Kurth-Nelson, Dolan, & Dayan, 2016) has focussed on the finding that human and non-human animals tend to value information beyond its instrumental value. People seek information about upcoming monetary (Bennett et al., 2016) and food rewards (Embrey et al., 2020), as well as show preferences for information about negative outcomes like electric shocks (Averill & Rosenn, 1972; Lanzetta & Driscoll, 1966, 1968) and aversive sounds (Embrey et al., 2020), all in circumstances where they cannot use the obtained information to influence the eventual outcome.

While the aforementioned research has primarily used abstract laboratory tasks, the phenomenon of non-instrumental information seeking (or avoidance) is not wholly divorced from real world behaviour. For instance, many people want to know the answers to obscure trivia questions, seek answers to the newspaper's weekly crossword, or check emails for typos after they've been sent. Anecdotally, much of the information we actively

consume has little 'instrumental value' and wouldn't be considered valuable by classical information-seeking theories.

The majority of the literature investigating non-instrumental information seeking has focused on a simple paradigm involving an event with two equiprobable outcomes (e.g., winning money or not). In day-to-day life however, outcomes are rarely equiprobable (unless you spend your days flipping coins) and we often have insight into how likely the outcomes of an event are (e.g., the likelihood of winning the lottery).

Recent research (Charpentier, Bromberg-Martin, & Sharot, 2018) finds that the probability of an outcome's occurrence influences people's desire for information about the corresponding event. Specifically, their results suggest that the more likely a positive outcome (negative outcome) is to occur, the more likely someone is to seek (avoid) information about it.

While it has been demonstrated that outcome probability plays a role in information preference (Charpentier et al., 2018), the generality of this influence is unclear. For instance, when gambles have a very high probability of reward (e.g., 99% chance of winning \$100 vs. 1% chance of winning nothing), it is hard to imagine people would seek advance information about the result given the near-certainty of a win. It is therefore unclear to what extent reward probability accurately tracks information seeking behaviour and whether other factors are at play, such as expected value (EV).

For instance, research pitting outcome variance against EV (Kobayashi, Ravaioli, Baranès, Woodford, & Gottlieb, 2019) found EV to be a stronger motivation to seek information than outcome variance. Specifically, Kobayashi and colleagues (2019) found participants would rather know about an event with a high EV (e.g., 300 points) with a low variance (e.g., ± 10 points) than an event with a lower EV (e.g., 210 points) and high variance (e.g., ± 80 points). A finding which suggests EV plays an important role in people's desire (or distaste) for information.

To capture these variance-based measures of events we apply two different formalizations of uncertainty. The first, is the entropy of the event (e.g., see Bennett et al., 2016), where entropy is highest when an event has two equally likely outcomes (50-50) and reduces as the likelihood of either outcome approaches 100%. Formally:

$$Entropy = - \sum_{i=1}^2 P(x_i) \log_2 P(x_i) \quad (1)$$

where $P(x_1)$ is the probability of the first possible outcome (e.g., winning points) and $P(x_2)$ the probability of the alternate outcome (e.g., 0 points) in a binary outcome task.

Another way to capture the uncertainty of an event is to measure *outcome uncertainty* (OU) as formalized by van Lieshout and colleagues (2018); a measure which aims to capture the subjective experience of uncertainty. It considers not only the raw entropy of the event, but also the range of possible outcomes. Formally:

$$OU(X) = (- \sum_{i=1}^2 P(x_i) \log_2 P(x_i)) \times |R_1 - R_2| \quad (2)$$

where R_1 and R_2 are the point values associated with the two possible outcomes.

This formulation of OU by van Lieshout and colleagues (2018) is conceptually distinct from both EV and outcome probability as it combines the entropy of an event with the values of the event's possible outcomes. For example, an event with a 50-50 chance of winning \$1000 (or nothing) would have a greater OU value than an event with a 50-50 chance of winning \$100 (or nothing). This consolidation leads to novel predictions about human information seeking and previously outperformed alternative models in predicting information seeking behaviour (van Lieshout et al., (2018).

The current set of experiments aim to delineate the possible roles that the probability of outcome, EV, and entropy have on people's preference for non-instrumental information. Furthermore, we examine whether the specific formulation of subjective OU offered by van Lieshout et al. (2018) is sufficient for explaining observed behaviour.

Experiment 1

The first experiment aimed to establish the role that outcome probability has on people's preference for advanced, non-instrumental information. The experiment was a modification of the secrets task where participants have the option to either find out or keep secret the outcomes of a future event: in this instance, winning 100 points or receiving 0 points. Rather than there being an equal chance of either winning points or no points (i.e., a 50-50 chance), the probability of a win was manipulated within-participants (see also Iigaya et al., 2020 for a similar design).

Following Charpentier et al. (2018), we expected to see an increase in information preference as the probability of reward increased. However, if OU is instead the main motivator of information preference, we would expect to see greatest information preference on trials with close to 50% probability of winning (where OU is greatest) and least information preference at the probability extremes (where OU is the smallest) (see Eq. 1). If entropy drives information preference, maximum information preference will occur when both outcomes are equiprobable (i.e., 50-50) and reduce as the probability of either outcome increases towards 100%. The predictions for OU and entropy are therefore equivalent in Experiment 1 (as they would be for any design with fixed outcome values).

Method

Participants Fifty two participants were recruited via Amazon's Mechanical Turk (mTurk), six did not complete the experiment, leaving 46 who were included in the analysis ($M_{age} = 35.17$, 31 males, 15 females). Participants were paid \$6 USD for completing the experiment and an

added bonus for the points earned during the task ($M = \$3$ USD).

Materials The experiment was coded in jsPsych (de Leeuw, 2015) and participants completed the experiment on their own laptop or desktop. We also used the Intolerance of Uncertainty Scale (Carleton et al., 2007) as a measure of potential individual differences in information-preference.

Design On each trial in the task participants were required to choose whether to 'find out' or 'keep secret' the outcome of an event: whether they will win 100 points or receive 0 points after a 20 second delay.

When choosing Find Out Now (FON) participants were immediately presented with either a smiley face or sad face. A smiley face indicated they would win 100 points and a sad face indicated 0 points would be received after the delay. Conversely, if Keep It Secret (KIS) was chosen a confused face was always presented, providing no information as to whether points will be won or not. Regardless of the participant's choice, the number of points received was displayed on screen after 20 seconds.

The likelihood of winning points (and conversely, receiving 0 points) changed between trials. There were five probability levels (win:no-win) in the experiment (1:99, 25:75, 50:50, 75:25, 99:1) which participants were informed of via a pie chart on each trial (Figure 1). There were 10 trials of each probability level in the experiment (50 trials in total) and trials were randomised within-participants.



Figure 1. The participants' view of a trial with a 99% chance of winning 100 points and a 1% of zero points.

Procedure Participants began the experiment with 0 points and were told how the two options allowed them to receive or avoid early information about the delayed outcome. They were also told (truthfully) that their choice would not influence the outcome. After reading the instructions participants completed a short three-item multiple choice questionnaire to ensure they understood the instructions. Questions were simple, such as, "How long is the delay length?"

Participants completed the 50 trials in a randomised order. Once the decision-making task was finished, participants completed the Intolerance of Uncertainty Scale (IUS) comprised of 12 Likert-scale questions. Once both tasks were finished, participants received a written debrief and a code to return to mTurk where they were compensated for their time. Bonus payments were paid dependent on the points earned: 1000 points = \$1 USD.

Results

The results for Experiment 1 are displayed in Figure 2. The graph shows a positive relationship between the probability of winning 100 points and the average proportion of Find Out Now choices across participants.

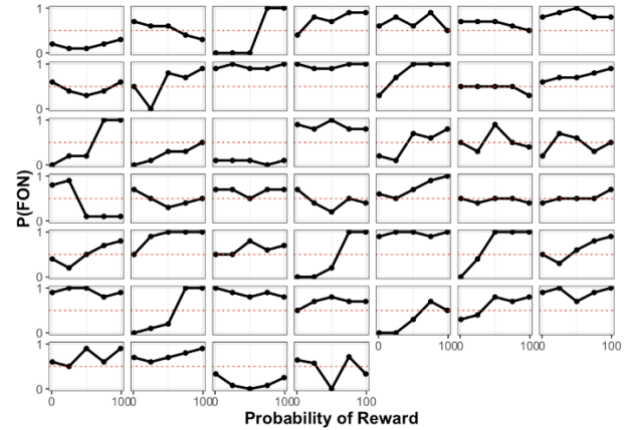



Figure 2. Left: The average proportion of FON choices for the five probability levels across all participants. Right: Individual FON preference across the five probability levels.

To analyse the observed relationship between outcome probability and choice proportion, we used generalised linear mixed models (GLMM) at the individual level. The baseline model contained only random intercepts and was compared to a model using intercepts and slopes, and a model using random intercepts and slopes, plus probability as a fixed factor.

Compared to the baseline model, the addition of random slopes significantly improved the fit ($p < .001$). Furthermore, the addition of probability as a fixed factor further improved fit to the data ($p < .001$, implying probability of reward was a strong determinant of people's choice preferences. No relationship was found between individual's choice preferences and their score on the IUS. The Pearson correlation between FON choice proportion (averaged across probability levels) and the IUS was $r = -0.05$ ($p = 0.78$).

Discussion

Experiment 1's results are commensurate with earlier findings (Charpentier et al., 2018) in that people's information preferences are positively related to the likelihood of winning reward. Furthermore, these results do not appear to support theories which argue demand for non-instrumental information is driven by an aversion to uncertainty/entropy (Bennett et al., 2016), nor are the results commensurate with the predictions of OU (van Lieshout et al., 2018).

The above results, at a qualitative level, are best predicted by the probability of the reward occurring: people's preference for information increases as the probability of the best outcome (winning 100 points) increases. In other words, participants seem to want information about events which are likely to be good (which they can savour) and are less likely to want information about events which are expectedly bad (which they may dread; see, Sharot & Sunstein, 2020).

While these results suggest a general dependency of information preference on outcome probability, Experiment 1's design conflates the influence of reward probability and EV on information preference. Specifically, as the probability of reward increases so does the EV of the event, making it difficult to determine which

one is the primary driver of the observed pattern of information preference. The design is also limited, as noted earlier, because it predicts indistinguishable patterns of information preference for OU and entropy.

Experiment 2

Experiment 2 builds on Experiment 1 but delineates probability of outcome from EV, allowing us to determine which has greater influence on people's information preferences. Here, expected value was held constant at 500 points while the probability of outcome fluctuated from 1% to 99%, analogous to experiment 1. By changing the probability of outcome while maintaining EV, we also systematically alter the OU value between each probability condition, such that OU decreases as the probability of the outcome increases. This design gives us four contrasting predictions, depending on which factor (EV, OU, entropy, or probability of outcome; Figure 4) is assumed the primary driver of outcome-predictive information seeking.

If EV drives preference, then we should not observe information preference differences between the five probability levels. If OU is important, then people's preference for information should reduce as the probability of outcome increases (see Design section for clarification). If entropy is the main driver, we should see an inverse U shape where information seeking is maximised at 50% and drops as either outcome becomes more likely. Finally, if the probability of outcome influences information preference, the results should be analogous to Experiment 1.

A loss condition was also added to test the prediction that information preference should diverge for win and losses of equivalent magnitude and probability. The win and loss conditions were run sequentially with different participants due to the length of the task (approximately 45 minutes). Participants in the win condition started with 0 points, whereas participants in the loss condition started with 50,000 points.

The IUS was not administered as no relationship was observed with FON preference in Experiment 1.

Participants Fifty participants signed up for Experiment 2 (win) via Amazon Mechanical Turk. Four participants were excluded from the analysis for failing the post-instruction quiz more than three times leaving 46 participants ($M_{age} =$

36.80; 16 females, 29 males, 1 other). Fifty-one participants enrolled in Experiment 2 (loss), but the data for five participants was excluded for failing the instruction-based quiz more than three times. Forty-six participants were therefore included in the analysis ($M_{age} = 35.73$; 18 females, 28 males). All participants were paid \$6 USD for completing the task and a bonus depending on their earnings throughout the task ($M = \$1.99$ USD).

Materials The materials in the second experiment were the same as Experiment 1 except the IUS was not administered.



Figure 3. The participants' view of a loss condition trial with a 50% chance of losing 1000 points and a 50% chance of zero points.

Design The five probability levels and their corresponding outcomes were as follows: a 1% chance of 50,000 points, 25% chance of 2,000 points, 50% chance of 1,000, 75% chance of 667 points, and 99% of 505 points. These points were either won or lost depending on which condition

participants were in (i.e., win or loss). The alternative outcome was 0 points. The expected value was 500 (or -500) points regardless of the probability of winning (or losing).

The numerical relationships between OU, entropy, EV, and the probability of win/loss are collated in Table 1.

Procedure In both the win and loss conditions the procedure was analogous to Experiment 1. Participants did not complete the IUS scale after the decision task.

Results

The results of Experiment 2 (for both win and loss conditions) are summarised in Figure 5. Inspection of the figures suggests that preference for information increases with probability as points are won and decreases with probability as points are lost. In other words, preference for information increases with the probability of the better outcome occurring (i.e., winning points in the win condition; 0 points in the loss condition).

Analogous to Experiment 1, a GLMM analysis was conducted. For the win condition, the addition of probability did not significantly improve the fit over a model with random slopes and intercepts ($p = .052$), but a model including random slopes and intercepts was better than a baseline model with only random intercepts. For the loss condition, a model including probability significantly improved fit over a model with random slopes and intercepts ($p = .024$).




Figure 4. Left: Predicted relationship (not to scale) between probability of win and OU, entropy, and EV for the win condition. Right: Predicted relationship (not to scale) between probability of loss and OU, entropy, and EV for the loss condition.

P(Win/Loss)	Entropy (bits)	EV (points)	OU
1%	0.08	500 / -500	4040
25%	0.81	500 / -500	1623
50%	1.00	500 / -500	1000
75%	0.81	500 / -500	540
99%	0.08	500 / -500	41

Table 1. The relationship between probability of a win (or loss) and entropy, EV, and OU in numerical terms. If separated by a slash (/) the left value corresponds to the win condition and the right to the loss condition, otherwise the values are the same for both (e.g., entropy).

Discussion

The results of Experiment 2 are similar to Experiment 1 in that proportions of FON choices fluctuate across probability levels, with a switch in outcome valence causing a switch in the relationship between probability of outcome (win or loss) and FON preference. These results however are less pronounced than in Experiment 1 for the win condition (i.e., differences in FON preference across probability levels was statistically non-significant).

Despite this, the experiments still illuminate which factors are important in driving information preference in the secrets task paradigm. Qualitatively, it appears that probability of outcome is best predictor of information preference. Moreover, van Lieshout and colleagues' (2018) formulation of OU does not appear to provide a reliable prediction of information preference in our participants. Uncertainty, formulated as entropy here, also appears to play little role in individuals' information preferences. The potential role of EV is considered further in the General Discussion.

General Discussion

The results of Experiments 1 and 2 are further evidence of people's willingness to seek information, even when it possesses no instrumental use. Furthermore, these two experiments enlighten and demarcate the factors important in determining whether people seek, avoid, or are indifferent to non-instrumental information about future events.

Earlier work has shown the importance of delay length (Iigaya et al., 2016), reward type (money or food; Embrey et al., 2020), and the probability of reward (Charpentier et al., 2018) on non-instrumental information preferences. Experiments 1 and 2 aimed to further assess the last of these findings, in particular, how and why fluctuating probabilities of rewards (or losses) influences people's desire for advanced information.

The results of Experiment 1 reflect similar findings by Charpentier and colleagues (2018) in that people's desire for information increases as the likelihood of reward increases and decreases as the likelihood of a negative outcome increases. Experiment 2 builds on this finding and also incorporates predictions by van Lieshout et al. (2018) which use a novel formulation of outcome uncertainty (OU). Here, rather than maximum uncertainty being experienced when outcomes are equally likely, subjective uncertainty is fundamentally tied with the magnitude of the possible outcomes and the event's entropy.

Experiment 2's results, while not entirely dismissive of the possible role of entropy and outcome value, are at odds with the predictions obtained by OU. If OU was a key driver of FON preference here, the relationship observed in the win condition would be inverted – most information seeking at a 1% chance of winning, least at 99% chance.

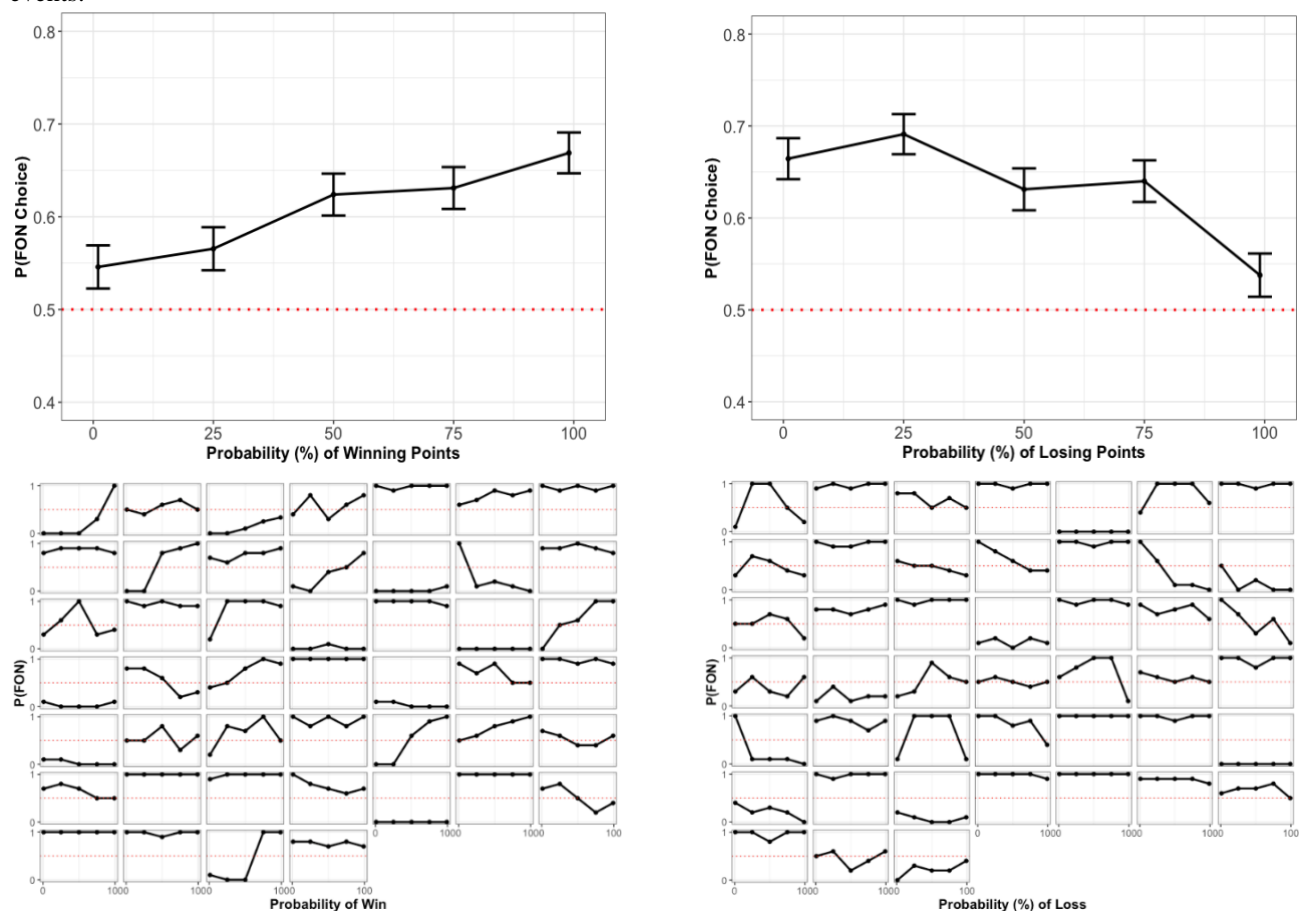


Figure 5. Left: average (top) and individual (bottom) FON preferences across levels of probability for the win condition. Right: average (top) and individual (bottom) FON preferences across levels of probability for the loss condition.

Moreover, while previous work has found entropy to play a major (Bennett et al., 2016) or minor role (Charpentier et al., 2018) in non-instrumental information seeking, no such role is observed here.

The influence of expected value (EV) however is less clear. Despite EV being held constant across probabilities, people's preferences for information still fluctuated when the likelihood of either winning or losing changed (Figure 5). Probabilities' influence however was not as marked in Experiment 2 as Experiment 1 which could be interpreted as EV having some role in people's preference for early outcome-predictive information. Comparing Experiment 1 to Experiment 2's win condition, FON preferences are higher at low probabilities of reward and lower at high probabilities; in Experiment 1, low probability trials had comparatively low EVs (1 point at 1% chance of winning) and high probability trials had comparatively high EVs (99 points at 99% chance of winning), whereas EV was constant in Experiment 2. Consistent with Kobayashi and colleagues' findings (2019),¹ this suggests increases in EV may boost information seeking. EV's possible role however would need to be relative to other experienced trials, as average information preference does not differ between win (EV = 500) and loss (EV = -500) conditions. Further research of the role of EV on non-instrumental information preference is therefore warranted.

With regard to the cognitive modelling of human information seeking behaviour, these results present new challenges. Liew and colleagues' (2021) analyses of different information seeking models find uncertainty aversion-based accounts better explain non-instrumental information seeking behaviours than anticipation-based models in a secrets task. It is not clear however how such a model would account for the behaviours observed here. Regardless of whether uncertainty is formalised as OU or entropy, our results are seemingly incompatible with such uncertainty-aversion based models. Conversely, we find information seeking is positively correlated with the likelihood of the better outcome.

The behaviour observed here is intuitively better aligned with anticipatory accounts. Broadly, anticipation is dependent on the expected outcome: good outcomes are savoured, and negative outcomes are dreaded. Formal anticipation-based accounts such as that proposed by Iigaya and colleagues (2020) can therefore account for the pattern of information seeking observed here: as the likelihood of a good outcome increases so does preference for information about the outcome. Adjustments however need to be made to account for the general lack of information avoidance observed in experimental tasks such as the secrets tasks.

In our current experiments and others (Iigaya et al., 2020; Charpentier et al., 2018) information seeking, on average, never drops below 50% across all levels of probability. In the current task for example, on average, participants show above 50% preferences for information, even when the likelihood of a win is 1% and the likelihood of a loss 99% (Figure 5). Furthermore, when outcomes are

negative primary reinforcers (i.e., electric shocks, aversive sounds), participants still show preferences for information (Averill & Rosenn, 1972; Embrey et al., 2020; Lanzetta & Driscoll, 1966, 1968) and at least indifference when the outcomes are 'mutilation' images (Zhu et al., 2017). Presumably, information seeking in these instances serves some function – such as bracing for the aversive outcome – nevertheless, evidence for information avoidance is often observed in case studies (Golman, Hagmann, & Loewenstein, 2017; Hertwig & Engel, 2016) and hypothetical questionnaires (Gigerenzer & Garcia-Retamero, 2017). Future work should investigate whether information avoidance can be induced in the lab and the specific factors responsible for strong information avoidance.

Future research should also aim to develop cognitive models which describe how factors such as entropy, probability, and expected value influence people's information preference and fit models at the individual level. Such work might also assess whether individual participants have stable traits which can be probed by questionnaires (other than the IUS, that we tried here) which might account for the observed heterogeneity in information preferences (see right panel of Figure 2 and bottom panels of Figure 5). For example, Bennett et al. (2020) found both anxiety and obsessive-compulsive traits to be associated with non-instrumental information seeking, and Ho, Hagmann, and Loewenstein (2021) have developed a scale specifically designed to measure people's desire or aversion to potentially unpleasant information about future events.

It should also be emphasised the analysis of the experiments presented here is largely qualitative and run at the aggregate level. A thorough analysis of the role probability, EV, entropy, and OU have on information preference should incorporate such factors into more sophisticated models (whether cognitive or not) and also assess behaviour at the individual level (which is highly heterogeneous).

Overall, these experiments strengthen the previously observed effect that probability of outcome has on people's information preferences in the secrets task. Moreover, it adds to the literature by demarcating the predictions of expected value from probability of outcome, and further tests the novel predictions of van Lieshout and colleagues' (2018) formulation of outcome uncertainty.

Acknowledgements

We would like to thank all members of the UNSW Cognition Lab for their feedback on this research during weekly lab meetings. Funding was provided by the Australian Research Council (grant number DP190101076).

References

Averill, J. R., & Rosenn, M. (1972). Vigilant and nonvigilant coping strategies and psychophysiological stress reactions during the anticipation of electric

¹ It should also be noted that Kobayashi and colleagues' (2019) design pitted two events (with different EVs) against one another – participants could only 'find out' the result of one event. In

this design, they found EV to be a strong predictor of which event participants would prefer information about.

- 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.
- 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*, 150(4), 739-755.
- 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.
- 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.
- Embrey, J. R., Liew, S. X., Navarro, D., & Newell, B. R. (2020). Do you want to know a secret? The role of valence and delay in early information preference. (pp. 273 – 279). Proceedings of the 42nd annual conference of the cognitive science society. Virtual: Cognitive Science Society.
- Gigerenzer, G., & Garcia-Retamero, R. (2017). Cassandra's regret: The psychology of not wanting to know. *Psychological Review*, 124(2), 179-196.
- Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1), 96-135.
- Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359-372.
- Hirshleifer, J., & Riley, J. G. (1979). The analytics of uncertainty and information-an expository survey. *Journal of economic literature*, 17(4), 1375-1421.
- Ho, E. H., Hagmann, D., & Loewenstein, G. (2021). Measuring information preferences. *Management Science*, 67(1), 126-145.
- 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.
- Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature human behaviour*, 3(6), 587-595.
- 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.
- Liew, S. X., Embrey, J. R., Navarro, D., & Newell, B. R. (2021, February 3). Comparing Anticipation and Uncertainty-Penalty Accounts of Non-Instrumental Information Seeking. <https://doi.org/10.31234/osf.io/nefpz>
- Raiffa, H., & Schlaifer, R. (1961). Applied statistical decision theory.
- Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. *Nature Human Behaviour*, 4(1), 14-19. doi:10.1038/s41562-019-0793-1
- van Lieshout, L. L., Vandenbroucke, A. R., Müller, N. C., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. *Journal of Neuroscience*, 38(10), 2579-2588.
- 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.