



Costly curiosity: People pay a price to resolve an uncertain gamble early

J.A. Max Rodriguez Cabrero, Jian-Qiao Zhu, Elliot A. Ludvig*

University of Warwick, United Kingdom



ARTICLE INFO

Keywords:

Sub-optimal choice
Curiosity
Information
Decision-making
Gambling

ABSTRACT

Humans are inherently curious creatures, continuously seeking out information about future outcomes. Such advance information is often valuable, potentially allowing people to select better courses of action. In non-human animals, this drive for information can be so strong that they forego food or water to find out a few seconds earlier whether an uncertain option will provide a reward. Here, we assess whether people will exhibit a similar sub-optimal preference for advance information. Participants played a card-flipping task where they were probabilistically rewarded based on the pattern of 3 cards that were revealed after a 5-s delay. During this delay, participants could instead pay a cost to find out the next card's identity immediately. This choice to find out early did not influence the eventual outcome. Participants preferred to find out early about 80% of the time when the information was free; they were even willing to incur an expense to get advance information about the eventual outcome. The expected magnitude of the outcome, however, had little impact on the likelihood of finding out early. These results suggest that humans, like animals, value non-instrumental information and will pay a price for such information, independent of its utility.

1. Introduction

Many animals are inherently curious, information-seeking creatures. Animals often gather information in close to an optimal fashion, where such information can lead to better reward harvests in the future (e.g., Charnov, 1976; Krebs et al., 1978). When faced with delayed, probabilistic rewards, however, many animals behave in a significantly sub-optimal manner. They sometimes prefer an informative option that reveals immediately whether the reward will pay off eventually, even if that information cannot be acted upon, and even to the point of giving up food or water to find out early (e.g., Blanchard et al., 2015; Bromberg-Martin and Hikosaka, 2009; Dunn and Spetch, 1990; McDevitt et al., 2018; Pisklak et al., 2015; Spetch et al., 1990; Stagner and Zentall, 2010; see McDevitt et al., 2016; Zentall, 2016 for reviews). The costs for obtaining such non-instrumental information can be broad, including food, water, energy, and time. In these cases, the animals seem to be behaving as though the information about future outcome is in itself rewarding.

One striking example of such suboptimal behaviour comes from experiments with pigeons working for signalled or unsignalled rewards (Dunn and Spetch, 1990; see also Spetch et al., 1990, 1994). In these experiments, pigeons were trained to choose between an option that led to food after a delay 100% of the time and a second option that led to food after the same delay only 50% of the time. In the unsignalled

condition, when there were no additional cues, the pigeons, unsurprisingly, showed almost exclusive preference for the optimal 100% option. In the signalled condition, however, additional cues were provided immediately after choice that indicated whether or not food would be available after the delay on that trial. In this case, pigeons showed up to a 50–60% preference for the sub-optimal 50% rewarded option—with a stronger preference for the suboptimal option with longer delays (Spetch et al., 1990; see also McDevitt et al., 2018). Similar sub-optimal choices have been observed with a range of other probabilities (e.g., Stagner and Zentall, 2010) and in several other species, including starlings (Vasconcelos et al., 2015) and monkeys (Blanchard et al., 2015).

Like other animals, humans also prefer to know in advance about future positive events (Loewenstein, 1987; Golman and Loewenstein, 2018). For example, in a pair of recent experiments where the rewards were images of attractive female models, people showed a strong preference to have information in advance about these rewards, a tendency which increased as the waiting time grew longer (Iigaya et al., 2016; Zhu et al., 2017). One possible explanation for this preference is that animals like signals for good news, but disregard signals for bad news (e.g., McDevitt et al., 2016). Similarly, when learning from such rewards, people tend to weight more heavily the positive information, and disregard negative information (Kuzmanovic et al., 2016; Lefebvre et al., 2017). This preference for good news in people may derive from

* Corresponding author at: Department of Psychology, University of Warwick, Coventry, CV4 7AL, United Kingdom.
E-mail address: E.Ludvig@warwick.ac.uk (E.A. Ludvig).

the savouring of future rewards, whereby people gain anticipatory utility from the experience preceding the actual reward (Loewenstein, 1987). Thus, knowing in advance about a future reward is preferable because that knowledge can be savoured, such as when people report a preference to kiss their favourite movie star with a delay of three days, rather than immediately. In contrast, people will avoid information about negative outcomes (e.g., distressful images; Zhu et al., 2017) as though they gain negative anticipatory utility from the potential dread elicited by bad news about future aversive outcomes (Charpentier et al., 2018; Golman et al., 2017; Karlsson et al., 2009).

In a probabilistic reward task similar to what has been used with monkeys (Bromberg-Martin and Hikosaka, 2009), people also preferred to get advance, but non-instrumental, information about future rewards (Bennett et al., 2016; Brydevall et al., 2018). Furthermore, people even paid a significant monetary cost for that information, while exhibiting sensitivity to the cost of that information. When information was free, people were willing to obtain it on average 90% of the time. As the cost of the information increased, the proportion of times participants paid for information decreased monotonically in the opposite direction to the cost, with the lowest point at 20% (Bennett et al., 2016). In a recent study, people were even willing to endure the risk of shocks to discover the secret behind a magic trick they had observed (Lau et al., 2018).

As with the pigeons above, this willingness to pay a cost for obtaining non-instrumental information might at first glance seem sub-optimal. These choices, however, were observed in situations deliberately constructed for information to be non-instrumental and perhaps reflect psychological mechanisms that would otherwise be adaptive. In different contexts, such a strong bias towards information gathering and exploration can actually lead to optimal behaviours, either at the individual level or at a societal level (Hills et al., 2015; Wilson et al., 2014a; Vasconcelos et al., 2015). In the experimental context with humans, perhaps participants were bored or disinterested and gained utility from simple engagement with the task, even if the only engagement resulted in an explicit cost or penalty (see Wilson et al., 2014b).

Here, similar to Bennett et al. (2016), we designed an experiment to further assess when and how people are willing to pay for information that does not have direct instrumental value—but can merely satisfy their curiosity a little earlier. In a virtual slot machine task (see Fig. 1), participants had to choose between revealing in advance the identity of up to three different cards, or waiting five seconds until the cards were automatically revealed. On most trials, revealing the cards before the end of the trial had a cost, and waiting was free. The reward on each round depended on the combination of cards obtained (see Table 1), and any action taken by the participants did not influence the outcome; therefore, all advance information was strictly non-instrumental. In

Table 1

Task specification including card combination, points rewarded, and odds of occurrence.

Cards	Points	Odds
H-H-H	500	1/27
M-M-M	300	1/27
L-L-L	100	1/27
Three Different	50	6/27
Any other combination	0	18/27

Note. H=High; M = Medium; L = Low.

contrast to the work by Bennett et al. (2016), in the current experiment, there were multiple possible reward levels and different probabilities associated with different possible outcomes. This enhancement allowed us to assess the impacts of reward magnitude and probability on information seeking. In the experiment, the first card was always free to reveal, so we could assess the impact of the reward indicated by that first card on the likelihood of revealing the second or third cards.

Based on the significant sub-optimal choice displayed by non-humans when faced with delayed probabilistic rewards, we expected that participants would exhibit information-seeking behaviour (i.e., reveal cards in advance). Second, we expected participants to be sensitive to the cost of the information: the higher the cost, the less often participants would choose to reveal a card in advance (Bennett et al., 2016). Third, if the magnitude of the reward impacts information-seeking, as would be the case if the information were being used to maximize rewards, then participants would be more likely to reveal a card in advance when a higher potential reward is available (see also Golman and Loewenstein, 2018). In contrast, recent research on curiosity in humans indicates that reward magnitude does not impact the degree of curiosity exhibited (van Lieshout et al., 2018). Applied here, only uncertainty about outcome would matter, and the potential reward magnitude on a trial should not impact the likelihood of revealing the cards early.

2. Method

2.1. Participants

75 participants (Mean age = 24; SD = 3.6; 3 undisclosed) were recruited through the University of Warwick (Coventry, UK) SONA system for participant recruitment. Participants consisted of 22 males, 49 females, 2 neither gender, and 2 undisclosed gender. Written informed consent was given by all participants prior the study. Ethical approval was provided through the psychology department at the University of Warwick. Participants were paid a show-up fee of £2 plus

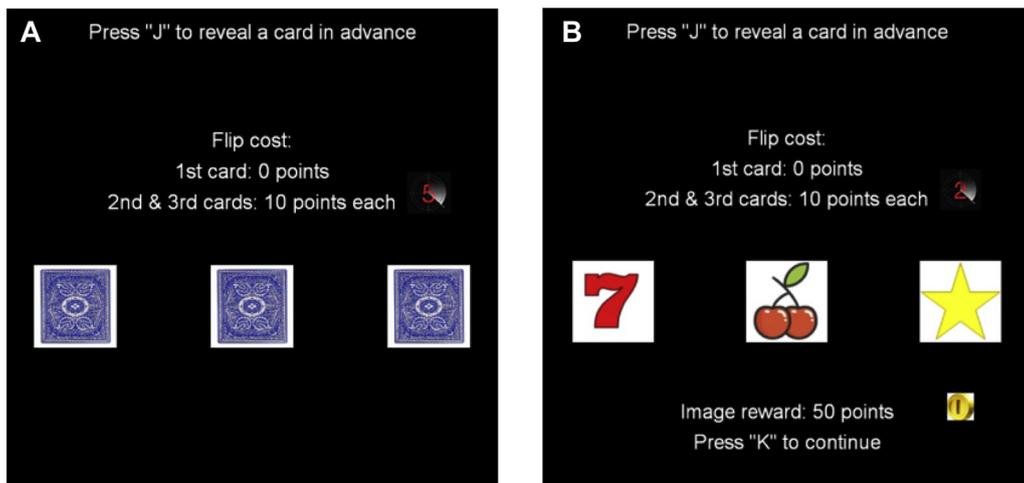


Fig. 1. (A) Screenshot of the stimuli at the beginning of each trial and (B) after the stimuli have been revealed on a trial with three different stimuli.

a variable payment of £3 to £6 based on choices. The average total pay including the show-up fee was £7.73 ($SD = £0.45$).

2.2. Materials and Stimuli

The experiment was implemented in MATLAB 2017b (The MathWorks, Inc.; Natick, Massachusetts, US), using the Psychophysics Toolbox extensions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). The experiment took place on a Windows desktop computer with a 32" monitor with resolution of 1920 × 1080 pixels. Annotated raw data, summary spreadsheets, and the full code for the experiment are available on the Open Science Framework: osf.io/kh8ps.

Fig. 1 shows screenshots of the task before and after the stimuli and rewards have been revealed. Stimuli were three different square cards with the back coloured in blue and white in an irregular pattern. The front of each card was one of the following images inspired by real-world slot machines: a yellow star, the number seven coloured in red, and a cherry with a leaf painted in red and green. The cards were aligned horizontally from left to right with a symmetrical space between them. A countdown timer at the top right corner of the screen displayed the remaining time in each round in seconds. Yellow coloured coins of one-third of the scale of the cards, with symmetrical distance between them, were displayed below the cards when a reward was received. The number of coins depended on the size of the reward: zero points showed no coins, 50 points showed one coin, 100 points showed two coins, 300 points showed three coins, and 500 points showed four coins.

2.3. Design and procedure

The experiment was conducted in a computer laboratory at the University of Warwick during weekdays from 10:00 am to 4:30 pm in 30-minute sessions. Sessions had 3–9 participants, each seated at their own computer.

The main task consisted of 221 trials, including the first five which were practice trials and excluded from the data analysis, leaving 216 trials in the main analysis. On each trial, participants played a virtual card game where they were rewarded based on the cards on that trial. They had the option of revealing the cards earlier (and sometimes paying a cost to do so) or waiting 5 s to see the outcomes. Revealing the cards early had no effect on either the rewards to be displayed or the timing of the task, which was made explicit to participants.

Basic instructions for the task were displayed for each participant on the computer screen prior to the beginning of the experiment. As shown in Fig. 1A, at the start of each trial the backs of the three cards were shown in addition to a five-second countdown timer on top of the third card, which displayed the number of seconds until the trial ended. At the end of the trial, the identities of all the cards were revealed at the same time, automatically and at no cost. Participants had the option to sequentially reveal in advance (i.e., before the end of the trial) as many cards as they wished: 0, 1, 2, or 3 (all) by pressing a key on the keyboard. Each keypress revealed another card. When participants chose to reveal cards in advance, they were always revealed in the same order: left, centre, and then right. Otherwise, all cards were revealed at the same time at the end of the 5-s delay.

For each participant, the three card images (the number seven, a cherry, or a star; see Fig. 1B) were randomly assigned to one of the following identities: H (High), M (Medium), or L (Low). Table 1 details how the card combinations mapped onto points rewarded. Getting three identical cards yielded the biggest rewards: three H cards = 500 points, three M cards = 300 points, and three L cards = 100 points. Any combination of three different cards yielded 50 points. Any other combination provided no reward.

Four different cost levels were assigned with equal probability to each trial: 0, 5, 10, and 15 points. The cost level was communicated to the participants before the beginning of each trial and was also stated at

the top of the cards during each trial (see Fig. 1). The first card could always be revealed at no cost. The second and the third cards could be revealed by paying the cost level assigned to that particular trial.

Excluding the practice trials, participants saw each condition exactly twice (216 trials/ [27 card identity permutations x 4 cost levels]). On each trial, the probability of obtaining three identical cards was 1/27 (see Table 1). Therefore, the total probability of getting any three identical cards and earning either 500, 300, or 100 points was 3/27. The probability of obtaining three cards with different identities was 6/27, which provided a reward of 50 points. Consequently, the probability of getting any other combination of cards and not getting a reward on that trial was 18/27 or 2/3.

Participant payments depended on the total number of points they held at the end of the experiment. Points were converted to GBP at a fixed exchange rate (1630 pts = £1).

2.4. Data analysis

The primary dependent measure was the proportion of time that people revealed each of the cards. All descriptive statistics are provided with 95% confidence intervals across participants. The primary analyses were repeated-measures ANOVAs where 0.05 was taken as the nominal threshold for statistical significance. Inferential statistics were calculated using SPSS 24 for the Mac (IBM Corp., Armonk, NY). All reported statistics were double-checked for consistency with statcheck.io.

3. Results

Fig. 2 displays the proportion of times the card was revealed as a function of card number (first, second, or third) and cost (0–15). When the information was free (cost = 0), participants chose to find out early about the ensuing outcome around 80% of the time. They revealed the first card $81.3 \pm 6.3\%$, the second card $77.0 \pm 6.4\%$, and the third card $72.6 \pm 6.8\%$ of the time. Moreover, on trials with a cost, participants were still willing to reveal the cards between 9–25% of the time, pending the condition (see Fig. 2). In total, participants sacrificed on average $7.4 \pm 2.6\%$ of their total earnings from the experiment to obtain advance information about the outcomes, even though this information could not improve their rewards.

Despite these seemingly irrational selections, people were still cost sensitive, choosing to reveal cards less often when the cost for revealing cards was higher, even for the first card when the cost was always zero.

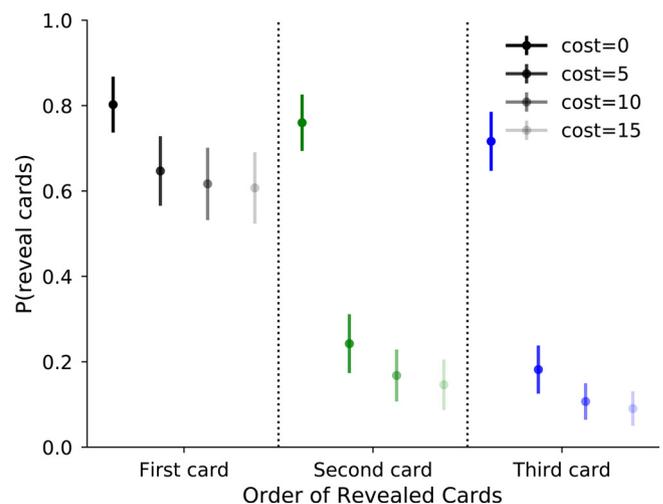


Fig. 2. Probability of revealing a card as a function of the order of the card and the cost of revealing a card. Opacity represents the cost, increasing from left to right in each sub-panel. Error bars represent 95% CIs across participants.

This pattern was corroborated by a two-way repeated-measures ANOVA (Card x Cost) where there was a main effect of both Card ($F(2,148) = 154.2, p < .001; \eta_p^2 = .68$) and Cost ($F(3,222) = 211.5, p < .001; \eta_p^2 = .74$) as well as a significant interaction ($F(6,444) = 86.0, p < .001; \eta_p^2 = .54$). Pairwise comparisons confirmed reliable differences (all p s $< .007$) between all neighbouring pairs except for costs 10 and 15 for the first card.

Because of the sequential nature of the choices, some participants never had the opportunity to reveal the second or third cards for a given cost level (i.e., because they didn't reveal the earlier cards). A two-way repeated-measures ANOVA with zero flips treated as missing values for unobserved instances (with participants thus removed) yielded similarly robust results with a main effect of both Card ($F(2,88) = 135.6, p < .001; \eta_p^2 = .76$) and Cost ($F(3,132) = 127.3, p < .001; \eta_p^2 = .74$) as well as a significant interaction ($F(6,264) = 66.6, p < .001; \eta_p^2 = .60$). Pairwise comparisons still confirmed reliable differences (all p s $< .006$) between all neighbouring pairs except for costs 10 and 15 for the first card.

As noted, the reveal probabilities for the second and third cards are bounded by the probabilities of the earlier cards, violating the ANOVA assumption of independence. Thus, as a further robustness check, to confirm the effect of cost sensitivity observed in Fig. 1, we also separately analysed the conditional probabilities of revealing a card given that the previous card was revealed in three one-way repeated-measures ANOVA. There was a significant main effect of Cost for each card (First: $F(3, 222) = 37.0, p < .001; \eta_p^2 = .33$; Second: $F(3, 213) = 217.8, p < .001; \eta_p^2 = .75$; Third: $F(3, 132) = 139.7, p < .001; \eta_p^2 = .76$). Note how each subsequent card includes fewer participants (and thus degrees of freedom) as the conditional probability is only defined for cards where there was an opportunity to reveal the card, and some participants never revealed either the first or second card for some cost levels.

Fig. 3 shows how, for the second card, there was no effect of potential reward magnitude, as indicated by the identity of the first card, on the probability of revealing the second card. The cost, however, made a big difference (as in Fig. 2) whereby higher cost options were revealed less frequently for all reward magnitudes. There was a mild interaction between these two variables, reflecting a larger difference between the two highest costs (10 and 15) in the high-magnitude option. These results were again corroborated with a two-way repeated-measures ANOVA (Magnitude x Cost). There was no reliable effect of Magnitude ($F(2,148) = .33, p = .72; \eta_p^2 = .004$), a main effect of Cost ($F(3,222) = 194.0, p < .001; \eta_p^2 = .724$), and a marginally significant

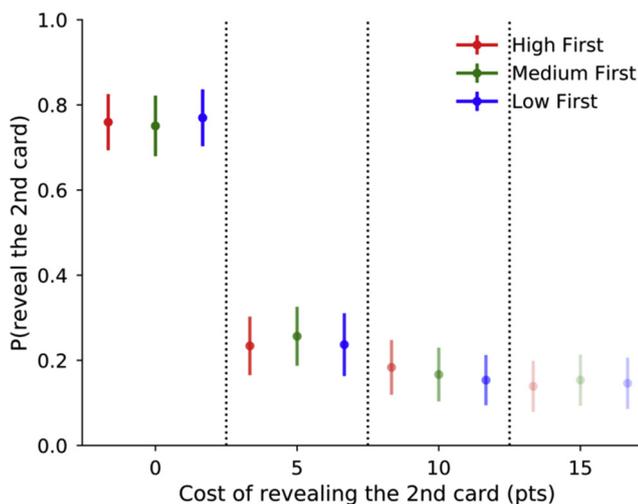


Fig. 3. Probability of revealing the second card given the cost and the potential reward magnitude, as indicated by the identity of the first card. Error bars represent 95% CIs across participants.

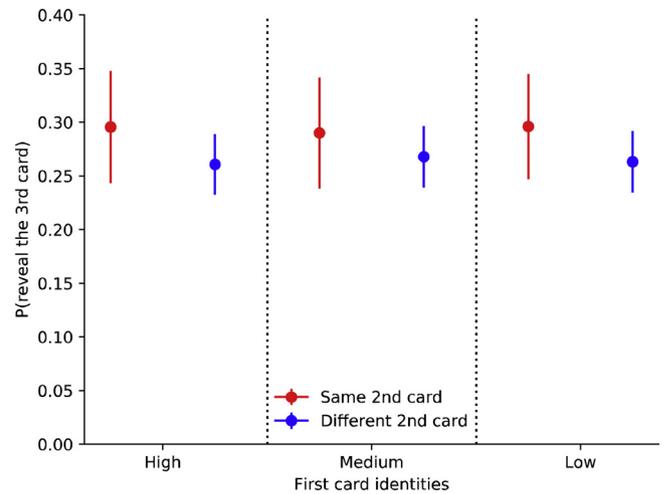


Fig. 4. Probability of revealing the third card as a function of the identities of the first two cards.

interaction ($F(6,444) = 2.18, p = .044; \eta_p^2 = .029$).

Fig. 4 shows how when the identities of the first and second cards were the same (i.e., H-H, M-M, and L-L) the probabilities of revealing the third card were significantly higher than when the identities of the cards were different (e.g., H-L, H-M). On average, participants revealed the third card $29.8 \pm 5.0\%$ of the time when the first two cards were the same, indicating a potential big win, but only $26.7 \pm 3.8\%$ of the time when the first two cards were different, indicating only a potential small (50-pt) win. The magnitude of the first card, however, made little difference, meaning that observing two stars in a row, two cherries in a row, or two sevens in a row yielded the same probability of revealing the third card. This pattern was corroborated with a 2×3 ANOVA (Card Match x Magnitude). There was a main effect of Card Match ($F(1,74) = 6.65, p = .012; \eta_p^2 = .082$), but no effect of Magnitude ($F(2,148) = .03, p = .97; \eta_p^2 = .000$) and no interaction ($F(2,148) = .60, p = .55; \eta_p^2 = .008$).

A follow-up analysis which only examined trials that followed learning (defined as at least 4 exposures to each cue-reward contingency) yielded qualitatively similar results for all three sets of analyses.

4. Discussion

In this study, people engaged in significant, sub-optimal, information seeking. They were willing to pay a cost for resolving a gamble early even though doing so did not influence the eventual outcome. They let their curiosity get the better of them—giving up nearly 8% of their total earnings in the experiment for this early information. These findings support the hypothesis that people like signals about potential good news, paralleling the sub-optimal choice exhibited by other animals (e.g., McDevitt et al., 2016; Blanchard et al., 2015; Zentall and Stagner, 2011). People may perhaps even obtain subjective utility from savouring the advance information of the future rewarding event (Iigaya et al., 2016; Loewenstein, 1987). Notably, the magnitude of the upcoming reward had little impact on the likelihood of selecting early information (see Figs. 3 and 4), in agreement with recent work on the triggers for curiosity (van Lieshout et al., 2018). This lack of sensitivity to magnitude suggests that some property of the information itself, rather than the outcomes it foretells, drives the sub-optimal information seeking.

In the experiment, when only the identity of the first card was known, participants were clearly not sensitive to potential reward magnitude (see Fig. 3). This result runs counter to one recent theory of information preferences, which predicts that the information-seeking is modulated by the importance of that information, as determined in part

by the reward magnitude (Golman and Loewenstein, 2018). One possible explanation for this insensitivity is that there was too large of an information gap between what people knew (after one card) and the final outcome (after three cards) to trigger people's curiosity (Golman and Loewenstein, 2018; Loewenstein, 1994). After one card, only 1/3 of the uncertainty was resolved, and there were still nine different possible outcomes, which even a further card would still not fully reduce to a certain outcome. On the other hand, participants exhibited some sensitivity to the magnitude of potential rewards when the identities of two cards were known (see Fig. 4). At this point, 2/3 of the uncertainty was resolved, and there were only three possible different outcomes, which a further card would completely resolve. The smaller information gap (as compared to when only the identity of one card is known) could have more strongly triggered participant's curiosity.

When the identity of two cards was known, participants were partially sensitive to the magnitude of the future potential rewards. They were more likely to reveal cards when the potential reward was higher (500, 300, or 100 points) than the lowest potential reward (50 points). The probability of revealing the third card, however, was similar among the three higher rewards (see Fig. 4). This partial sensitivity to the magnitude of rewards could potentially be due to insufficient learning of the differences between the high-value rewards. Across the task, getting three different cards was experienced 48 times, whereas the result of getting three cards with the same image was experienced only eight times for each triplet of outcomes. Therefore, the lowest reward combination was experienced six times more often than each combination of the larger rewards.

Another possibility is that perhaps participants were indeed indifferent to reward magnitude even for that third card, but the slight preference for information after two identical images emerged due to a desire for pattern completion. People may have preferred the information early when it had a chance to complete the three-in-a-row pattern than when the first two images were different from each other. In other situations, people do exhibit preference for patterns and symmetry, as in getting the same image three times in a row, as well as disgust for asymmetrical stimuli, such as obtaining three different pictures in a row (Enquist and Arak, 1994; Evans et al., 2012; Plonsky et al., 2015). Participants may have been willing to pay for the information as the pattern itself (not the potential reward) provided them with subjective utility. In the current experiment, value and pattern were confounded, but future work could reverse the ordering and assess whether people were indeed preferentially aiming for pattern completion.

Despite the seemingly irrational payments for early information, people were still cost sensitive. They were significantly less willing to reveal information as the cost of the information increased (as in Bennett et al., 2016). Interestingly, people were even less willing to reveal the first, free card when subsequent cards had a cost. This shift could potentially be due to some confusion about the cost of that first card. Alternatively, one card alone is not very informative (it rules out two of the high-reward options, but that's all), so participants may have been assessing the cost of revealing the whole sequence. They may even have been anticipating their future curiosity and limiting the cost of the sub-optimal choices their future selves would make—like a form of pre-commitment (Ariely and Wertenbroch, 2002; Rachlin and Green, 1972).

Curiosity is often an adaptive trait, prompting the gathering of information that may help improve future choices. Curiosity also amplifies learning and memory—leading to better memory for trivia questions (Kang et al., 2009) and even incidental faces (Gruber et al., 2014). One defining feature of curiosity, however, is that it does not strictly depend on whether that knowledge can be applied or become useful in the future, but rather seems to be an intrinsic drive in itself (e.g., Kidd and Hayden, 2015). In our study, indeed, there was no possibility that the information gleaned would improve the reward harvest, yet people still were willing to incur a cost for that information.

Several flavours of curiosity may have been involved in the information-seeking behaviour in this task, including diverse perceptual curiosity—the desire to evade tedium—and epistemic curiosity, which is the wish to obtain knowledge for its own sake (e.g., Berlyne, 1954). The high degree of uncertainty in the task likely amplified the willingness to pay for information, as uncertainty correlates with curiosity (Kang et al., 2009), even if reward magnitude often does not (van Lieshout et al., 2018; Fig. 2).

In the brain, curiosity elicits activity in midbrain regions overlapping with those involved in reward processing, such as the striatum (e.g., Gruber et al., 2014; Kang et al., 2009; Lau et al., 2018; see Gottlieb et al., 2013, and Kidd and Hayden, 2015, for reviews). Similarly, when monkeys can acquire information to resolve an uncertain gamble early, dopamine neurons, often thought to encode prediction errors (e.g., Schultz et al., 1997) are influenced both by the reward value and the informational value of upcoming stimuli (Bromberg-Martin and Hikosaka, 2009). This twin sensitivity to both rewards and information is paralleled during a similar task in the feedback-related negativity in humans (Brydevall et al., 2018). This direct overlap between reward and information, however, seems restricted to information about future positive outcomes, but not negative ones (e.g., Charpentier et al., 2018), suggesting a plausible mechanism as to why information about aversive outcomes may instead be avoided (Golman et al., 2017; Karlsson et al., 2009; Zhu et al., 2017). Overall, these results suggest that the brain processes extrinsic rewards and the value of non-instrumental information in the same set of overlapping neural structures—highlighting how curiosity and information-seeking may be best understood as reward-driven processes.

One limitation of the current work is that all the contingencies in the task had to be learned from experience. Though that makes the study more comparable to related animal work (e.g., McDevitt et al., 2016; Zentall, 2016), it raises the possibility that people did not fully understand the cue-outcome contingencies or the lack of an action-outcome contingency. In other areas of decision making under uncertainty, there can be wide differences in preference when people learn from their experience as opposed to making decisions based on explicit descriptions (Hertwig and Erev, 2009; Ludvig and Spetch, 2011). Future work could inform and train people on the underlying reward structure, rather than expecting them to learn them solely by experience.

Early information is often valuable. Knowing in advance whether or not you will get a raise this year would be useful to plan your finances and to take all sort economic decisions, such as buying a new car, getting a different mortgage, or tightening your spending. And human and other animals, quite sensibly, will seek out information that better their future outcomes (Hills et al., 2015; Wilson et al., 2014a). The current results, however, highlight the darker side of such information seeking. Humans are curious to the point that they are willing to pay a cost for useless information (see also Powdthavee and Riyanto, 2015), paralleling other animals who forego food and water just to find out a few seconds earlier whether a reward will be arriving (e.g., Blanchard et al., 2015; Spetch et al., 1990; Stagner and Zentall, 2010). This broad comparative convergence suggests that the roots of this sub-optimal information-seeking lies in the shared fundamental mechanisms that drive reward-based choice.

Acknowledgments

The data in this paper were included in the first author's M.Sc. dissertation at the University of Warwick and were presented at the 2018 Annual Meeting of the Comparative Cognition Society in Melbourne, FL, USA.

References

Ariely, D., Wertenbroch, K., 2002. Procrastination, deadlines, and performance: self-

- control by precommitment. *Psychol. Sci.* 13, 219–224.
- Bennett, D., Bode, S., Brydevall, M., Warren, H., Murawski, C., 2016. Intrinsic valuation of information in decision making under uncertainty. *PLoS Comput. Biol.* 12, e1005020.
- Berlyne, D.E., 1954. A theory of human curiosity. *Br. J. Psychol.* 45, 180–191.
- Blanchard, T.C., Hayden, B.Y., Bromberg-Martin, E.S., 2015. Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 85, 602–614.
- Brainard, D.H., 1997. The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- Bromberg-Martin, E.S., Hikosaka, O., 2009. Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63, 119–126.
- Brydevall, M., Bennett, D., Murawski, C., Bode, S., 2018. The neural encoding of information prediction errors during noninstrumental information seeking. *Sci. Rep.* 2018.
- Charnov, E.L., 1976. Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* 9, 129–136.
- Charpentier, C.J., Bromberg-Martin, E.S., Sharot, T., 2018. Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci.* 115 (31), E7255–E7264.
- Dunn, R., Spetch, M.L., 1990. Choice with uncertain outcomes: conditioned reinforcement effects. *J. Exp. Anal. Behav.* 53, 201–218.
- Enquist, M., Arak, A., 1994. Symmetry, beauty and evolution. *Nature* 372, 169–172.
- Evans, D.W., Orr, P.T., Lazar, S.M., Breton, D., Gerard, J., Ledbetter, D.H., Batchelder, H., 2012. Human preferences for symmetry: subjective experience, cognitive conflict and cortical brain activity. *PLoS One* 7, e38966.
- Golman, R., Loewenstein, G., 2018. Information gaps: a theory of references regarding the presence and absence of information. *Decision* 5, 143–164.
- Golman, R., Hagmann, D., Loewenstein, G., 2017. Information avoidance. *J. Econ. Lit.* 55, 96–135.
- Gottlieb, J., Oudeyer, P., Lopes, M., Baranes, A., 2013. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends Cogn. Sci.* 17 (11), 583–593.
- Gruber, J.M., Gelman, B.D., Ranganath, C., 2014. States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 84, 486–496.
- Hertwig, R., Erev, I., 2009. The description–experience gap in risky choice. *Trends Cogn. Sci.* 13, 517–523.
- Hills, T.T., Todd, P.M., Lazer, D., Redish, A.D., Couzin, I.D., Group, C.S.R., 2015. Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* 19, 46–54.
- 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.
- Kang, M.J., Hsu, M., Krajbich, I.M., Loewenstein, G., McClure, S.M., Wang, J.T.-y., Camerer, C.F., 2009. The wick in the candle of learning: epistemic curiosity activates reward circuitry and enhances memory. *Psychol. Sci.* 20 (8), 963–973.
- Karlsson, N., Loewenstein, G., Seppi, D., 2009. The ostrich effect: selective attention to information. *J. Risk Uncertain.* 38 (2), 95–115.
- Kidd, C., Hayden, B.Y., 2015. The psychology and neuroscience of curiosity. *Neuron* 88 (3), 449–460.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., 2007. What's new in Psychtoolbox-3. *Perception* 36 (14), 1.
- Krebs, J.R., Kacelnik, A., Taylor, P., 1978. Test of optimal sampling by foraging great tits. *Nature* 275 (5675), 27–31.
- Kuzmanovic, B., Jefferson, A., Vogeley, K., 2016. The role of the neural reward circuitry in self-referential optimistic belief updates. *NeuroImage* 133, 151–162.
- Lau, J.K.L., Ozono, H., Kuratomi, K., Komiya, A., Murayama, K., 2018. Hunger for knowledge: how the irresistible lure of curiosity is generated in the brain. *bioRxiv*, 473975.
- Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S., Palminteri, S., 2017. Behavioural and neural characterization of optimistic reinforcement learning. *Nat. Hum. Behav.* 1, 0067.
- Loewenstein, G., 1987. Anticipation and the valuation of delayed consumption. *Econ. J.* 97, 666–684.
- Loewenstein, G., 1994. The psychology of curiosity: a review and reinterpretation. *Psychol. Bull.* 116, 75.
- Ludvig, E.A., Spetch, M.L., 2011. Of black swans and tossed coins: Is the description–experience gap in risky choice limited to rare events? *PLoS One* 6, e20262.
- McDevitt, M.A., Dunn, R.M., Spetch, M.L., Ludvig, E.A., 2016. When good news leads to bad choices. *J. Exp. Anal. Behav.* 105, 23–40.
- McDevitt, M.A., Pisklak, J.M., Spetch, M., Dunn, R., 2018. The influence of outcome delay on suboptimal choice. *Behav. Processes* 157, 279–285.
- Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442.
- Pisklak, J.M., McDevitt, M.A., Dunn, R.M., Spetch, M.L., 2015. When good pigeons make bad decisions: choice with probabilistic delays and outcomes. *J. Exp. Anal. Behav.* 104 (3), 241–251.
- Plonsky, O., Teodorescu, K., Erev, I., 2015. Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychol. Rev.* 122, 621–647.
- Powdthavee, N., Riyanto, Y.E., 2015. Would you pay transparently unless advice? A test of boundaries of beliefs in the folly of predictions. *Rev. Econ. Stat.* 97 (2), 257–272.
- Rachlin, H., Green, L., 1972. Commitment, choice and self-control. *J. Exp. Anal. Behav.* 17, 15–22.
- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Spetch, M.L., Belke, T.W., Barnet, R.C., Dunn, R., Pierce, W.D., 1990. Suboptimal choice in a percentage-reinforcement procedure: effects of signal condition and terminal-link length. *J. Exp. Anal. Behav.* 53 (2), 219–234.
- Spetch, M., Mondloch, M., Belke, T., Dunn, R., 1994. Determinants of pigeons' choice between certain and probabilistic outcomes. *Anim. Learn. Behav.* 22, 239–251.
- Stagner, J.P., Zentall, T.R., 2010. Suboptimal choice behavior by pigeons. *Psychon. Bull. Rev.* 17, 412–416.
- 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. *J. Neurosci.* 2816–2817.
- Vasconcelos, M., Monteiro, T., Kacelnik, A., 2015. Irrational choice and the value of information. *Sci. Rep.* 5, 13874.
- Wilson, R.C., Geana, A., White, J.M., Ludvig, E.A., Cohen, J.D., 2014a. Humans use random and directed exploration to solve the explore-exploit trade-off. *J. Exp. Psychol. Gen.* 143, 2074–2081.
- Wilson, T.D., Reinhard, D.A., Westgate, E.C., Gilbert, D.T., Ellerbeck, N., Hahn, C., Shaked, A., 2014b. Just think: the challenges of the disengaged mind. *Science* 345 (6192), 75–77.
- Zentall, T.R., 2016. When humans and other animals behave irrationally. *Comp. Cogn. Behav. Rev.* 11, 25–48.
- Zentall, T.R., Stagner, J., 2011. Maladaptive choice behaviour by pigeons: an animal analogue and possible mechanism for gambling (sub-optimal human decision-making behaviour). *Proc. R. Soc. Lond. B: Biol. Sci.* 278, 1203–1208.
- Zhu, J., Xiang, W., Ludvig, E.A., 2017. Information seeking as chasing anticipated prediction errors. *Proceedings of the Cognitive Science Society Meeting (CogSci-17)*, vol. 39, 3658–3663.