Generalizing the Effect of Extreme Outcomes in Risky Decision Making:

A Cross Species Comparison of Pigeons and Humans

by

Jeffrey M. Pisklak

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Psychology

University of Alberta

Abstract

Research has shown that when decisions between risky and fixed options are based on repeated exposure to the governing contingencies, preference is often influenced by an overweighting of the most extreme outcomes appearing in the decision context. Known formally as the extreme-outcome rule, this predicts that organisms will be more risk seeking for gains relative to losses. This thesis examines the extent to which the extreme-outcome rule can be considered to generalize both within and across two evolutionary distant species: pigeons (*Columba livia*) and humans (*Homo sapiens*). In Chapter 2, an operant chamber analogue of a foraging task was conducted with pigeons and humans to test the predictions of the extreme-outcome rule. Risk-preference results across both species corresponded to the predictions of the extreme-outcome rule; however, a detailed analysis suggested that the risk-preference results might be confounded by an effect of probabilistic discounting. Controlling for this, Chapter 3 re-examined the predictions of the extreme-outcome rule. Pigeons showed no effect of overweighting the most extreme values. Humans who passed the catch trials did show a clear effect of extreme-outcomes; though, this only represented approximately half of the sample tested, as the other half failed to meet the set passing criteria for catch trial performance. Further testing revealed that this poor performance was likely the result of poor instructional (i.e., discriminative) control.

Preface

The research conducted herein will be published with the following co-authors: Christopher R. Madan of Boston College, Elliot A. Ludvig of the University of Warwick, and Marcia L. Spetch of the University of Alberta. The pigeon operant chambers referred to in Experiment 1a and the testing arena in Experiment 2a were designed and built by Isaac Lank of the University of Alberta's Psychology Department. I was responsible for data collection and analysis. Christopher R. Madan assisted me in the development of the computer programs used to present the experimental stimuli and log the behavioural responses. Marcia L. Spetch was the supervisory author and, along with Elliot A. Ludvig and Christoper R. Madan, was also involved in the concept formation of the work itself.

Acknowledgements

I would like to thank my supervisor, Marcia Spetch, for her assistance which has been invaluable not only in the development of this thesis but also in my University career more generally. In addition, I want to acknowledge my other two fellow authors on these projects, Christopher Madan and Elliot Ludvig. Their input and direction on these projects have been immensely helpful throughout. Finally, I would like to recognize the support of the various research assistants who aided in the collection of both pigeon and human data: Ariel Greiner, Jason Long, Nuha Mahdi, and Josh Yong.

All animal research was supported by a Natural Sciences and Engineering Research Council of Canada grant belonging to Marcia Spetch. Human research was supported by a grant from the Alberta Gambling Research Institute also belonging to Marcia Spetch.

Table of Contents

List of Figures

Chapter 1: Introduction

Introduction

In common parlance the term “risky” can be used to describe a wide array of behaviours. At its core though lies the notion of possibility. For instance, if there is a possibility of severe financial loss or great financial gain, then the behaviour of making an investment might be said to be a “risky” one. If there is a possibility of the parachute not opening, then the behaviour of skydiving might also be considered risky. Of course risk need not be relegated to solely human affairs. Animals too constantly make high-stake decisions that incorporate many variables. Choosing wisely or unwisely could result in any number of desirable or undesirable outcomes, not least of which could include starvation and predation. In research, however, the stakes are rarely quite so severe.

Experiments examining risk-preference, the tendency to choose a riskier alternative out of a set of other safe or less risky alternatives, is often studied in humans using fictional amounts of money in fictional scenarios on the assumption that this will generalize to the real world. Notable in this regard has been Kahneman and Tversky’s (1979) work on Prospect Theory. In one of many demonstrations of Prospect Theory’s potential, they posed a series of questions to people ($N = 644$) asking them to choose between two alternatives. For example, a subject might be asked “would you prefer a 100% chance of receiving $50 or would you prefer a 50% chance of receiving either $100 or $0?” When confronted with a question like this, most people tended to select the certain gain of $50. However, when the same question was framed in terms of losses instead of gains — “would you prefer a 100% chance of losing $50 or would you prefer a 50% chance of losing either $100 or $0?” — most people tended to pick the variable or “risky” option. Usually referred to as the *reflection effect*, but also known as the *payoff domain effect,* this has been a relatively common finding in the human literature on risky choice (Payne, Laughhunn, & Crum, 1980; Schoemaker, 1990; Tversky &

Kahneman, 1986). More recent work, however, has suggested that this finding, as well as others supporting Prospect Theory, may not be so general as once thought. For instance, Barron and Erev (2003) found that when people are made to learn the consequences of their choices through experience, as opposed to having the outcome probabilities verbally described to them as was the case above, the *certainty effect* (an overweighting of outcomes believed to be certain relative to outcomes believed to be probable), *the payoff domain effect* (greater risk aversion for gains relative to losses), and the overweighting of rare events are all observed to be reversed. Subsequent work has corroborated these findings and created what has become known in the decision-making literature as the *description-experience gap* (e.g., Camilleri & Newell, 2009; Hertwig & Erev, 2009; Ungemach, Chater, & Stewart, 2009).

One hypothesis put forward to explain, in part, the experience-based side of this 'gap' has been that of the *extreme-outcome rule* (Ludvig, & Spetch, 2011; Ludvig, Madan, & Spetch, 2014). This states that when the outcome probabilities of a decision are learned through repeated exposure to the governing contingencies, people overweight the most extreme values encountered inside the decision context. Consequently, a disproportional influence on choice behaviour is exerted by the largest and smallest outcomes encountered. As a practical example of how this plays out, suppose a person is repeatedly confronted with choices that lead to the outcomes described in the Prospect Theory scenario mentioned above. The extreme-outcome rule predicts that people will overweight the best and worst possible outcomes when making decisions. Therefore, contrary to Prospect Theory, people should be more likely to select the risky option that might provide a gain of $100 and select the safe option that guarantees a loss of $50 — assuming the gains and losses are all found within the same context and repeated exposure to the outcomes occurs. Computer-based experiments in which participants are made to search behind contrasting doors for differing amounts of gains and losses of points exchangeable for money have corroborated this prediction (Ludvig, & Spetch, 2011; Madan, Ludvig, & Spetch, 2014). Moreover, when the decision context is restricted to all gains or all losses, more risk seeking can still be observed for choices involving relatively better outcomes, than for choices involving relatively worse outcomes (Ludvig, Madan, & Spetch, 2014).

Importantly, the application of this rule may not be specific to humans. A cross-species comparison by Ludvig, Madan, Pisklak, and Spetch (2014) gave pigeons a foraging analogue of a procedure typically used to examine human decision making. In this study pigeons had to walk behind four distinct barriers to obtain food rewards. Two barriers contained either high or low value risky gain outcomes: 1 or 3 cups of food with a 50% chance of either occurring and 0 or 2 cups of food also with a 50% chance of occurring. The other two barriers contained either high or low safe gain outcomes: 2 cups of food with a 100% chance of occurring and 1 cup of food also with a 100% of occurring. The pigeons' choice behaviour was compared to that of humans who participated in a computer-based task where differing doors had to be selected for similarly varied amounts of points exchangeable for money. The results of both species conformed nicely to the predictions of the extreme-outcome rule, showing greater risk-seeking for high-value gains relative to low-value gains. In terms of the pigeons behaviour, this implies their observed distribution of choices really reflected, on the one hand, a drive to obtain the 3 cups of food presented behind the risky high-value option and, on the other hand, an avoidance of the 0 cups presented by the risky low-value option.

While these findings offer some promise of a correspondence between the human and animal literatures on risk preference, caution is nonetheless warranted because a broader scope of the animal risk-preference literature reveals large inconsistencies. As noted in Lagorio and Hackenberg (2012), when risk is assessed in terms of reinforcer amount (as opposed to reinforcer delay) the most common finding in the animal literature favours risk-aversion; that is to say, a tendency to prefer fixed (safe) amounts of food to variable (risky) amounts of food (e.g., Clements, 1990). Though, findings of risk-seeking (e.g., Barnard, Brown, Houston, & McNamara, 1985) and risk-neutrality (e.g. Mazur, 1989) are not wholly uncommon either. Lagorio and Hackenberg note that, even when the literature review is confined to a single species, pigeons, operating under similar laboratory conditions, there tends to be little in the way of consistent behavioural effects across experiments. Further complicating matters are experiments indicating that the distribution of choices between risky and safe outcomes may be contingent on factors such as energy budget (e.g., Stephens, 1981) or the effects of the most recent reinforcing and punishing outcomes that are incurred as a result of the choices made (e.g., Marshall & Kirkpatrick, 2013).

Given the immense diversity of results that have been observed in the animal literature on risky choice, it was necessary that some effort be made to establish the reliability of Ludvig, Madan, Pisklak, and Spetch's (2014) findings. The series of experiments that follow were specifically designed to assess both the replicability and generalizability of this work, seeing to what extent the predictions of the extreme-outcome rule hold, not only within species, but across species as well.

Chapter 2: Generalizing the Extreme-outcome Rule

Generalizing the extreme-outcome rule

In order to test the generalizability of Ludvig, Madan, Pisklak, and Spetch's (2014) findings, a novel pigeon and human operant conditioning task was employed. The task was modelled on the open-field pigeon foraging procedure used in Ludvig, Madan, Pisklak, and Spetch (2014), and consisted of a quasi-token economy. Subjects selected between pairs of four possible circles of varying colours. Each colour corresponded to either a risky or safe outcome that contained either a high or low amount of rewards. Once a coloured circle was chosen, a particular number of black circles (tokens) would appear on screen. Each black circle could be selected for a fixed amount of reward. For pigeons, a response to each black circle resulted in 1 second of food access via an automated hopper. For humans, a response to each black circle resulted in a 1 point increase on a point tally. Points earned by the humans could be exchanged at the end of the experiment for money. If the extreme-outcome rule is correct, then significantly greater risk-seeking should be observed when a choice between a risky and fixed high-value outcome (of equal expected value) has to be made, than when a choice between a risky and fixed low-value outcome (also of equal expected value) has to be made.

Experiment 1a: Pigeons

Subjects

As per Ludvig et al. (2014), six adult pigeons (*Columba livia*) — four Racing Pigeons and two Silver King – randomly selected from a University of Alberta pigeon colony were used. All the birds had extensive learning histories but none had ever been used in any studies of risky choice. They were individually housed inside a temperature controlled colony room with a 12 hour light-dark cycle. Free access to both grit and vitamin enriched water was provided, and each bird was maintained at approximately 85% of its free-feeding weight by means of post-session rations of Mazuri Gamebird

(PMI Nutrition International) food pellets. No birds sampled were excluded from analysis. All procedures were approved by the University of Alberta Biological Sciences Animal Care and Use Committee who follow the guidelines set forth by the Canadian Council on Animal Care.

Apparatus

A custom built 71.12 × 33.02 × 44.45 cm sound attenuating operant conditioning chamber, located in an isolated room, was used. Mounted centrally against the chamber's widest wall was a 17 in. ViewSonic LCD (1280 × 1024 screen resolution) monitor equipped with a Carrol Touch infrared touchscreen (Elo Touch Systems, Inc., Menlo Park, CA). Two feeding ports, each adjacent to the left and right side of the monitor, provided access to food pellets via a solenoid controlled food hopper containing Mazuri food pellets. Each feeding port was also equipped with a light to signal that the hopper had been raised, as well as an infrared beam to precisely detect entry into the feeding port.

Stimuli

Experimental stimuli consisted of coloured circles presented against a grey background on the chamber's internal monitor. Six different colours were used: white, black, purple, green, yellow, and orange. Each circle had a radius of 50 pixels (approximately 2.5cm).

Procedure

Training. Training sessions lasted for 45 min and were run six days a week at approximately the same time each day. The pigeons were first exposed to an autoshaping paradigm that presented only white and black circle stimuli. White circles were presented either centrally on the screen or to the right or left of center. Black circles could appear in any one of numerous spatial locations (see Figure 1-1). Each autoshaping trial began with a single white or black circle that remained onscreen for 60 seconds or until a single peck was made to it, at which point the stimulus would disappear and a hopper would raise allowing the bird 1-second of food access. Once the 1-second had elapsed, a 20-second intertrial interval would come into effect and a new trial would begin. This autoshaping procedure lasted until the birds responded to at least 75% of the stimuli over the course of 30 trials.

Once the autoshaping criteria had been met, a basic operant conditioning procedure was implemented in which the intertrial interval was shortened to 2 seconds and each circle had to be pecked once to raise a hopper. When the birds demonstrated that they could complete 50 trials within the session's 45 minute duration, they were given daily sessions consisting of only 16 trials. On these trials, the stimuli appeared in a fashion that mimicked what they would receive in testing. First, a white circle (the start-stimulus) would appear centrally on the screen. A single peck to this would then produce another white circle to either the left or right of center while simultaneously erasing the former. Pecking this new circle erased it and resulted in 1, 2, 3, or 4 black circles appearing on screen according to the layouts depicted in Figure 1-1. Each dot could be pecked for 1-second access of food. Once all the dots had been selected a 2-second intertrial interval ensued. Sessions were continued in this manner until the birds' weight stabilized at 85%.

Following stabilization of their weight, a pre-testing phase began. This phase exposed the pigeons to four sessions, each consisting of 16 single-option trials that now employed the other four coloured circles: yellow, green, orange, and purple. During these single-option trials, the birds were exposed to the various stimuli and contingencies they would receive in testing, but without forcing them to make any choices. The complete set of stimuli and reward contingencies (see Testing section below and Figure 1-1) were presented over the course of each session's 16 trials, with the order and side presentation randomized.

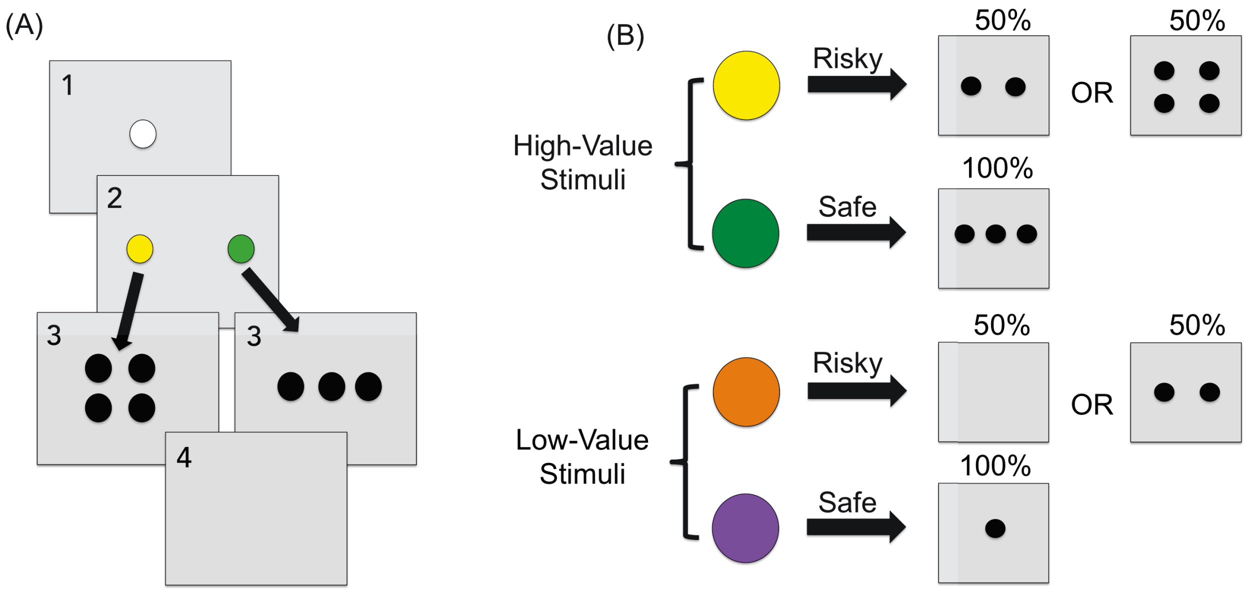

Figure 1-1. (A) A schematic of a choice trial displaying a risky high-value choice and a safe high-value choice, with possible outcomes. (B) A schematic illustrating the various reward contingencies in effect. The colour of the choice circles relative to their outcome was counter balanced across participants according to four different combinations of yellow, green, orange, and purple.

Testing. Testing lasted for 80 sessions of 45 min, with each session consisting of 16 trials. If a bird did not complete the 16 trials, the session was extended to 60 min at which point it was terminated regardless of whether the 16 trials had been completed. Sessions were run daily, six days a week, and at the same time each day. A schematic of an example trial can be seen in Figure 1-1. Each trial began with the white start-stimulus appearing centrally on the screen. A single peck to this stimulus caused it to disappear and simultaneously produce either one or two coloured (choice) circles on either side of it. A single peck to a coloured circle erased the screen and produced 0, 1, 2, 3, or 4 black circles. The number of black circles appearing depended on which coloured circle was pecked. Choice circles were classified according to the type of outcome they produced: *risky-high, risky-low, safe-high, or safe-low.* A risky-high outcome produced either 2 or 4 black circles with a 0.5 probability of obtaining either. A risky-low outcome produced either 0 or 2 black circles again with a 0.5 probability of obtaining either. Safe-high produced 3 black circles with a probability of 1, and safe-low produced one black circle also with a probability of 1. The spatial layout of the various circles can be seen in Figure 1-1. A single peck to any one of these black circles caused that particular circle to disappear and raise either the left or right side food hopper, randomly selected each time. Upon entry into the feeding port, each bird could obtain 1 second of food access. Following this, the pigeons were free to select from any remaining black circles for another 1 second of food access until all the black circles had been cleared from the screen. After a 2-second inter-trial interval had elapsed a new trial would resume.

Three types of trial were included in each session: four *risk-preference*, eight *catch* trials, and four *single-option* trials (see Table 1-1). Risk-preference trials presented a choice between a risky-high and a safe-high option, or a risky-low and a safe-low option. Catch trials presented a choice between a high and low option, altogether creating four different choice combinations: risky-high vs. risky-low, risky-high vs. safe-low, safe-high vs. safe-low, and safe-high vs. risky-low. The order of these trials was randomized within the session and counterbalanced such that each choice stimulus appeared equally often on either the left or right spatial positions. All possible combinations of stimuli and reward outcomes were counterbalanced across every four sessions.

Experiment 1b: Humans

Subjects

A total of 30 human participants (*Homo sapiens*) were recruited from the University of Alberta Psychology Subject Pool, each participating for course credit. Informed consent was obtained from all participants and all procedures were approved by the University of Alberta's Arts, Science, and Law Research Ethics Board. The age ($M \pm SD$) range was 20.5 ± 2.1 years, with 20 females. No other demographic information was collected. Seven participants were excluded from the analysis due to failure to learn the contingencies (see Results section for further details).

Apparatus

The experiment was conducted inside 15 small rooms that surrounded a larger common room. Each small room contained a single table, chair, computer, and booklet of seven mazes.

Procedure

Participants were first directed to a row of tables with an informed consent document, which they were asked to read and fill out. Upon completing the document, the following instructions were read aloud to them and projected onto a large lecture screen. "*The experiment consists of a computer portion and a maze portion. For the computer portion: Use the mouse cursor to click on the circles that you see. Try and get as many points as possible. For the maze portion: a message will appear on screen telling you when to complete a particular maze. Take a few minutes to complete the maze with a pen or pencil. All the mazes are solvable, but don't worry if you can't solve it. The experiment should take just under 2 hours to complete. You will earn 2 credits for participating, and receive a cash bonus of up to $5.00 depending upon the number of points you get in the computer portion. The more points you get the more money you will earn."* Participants were then each randomly assigned to one of the 15 rooms.

Inside the room participants were presented with a prompt on the computer screen asking them to input their age and sex. Completing the prompt then began the experiment. Visually, the computer portion of experiment was identical to what the pigeons received during testing with the exception that, at the bottom of the screen, a point tally was kept. Each click on a black circle raised this tally by 1. The experimental session was divided into 8 parts consisting of 64 trials. Because of the session's long length, each part was separated by on-screen instructions to complete one particular maze to prevent eyestrain. The first part consisted solely of 64 single-option trials. All subsequent parts of the experiment consisted of 16 risk preference trials, 16 single-option trials, and 32 catch-trials, which accounted for every possible choice, side, and reward, combination of the stimuli. Upon completion of the experiment, each participant's respective point values were used to calculate a cash bonus of up to \$5.00. Receiving the number of points expected by chance responding resulted in a \$2.50 bonus, whereas responding that produced either the minimum or maximum number of possible points resulted in a \$0.00 or \$5.00 cash bonus respectively. All intermediate point values resulted in a cash bonus that was determined linearly on the basis of the above three values, with the cash amount received always rounded up to the nearest \$0.25. Participants were not made aware of the point to cash scaling anytime prior to or during the experiment.

Results

To equate obtained experience of the contingencies across species, statistical analysis was conducted on testing trials 256-448 — this corresponds to the last 3 blocks of 64 trials completed by the humans and sessions 16-28 completed by the pigeons. Data from seven human participants were removed due to failure to learn the contingencies. Failure to learn is defined here as choosing the high value option, averaged across all catch trial types, less than 60% of the time.

Figure 1-2 shows the mean proportion of risk preference for high-value and low-value choice types, as well as the 95% confidence interval of the *paired difference* (Franz & Loftus, 2012) between the two choice types. On average, pigeons showed significantly more risk-seeking for the high-value options ($M = .54$; 95% $\mathrm{CI}^{\mathrm{betw}}$ [0.29, 0.78]) than the low-value options ($M = .15$; 95% $\mathrm{CI}^{\mathrm{betw}}$ [0.09, 0.21]), $t(5) = 4.83$, $p = 0.005$, $d = 1.91$, 95% CI [0.60, 3.89]. Similarly, the humans showed significantly more risk-seeking for the high-value ($M = .57$; 95% CI$^{\text{betw}}$ [0.43, 0.71]) than the low-value options ($M = .29$; 95% CI$^{\text{betw}}$ [0.18, 0.40]), $t(22) = 5.45$, $p < 0.001$, $d = 0.93$, 95% CI [0.51, 1.41]. All effect sizes in Experiments 1-3 were standardized using the averaged sample standard deviation and reported using an unbiased estimate of Cohen's d (see Equations 11.9 and 11.13 in Cumming, 2011).

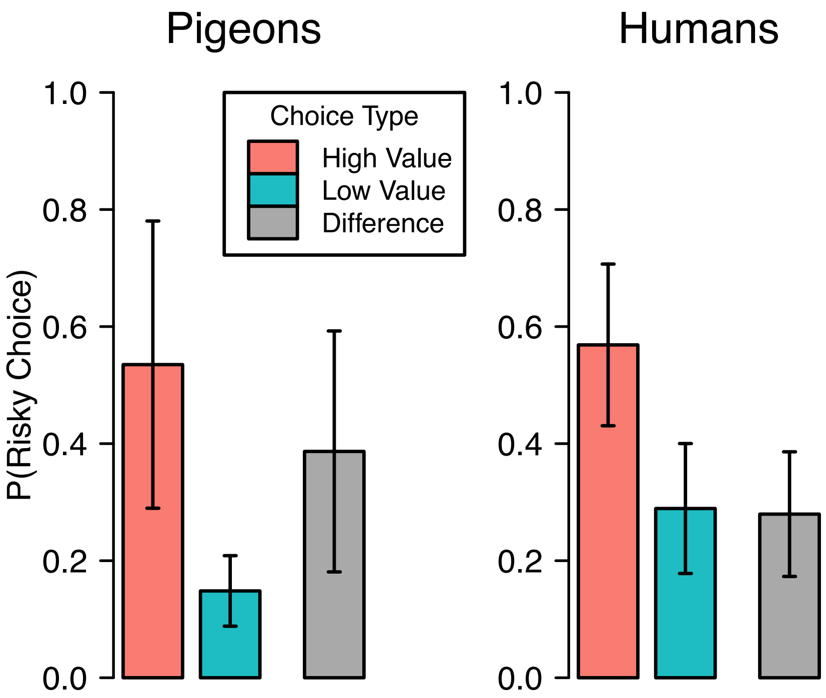

Figure 1-2. Mean proportion of choices to the risky option for pigeons (left) and humans (right) on high- and low-value risk preference trials. The mean difference between both groups is plotted in grey. Error bars on the high- and low-value choice types indicate the 95% between-subject confidence interval. The error bar on the differences indicates the 95% confidence interval of the paired difference between the high- and low-value choice types.

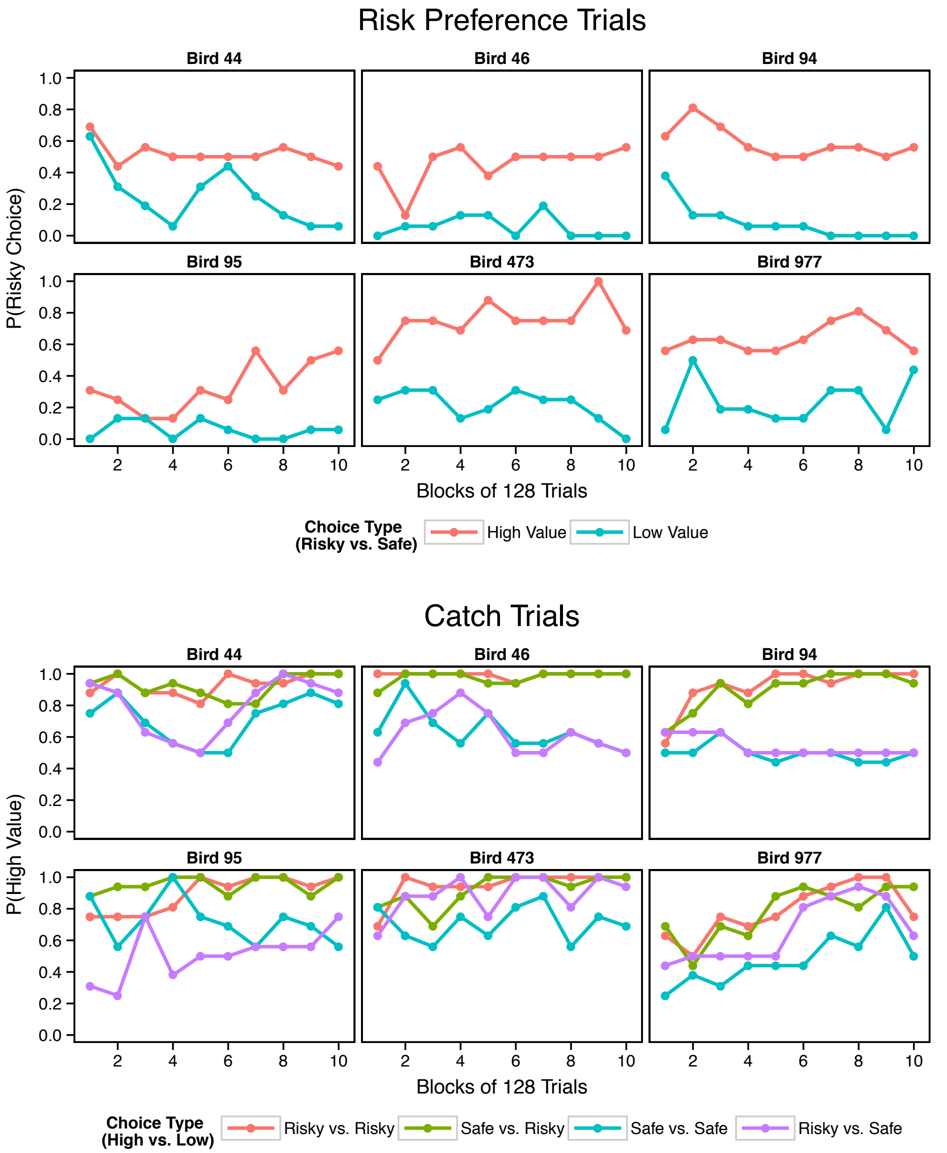

Figure 1-3. Individual learning curves for each pigeon on both risk-preference (top) and catch trials (bottom). Risk-preference graphs show proportion of risky choices made as a function of 128 trial blocks across both high-value and low-value choice types. Catch trial graphs show the proportion of high value choices made as a function of 128 trial blocks across the four types of catch trial. The catch trial choice types in the legend should be read as high-value options versus low-value options. E.g., 'Safe vs. Risky' refers to a choice between a safe high-value option and a risky low-value option.

Figure 1-3 depicts the individual learning curves (Brown & Heathcote, 2003) for pigeons on both risk-preference trials and catch trials. Both the high-value and low-value risk-preference remained relatively stable over the course of the experiment. However, catch trial performance was far more variable. Notably, this variability tended to be exclusively the result of choice types that had a safe low-value alternative. Choice types that contained a risky low-value alternative tended to be relatively stable and exhibit higher probability of selecting the larger reward outcome. To put it another way, catch-trials containing an option that could provide 0 black circles as a potential outcome were learned the most readily. Graphical inspection of the human catch-trial data (not shown) revealed no such differentiation, as ceiling effects were observed across all catch trial types.

Discussion

Using a novel operant conditioning task in pigeons and humans, we observed similar levels and patterns of risk-preference after equivalent amounts of training, replicating the core findings of Ludvig, Madan, Pisklak, and Spetch (2014). This supports the basic tenet of the extreme-outcome rule (Ludvig, Madan, & Spetch, 2014) stating that preference is disproportionately influenced by the most extreme values occurring in a context. In other words, the most extreme values are overweighed in the decision process. Consequently, the possibility of receiving 4 rewards pulls preference towards the risky high-value option, while the possibility of receiving 0 rewards pushes preference away from the risky low-value option.

However, when one considers the observation that catch trials containing a potential zero outcome were the most reliably learned by the pigeons, an interpretation in terms of probabilistic discounting seems appropriate. Many studies (e.g., Green, Myerson, & Calvert, 2010; Mazur, 1989; Rachlin, Logue, Gibbon, Frankel, 1986) have found that when a reinforcing event occurs probabilistically (e.g., food is sometimes provided after the operant response and sometimes not provided), the reinforcing value of the event is diminished in a manner analogous to that seen when a delay is imposed between a operant response and reinforcing outcome. Consider the safe low-value option and the risky low-value option of Experiment 1a. For the safe option, a reinforcing event is provided each time the operant response is made; consequently, there is little to no delay to food for the pigeon. For the risky option, since the probability of getting reinforced is .5, there is no guarantee that the pigeon will get an immediate delivery of food on any given trial. On some selections it may get food immediately, while on others it may have to wait through two or more attempts before any is forthcoming. Thus, on average the pigeon has to wait longer for food on the risky low-value option than on the safe low-value option. It is a well-established finding that the value (i.e., the effectiveness) of a reinforcer decreases as a function of delay (e.g., Chung, 1965; Mazur, 1984; McDiarmid & Rilling, 1965). Frequently referred to as delay-discounting, this decrease is often drastic enough that in many cases animals will prefer a smaller food reward presented sooner than a larger food reward presented later (for a review see Stevens & Stephens, 2010). In Experiment 1, because of its zero-value outcome, only the risky low-value alternative poses any kind of delay. By contrast, the other three alternatives always provide immediate reinforcement for each operant response. On the basis of this it seems likely that the risky low-value option's reinforcing event (i.e., the two black dots that appear 50% of the time) is being devalued by the delay imposed by zero occurring some of the time. If the reinforcing event is devalued then, by definition, the operant response its delivery is contingent upon (selecting the risky low-value circle) has to occur less frequently.

This interpretation has a strong appeal for a number of reasons. First, it accounts for why the risky high vs. risky low and the safe high vs. risky low catch trials demonstrated more selection of the larger reward option than the other two types of catch trials. The possibility of incurring the zero value meant that birds' preference could develop on the basis of both the amount of reinforcement delivered and the delaydiscounted value of one of the options. This is not the case for catch trials which had no possibility of zero. On those trials preference could develop only on the basis of reinforcement amount since delay was not a factor. Second, probabilistic discounting does not preclude the possibility that learning to choose the larger reward option can occur on catch-trials that contain no possibility of a zero value outcome. Third, and most importantly, probabilistic discounting predicts the pattern of results seen in the risk-preference trials, especially if one assumes that the equal expected values of the two high-value outcomes on risk-preference trials essentially indicate equally reinforcing values.

In Ludvig, Madan, Pisklak, and Spetch (2014), it was suggested that because there was a trend towards risk seeking for the high-value choice types by the pigeons, an interpretation in terms similar to those just described (i.e., an avoidance of zero) could be ruled out. However, such a trend was not observed in this study. For both the pigeons and the humans, choices between the risky high-value and safe high-value options hovered around indifference, whereas choices between the risky low-value and safe low-value options showed significant departures from chance.

One potential argument against the probabilistic discounting view in this case is the finding that when humans are given choices between gains and losses (e.g., +40/0 vs. +20 and -40/0 vs. -20) the predictions of the extreme outcome rule still hold, implying that discounting is not occurring even though zero is a possible outcome in both risky cases (Ludvig & Spetch, 2011, Ludvig, Madan, & Spetch 2014). Insofar as this gain and loss design is concerned, that is probably correct but it is not clear why it should be assumed that this would generalize to an all gains procedure of the sort used in Experiment 1. In Experiment 1, zero is the worst of all possible outcomes. In the gain and loss procedure, zero is among the best results that can be obtained relative to the two potential losses (-20 and -40). On the basis of this, it seems almost self–evident that the zero value outcome is not functioning the same in both procedures and hence a generalization of its function on behaviour in these two cases is probably not appropriate. A further, but nonetheless related, complication is that in the gain and loss procedures cited, a numerical amount is always presented with a corresponding image. Procedurally then, a reinforcing, punishing, or (possibly) neutral outcome of some kind is *always* being provided for each response. The same cannot be said of the current procedure, which only ever provides reinforcing outcomes, one of which has a probabilistically determined appearance. Hence some trials contain, in a more literal sense, no actual outcome. It could be that, in some circumstances, *seeing* an outcome of "0" may not be equivalent to what amounts to a trial of operant extinction.

One point worth considering is that probabilistic discounting and the extreme-outcome rule are not mutually exclusive. Both could be operating inside Experiment 1's procedure such that the difference between the high- and low-value risk-preference choice types are actually being inflated relative to what would occur if only the extreme-outcome rule or only probabilistic discounting was in effect. On the basis of the current design, it is impossible to know for certain the relative influence of either.

In summary, while these two experiments did technically uphold the predictions of the extreme-outcome rule, showing greater risk seeking for high-value rewards relative to low-value rewards, this conclusion needs to be tempered by the fact that the results of the experimental design may be explainable, in part or altogether, by probabilistic discounting of the risky low-value outcome.

Chapter 3: When Nothing Literally Matters

When Nothing Literally Matters

One plausible interpretation of the results seen in Experiment 1 is that of probabilistic discounting. Specifically, the probabilistic occurrence of the risky low-value's outcome might be diminishing its reinforcing value and subsequently be generating the risk-aversion seen on the low-value choices. Experiment 2 was designed to explicitly remove any confound of probabilistic discounting by eliminating the zero value from the range of possible outcomes encountered. Consequently, every outcome experienced would now present a positively reinforcing event.

This procedural change raises an obvious problem insofar as the pigeons are concerned. In Experiment 1a, learning effects were reliably obtained by some pigeons only on those trials that contained a possibility of receiving zero as an outcome. Trials containing no possibility of zero showed more variable amounts of learning. Any judgment about an organism's willingness to prefer a risky option to some other equally valued alternative requires an assumption that the contingencies involved have, in some sense, been learned or are "understood" by the organism being measured. In this design, catch trials are the only readily available means of assessing this in the pigeons. However, on the basis of Experiment 1a, we might predict that removing zero would have an altogether deleterious effect on their catch trial performance. Consequently, if the pigeons were to demonstrate, for instance, risk-neutrality on preference trials — a not uncommon finding in the animal literature (Behar, 1961; Mazur, 1989, Staddon Innis, 1966; Waddington, 1995; Wunderle & O'Brien, 1985) — that interpretation may reflect nothing more than an inability to learn the contingencies and not a genuine preference. As a solution to this problem, we opted to employ the more robust foraging procedure used in Ludvig, Madan, Pisklak, and Spetch (2014). This is the procedure upon which Experiment 1's operant protocol was based.

Experiment 2a: Pigeons

Subjects

As per Ludvig, Madan, Pisklak, and Spetch's (2014) and Experiment 1a, the subjects consisted of six pigeons (*Columba livia*) randomly selected from a University of Alberta colony room. Three of the birds were of Racing pigeon stock and the other three were of Silver King stock. All six birds had substantial learning histories with no history of risky-choice experiments and were kept on a 12 hour light-dark cycle. The birds were housed individually within metal cages, fed a diet of Mazuri Gamebird (PMI Nutrition International) food pellets, and had their food intake adjusted daily following the experimental sessions to keep them at approximately 85% of their free-feeding weight. Both grit and water containing a dissolved vitamin supplement were left freely available to the birds outside of the experiment. All the birds were run daily, five days a week, at the same time each day. The experimental procedures were all approved by the University of Alberta's Biological Sciences Animal Care and Use Committee, which follow the national guidelines set by the Canadian Council on Animal Care.

Apparatus

The experimental arena consisted of two compartments separated by a 50.8 cm long central wall. Each compartment was enclosed by a 91.4 cm long front and rear wall and a single 82.5 cm long side-wall. Both the central and rear walls were built of 1.27 cm thick plywood painted white. The front and side-walls were fashioned from thin white corrugated plastic.

Two 44.45 × 63.5 cm guillotine doors, each set at a 45° angle from the central wall, were positioned so as to form a small triangular *decision area* against the front of the arena. Both doors were fashioned from white corrugated plastic and operated via a single-string pulley system (Figure 2-1). A 16.51 × 20.32 cm entrance was cut into the front wall at the base of the decision area. This entrance led into plastic 48.26 × 39.37

× 101.6 cm start-box consisting of black opaque walls and a white translucent ceiling. In the start box, right and left compartments, and decision area, a thin layer of aspen chip bedding was laid down prior to the experiment. The room that housed the arena was illuminated by four fluorescent 40-W light bulbs mounted in the ceiling. With the exception of the start-box, the entire arena (ceilings included) was surrounded by white curtains to remove any extraneous visual stimuli that might be present in the surrounding room. The pigeon's behaviour was monitored and recorded by a closed circuit camera mounted centrally on the ceiling. A small hole was cut into the ceiling's curtain to permit the camera lens to view the arena.

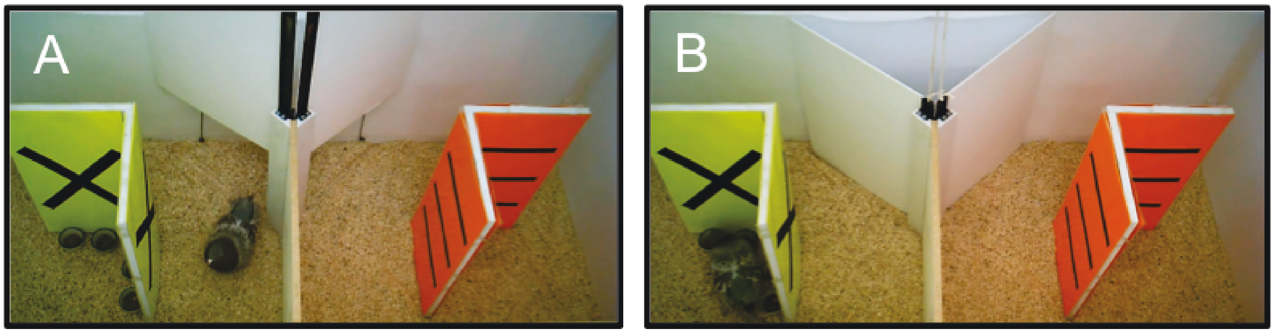

Figure 2-1. Shown is the testing arena during an example trial. (A) Pigeon exiting the decision area via the open guillotine doors. (B) The pigeon eating from the food cups concealed behind the choice stimulus.

Stimuli

Choice (i.e., discriminative) stimuli used for testing were built from two 30.48 × 58.42 cm planks of 2.54 cm thick plywood set at a 90° angle. Each was covered with four sheets of laminated paper that contained a distinct colour and design unique to that stimulus. In total, four different colours were used — green, orange, purple, and yellow. Their respective designs included a hollow black triangle, three horizontal black lines, four white squares, and a black "X" (see Figure 2-1 for an example). Two additional stimuli, coloured solid white, were also used for training purposes. The stimuli stood horizontally inside the centre of each of the arena's two main compartments to allow food rewards to be concealed. The food rewards were presented in 6.99 cm diameter ceramic cups located behind each choice stimulus. Each cup was operationally defined as 1 food reward and always contained two Mazuri Gamebird pellets on top of a small amount of grit.

Procedure

Phase 1. Preliminary training occurred over several days in a three-phase process. In the first phase, all pigeons were trained via successively reinforced approximations to enter each compartment from the start-box, through either the left or right guillotine door (randomly chosen), walk to the furthest corner to obtain three food rewards, and then return to the start box.

Phase 2. Once the pigeons reliably completed multiple sessions within Phase 1 the second phase began. In this phase, a white choice stimulus was placed in the centre of each compartment. The pigeons were gradually shaped to walk around this stimulus to obtain food rewards and return to the start-box as before. Once they could reliably complete 16 trials within a session, when no food rewards were directly visible from the decision area (i.e., with all food rewards concealed behind the stimulus), phase 3 was then introduced.

Phase 3. Phase 3 gave the pigeons 64 single-option trials randomly distributed over the course of four experimental sessions, with each session consisting of 16 trials. At the start of each trial, one guillotine door would open and then an opaque white plastic barrier blocking entry into the decision area from the start-box would be removed. Removing this barrier allowed a pigeon to enter the decision area and, subsequently, the compartment containing the food rewards. Once the pigeon had entered far enough into the compartment that the guillotine door could be safely closed without harm to the animal, the door was shut to prevent re-entry into the decision area. During this period, the pigeon could search behind the stimulus to obtain 1, 2, 3, 4, or 5 food rewards, with the amount contingent upon the stimulus' colour. Following consumption of the rewards, the guillotine door was re-opened allowing the pigeon to return to the start-box, which now contained a single food reward. Upon successful re-entry, the guillotine door would close and the barrier between the start-box and decision area would be put back in place.

The specific contingencies the birds were exposed to were designated as having either a *risky* or *safe* consequence, and of being either a *high-value* or *low-value* type. For instance, a stimulus with a safe high-value contingency meant that the pigeon would be guaranteed to receive four rewards, whereas a risky high-value contingency would result in a 50% chance of receiving three rewards or five rewards. Alternatively, a safe low-value contingency resulted in two guaranteed rewards, while the risky low-value contingency resulted in a 50% chance of receiving one or three rewards. The presentation of the stimuli was counter-balanced and randomized across the 64 trials such that no colour, reward value (risky vs. safe and high vs. low), or side was presented more than any other.

Testing. Following Phase 3, testing trials would begin and last for a total of 28 sessions. Each session lasted for 16 trials and, as per Experiment 1, contained four single-option trials and 12 choice trials. Within a session, the four single-option trials exposed the bird to each possible choice type (i.e., a risky high-value stimulus, a risky-low value stimulus, a safe high-value stimulus, and a safe low-value stimulus). The 12 choice trials consisted of eight *catch trials* and four *risk-preference trials.* Catch-trials required the pigeon to choose between two stimuli that lead to outcomes with unequal expected values (e.g., a high-value outcome vs. a low-value outcome) and were used to assess the pigeon's degree of learning in the task. Unlike Ludvig, Madan, Pisklak, and Spetch (2014), all possible types of catch trial were employed: risky vs. risky, safe vs. risky, safe vs. safe, and risky vs. safe. On risk-preference trials, the pigeon had to choose between two stimuli that lead to outcomes with equal expected values (e.g., a risky high-value contingency vs. a safe high-value contingency). Sessions were counterbalanced so that each stimulus appeared twice in both the right and left side compartments on choice trials. Ordering of the trials was randomized each session.

The testing procedure was identical to phase 3 with the exception that both guillotine doors now opened simultaneously on each trial. On single-option trials, one of the compartments would be left empty. If the pigeon happened to enter the empty side, both guillotine doors were left open until the pigeon entered the side with the stimulus.

Experiment 2b: Humans

Subjects

A total of 44 human participants (*Homo sapiens*) were recruited from the University of Alberta Psychology Subject Pool. Each subject was participating for course credit. Informed consent was obtained from all participants and all procedures were approved by the University of Alberta's Arts, Science, and Law Research Ethics Board. The age ($M \pm SD$) range was 20.3 ± 2.3 years, with 35 females. No other demographic information was collected.

Methods

Except for the outcome values encountered by the subject, Experiment 2b employed identical methods and procedures to those used in Experiment 1b. Risky high-value options could produce 3 or 5 rewards, whereas risky low-value options produced either 1 or 3 rewards. Safe high-value options guaranteed 4 rewards, and safe low-value options guaranteed 2 rewards. A schematic of the reward contingencies and their spatial layout on the screen can be seen in Figure 2-2.

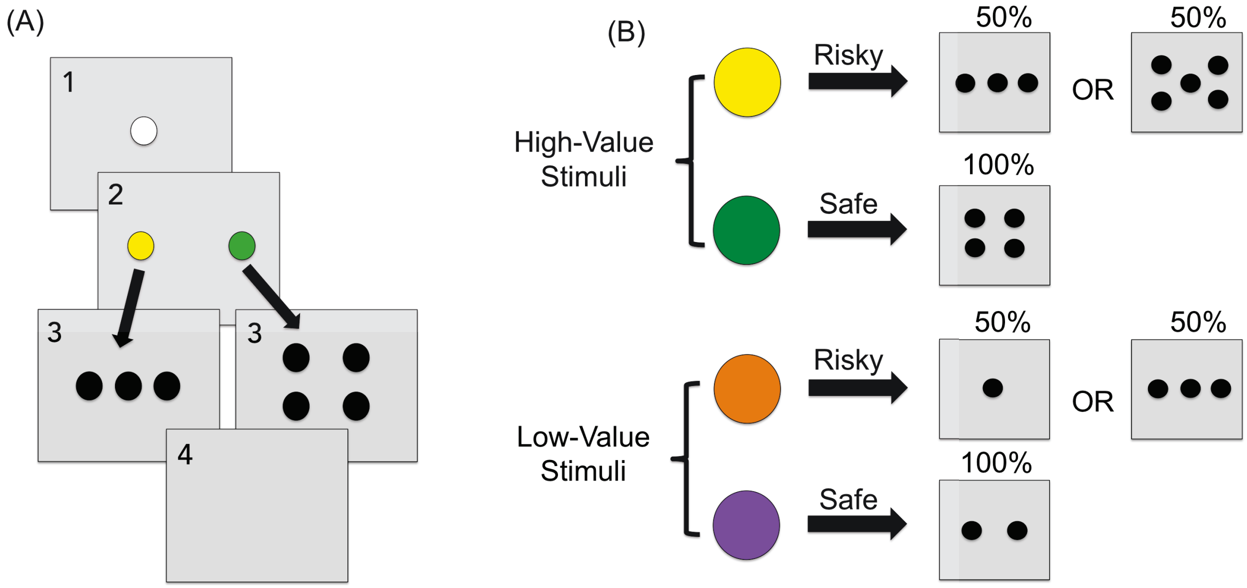

Figure 2-2. (A) A schematic of a choice trial displaying a risky high-value choice and a safe high-value choice, with respective outcomes. (B) A schematic illustrating the various reward contingencies in effect. The colour of the choice circles relative to their outcome was counter balanced across participants according to four different combinations of yellow, green, orange, and purple.

Results

For the pigeon data, statistical analysis was conducted on the last 20 instances of each choice type. As per Experiment 1b, statistical analysis of the human data was conducted on the last 3 blocks of 64 trials completed. Twenty participants who failed to learn the task — defined as choosing the high value option, across all catch trial types, less than 60% of the time — were excluded from the risk-preference analysis.

Figure 2-3 shows the mean proportion of risky choices on high- and low-value choice types along with the 95% confidence interval of the *paired difference.* Contrary to the findings of Experiment 1a and Ludvig, Madan, Pisklak, and Spetch (2014), pigeons showed no significant differences between the high-value (M = .44; 95% CI^{betw} [0.35, 0.53]) and the low-value (M = .43; 95% CI^{betw} [0.37, 0.50]) choice types, $t(5)$ =

0.16, $p = 0.876$, $d = 0.09$, 95% CI [-1.23, 1.44]). This is the case even in spite of near perfect performance on all catch trials tested (Figure 2-4).

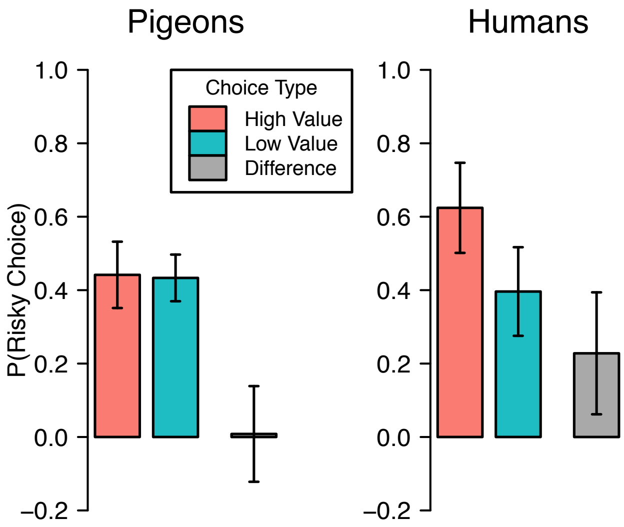

Figure 2-3. Mean proportion of choices to the risky option for pigeons (left) and humans (right) on high- and low-value risk-preference trials. The mean difference between both groups is plotted in grey. Error bars on the high- and low-value choice types indicate the 95% between-subject confidence interval. The error bar on the differences indicates the 95% confidence interval of the paired difference between the high- and low-value choice types.

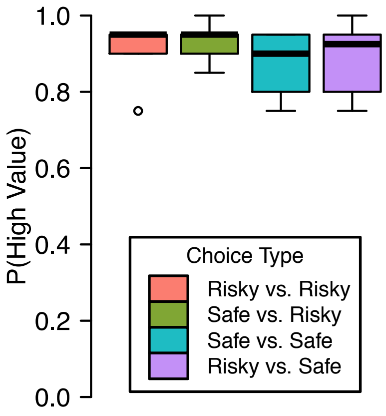

Figure 2-4. Boxplots displaying pigeon catch trial performance as a function of choice type.

In contrast to the pigeons, human risk-preference demonstrated significantly more risk-seeking for the high-value options ($M = .62$; 95% CI^{betw} [0.50, 0.75]) than the low-value options ($M = .40$; 95% CI^{betw} [0.28, 0.52]), $t(23) = 2.84$, $p = 0.009$, $d = 0.77$, 95% CI [0.19, 1.36] (observed power = 0.95). However, these results should be seen in the context of the entire sample, as they represent only about half (55%) of the participants tested. Of the 44 subjects run, 20 failed to pass the catch trials. When the distributions of catch trial scores are viewed for both the passes and the fails, a clear dichotomy emerges (see Figure 2-5a). The subset of the participants passing the catch trials scored at or near ceiling levels, while the subset of participants failing the catch trials consistently scored at or near chance levels. There is virtually no middle ground between these two sets of participants.

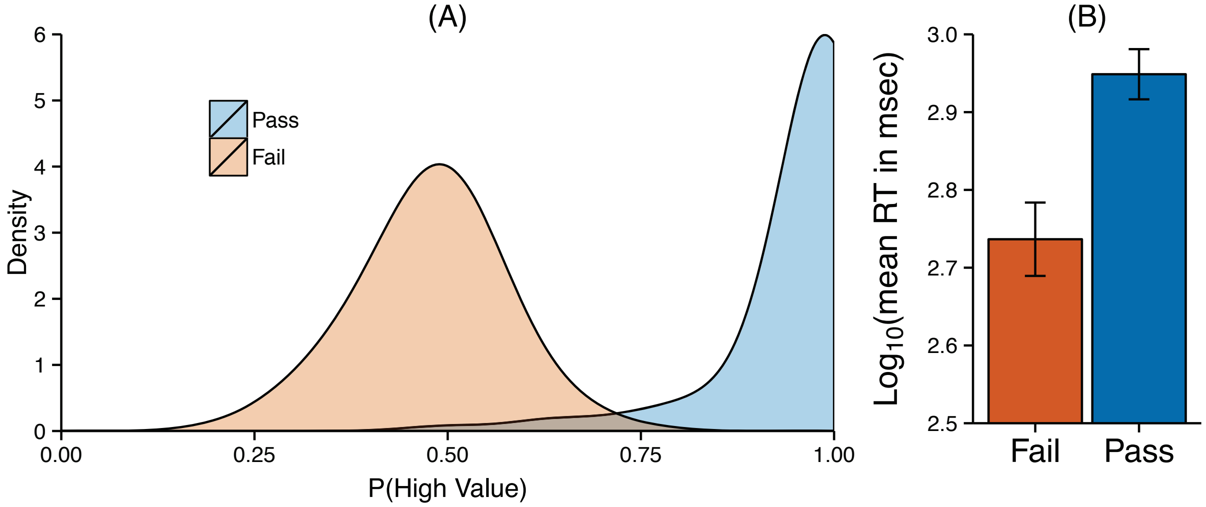

Figure 2-5. (A) Density plot illustrating the distributions of the obtained human catch trial scores using a Gaussian smoothing kernel. (B) The common log mean reaction time in milliseconds as a function of catch trial learning in the last three blocks of 64 trials. Error bars show the 95% between-subject confidence interval of the mean.

Discussion

In Ludvig, Madan, Pisklak, and Spetch (2014) it was argued that their obtained pigeon results were not likely due to any effect of zero value avoidance. However, using a near identical procedure, Experiment 2a unambiguously demonstrated no substantive risk-preference when zero value outcomes were removed. This finding, in conjunction with the results of Experiment 1a, suggest that the pigeon results of Ludvig et al. (2014) may represent an avoidance of zero explained in terms of probabilistic discounting (Mazur, 1989).

Unlike the pigeons, the subset of human participants who did manage to pass the catch trials still showed the predicted effects of the extreme-outcome rule, albeit with a rather diminished effect size relative to Experiment 1b, $d = 0.77$ versus $d = 0.93$. It is possible that the effects of probabilistic discounting were inflating the difference between the high- and low-value groups in Experiment 1b and correspondingly producing an exaggerated effect size. This idea is somewhat corroborated by the fact that the high-value risk-preference scores were relatively invariant across Experiment 1b and 2b, whereas the low-value risk-preference scores showed a noticeable increase in risky choice selection.

Interestingly, the human results also show an effect of removing the zero value outcome, but not in the same manner as that displayed by the pigeons. For the humans, removing zero markedly impaired learning, with 45% of the sampled population failing to meet the set catch trial criteria. By contrast, in Experiment 1b only 23% of the sample had failed. The numbers are striking and there is no obvious explanation as to why this should occur. One possibility, though, is that the participants who did fail the catch trials may have been employing an altogether different strategy of maximising their rewards. Unaware that there was a fixed amount of trials for them to complete, they may have been simply trying to get through as many trials as possible, in as short of time as possible. To assess the feasibility of this hypothesis, a *post hoc* independent t-test was used to examine the difference between choice response times across the two groups. Results are shown in Figure 2-5b and indicate a clear group difference between passes ($M = 889$ msec) and fails ($M = 545$ msec) when the data was logarithmically transformed to satisfy the test's distributional assumptions: $t(42) = 7.94$, $p < 0.001$, $d = 2.36$, 95% CI [1.61, 3.18]. While by no means definitive, these results are in accordance with the idea of an altered reward strategy. Moreover, these findings are consistent with those reported in studies that have assessed instructional control of operant responses in humans. Often, when instructions are given to the effect that a certain response needs to occur to obtain a reinforcer, the experimenter is drawing on — in a discriminative fashion — a complex history in which reinforcement is proportional to response output. This, not surprisingly, tends to generate high rates of responding inappropriate to the schedule of reinforcement in effect (Baron & Galizio, 1983; see also Skinner, 1969). An analogous situation in pigeons is described by Biglam and Kass (1977):

> Suppose we establish stimulus control over fixed interval and fixed ratio responding of a pigeon, using arbitrary stimuli. Let the controlling stimulus for the fixed interval schedule be the letters "FI" and let the stimulus for the fixed ratio schedules be the letters "FR." Now suppose that, once stimulus control is established, we present the stimulus "FI" when the pigeon is in fact on a fixed ratio schedule. It appears safe to assume, on the basis of available research, that the pigeon will respond "as though it believes that the schedule is a fixed interval." (p. 11)

The crucial point here is that, in the human case, when the instructions — "*use the mouse cursor to click on the circles that you see. Try and get as many points as possible . . . The more points you get the more money you will earn*" — are administered, we (the experimenters) may be failing to appreciate the type of discriminative control being imposed on the behaviour. A seemingly obvious, yet lamentably often overlooked, detail noted by Baron and Galizio (1983) is relevant here. When one studies discriminative function (i.e., stimulus control) in animal laboratory settings "the characteristics of the stimuli can be specified, and considerable information is available about the schedule and other aspects of procedures *leading to* discriminative control" (p. 508). This is not the case with human research on instructions "where the controlling properties of the stimuli have been established prior to the start of the experiment" and "little is known about the circumstances under which these properties have developed" (p. 508). Presumably then, if the poor performance exhibited by the humans in Experiment 2b is an artefact of inappropriate discriminative control, then instructions designed to be more evocative of the contingencies should improve the catch trial performance. If, however, the poor performance exhibited by the humans is the result of some unobservable construct such as boredom or apathy with the task, then we might expect a slight alteration of the instructions to have no meaningful effect on performance. Experiment 3 was designed to test precisely this hypothesis.

Experiment 3

In order to test the hypothesis that the large subset of catch trial fails seen in Experiment 2b was an artefact of poor discriminative control, a sample of 24 people (*Homo sapiens*) with an age range ($M \pm SD$) of 19.4 ± 1.5 years and 13 females was employed. All subjects were participating for course credit through the University of Alberta's Psychology Subject Pool. Informed consent was obtained from all participants and all procedures were approved by the University of Alberta's Arts, Science, and Law Research Ethics Board. Procedurally, the task was identical to that used in Experiment 2b with the exception that the instructions read aloud to the participants were altered to the following: *Your goal is to earn as many points as you can. After clicking on a white centre circle, you will see one or two coloured circles on the computer screen. You choose a coloured circle by clicking on it with the mouse. After clicking the coloured circle you will see one or more black dots, clicking each black dot will give you one point. When there are two circles you should choose the one you think will win you* *the most black dots (i.e. the most points). If there is only one circle on the screen, you must click on that one circle to continue. For the maze portion: a message will appear on screen telling you when to complete a particular maze. Take a few minutes to complete the maze with a pen or pencil. All the mazes are solvable, but don't worry if you can't solve it. The experiment should take just under 2 hours to complete. You will earn 2 credits for participating, and receive a cash bonus of up to $5.00 depending upon the number of points you get in the computer portion. The more points you get the more money you will earn.*

Results

As before, statistical analysis of the observed risk preference was conducted on the last 3 blocks of 64 trials completed. The criterion for task failure was also kept the same: choosing the high value option, across all catch trial types, less than 60% of the time.

In terms of catch trial performance, none of the participants tested failed the catch trials. In fact, performance was at or near ceiling levels across all subjects ($M =$ .99). A chi-squared test between the two modes of instruction, in Experiment 2b and 3 respectively, showed the odds of passing the catch trials are significantly lowered by receiving the original (Experiment 2b) instructions; $\chi^2(1) = 15.45$, $p < 0.001$. The observed levels of risk-preference were remarkably comparable to those seen in Experiment 2b's participants, again showing significantly more risk-seeking for the high-value ($M = .58$; 95% $\mathrm{CI}^{\mathrm{betw}}$ [0.42, 0.74]) than the low-value options ($M = .37$; 95% $\mathrm{CI}^{\mathrm{betw}}$ [0.25, 0.49]), $t(23) = 2.59$, $p = 0.02$, $d = 0.60$, 95% CI [0.11, 1.11].

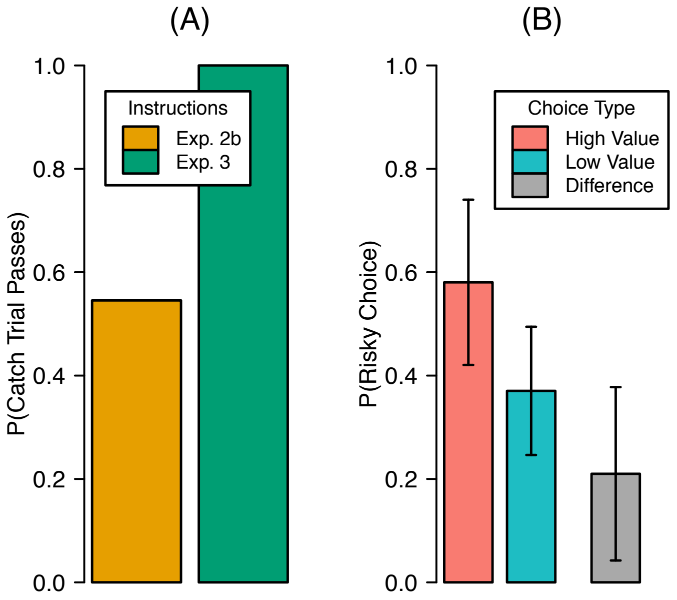

Figure 2-6. (A) Proportion of participants who passed the catch trial in Experiment 2b and 3. (B) Mean proportion of choices to the risky option for the humans in Experiment 3 on high- and low-value risk-preference trials. The mean difference between both groups is plotted in grey. Error bars on the high- and low-value choice types indicate the 95% between-subject confidence interval. The error bar on the differences indicates the 95% confidence interval of the paired difference between the high- and low-value choice types.

Discussion

One straightforward interpretation of the poor catch trial performance seen in Experiment 2b stems from a construct-based account which posits an intervening form of paramechanistic phenomena (Ryle, 2009) as the principle cause; boredom and apathy being two such examples (though any number of more elaborate cognitive processes could be appealed to as well). However, the near flawless performance observed in Experiment 3 seems to warrant a more parsimonious explanation in terms of discriminative function of the instruction stimuli. Simply put, the Experiment 2b instructions said to '*click the circles*' and '*get as many points as possible*' within the two hours, so that is what many tried to do. There were no details suggesting how best to go about this, so it is not unreasonable that some chose to go about this in one way, carefully responding on the basis of the differentially available black circles, while others chose to go about it in another way, getting through as many trials as fast as possible. Perhaps this latter method seems rather senseless on the face of it. But one can easily imagine a scenario where such a method would pay off more than the other — as in a case where the experiment had a fixed time limit, as opposed to a fixed trial limit. Why anyone should interpret these instructions one way as opposed to the other is probably best accounted for in terms of each individual's own unique history of reinforcement and punishment, something for which no ethical experimental control exists in humans. When the instructions were modified to be more evocative of the contingencies, including in particular the line '*when there are two circles you should choose the one you think will win you the most black dots (i.e. the most points),'* not surprisingly the catch trial performance improved markedly.

Alongside all of this is the fact that the altered instructions of Experiment 3 seemed to have no bearing on the observed patterns of risk-preference. Risky choice proportions were nearly identical to those seen in the catch trial passes of Experiment 2b, with the predictions of the extreme-outcome rule preserved.

Chapter 4: General Discussion

General Discussion

The first of the three experiments attempted to replicate the findings of Ludvig, Madan, Pisklak, and Spetch (2014), and by extension the predictions of the extreme-outcome rule, in a operant chamber analogue of their pigeon foraging task. The results were generally quite favourable to the extreme outcome-rule's predictions when tested in both pigeon and human species. However, some catch trial anomalies observed in the pigeon data warranted scepticism and suggested that an explanation in terms of probabilistic discounting was not unreasonable and should be explored further. The second of the three experiments did just that. When zero was removed as a possible outcome from choices that offered a risky low-value option, thus eliminating the possibility of probabilistic discounting taking place, the patterns of pigeon risk-preference observed in Experiment 1 disappeared completely revealing a stark level of risk-neutrality. The pattern of human data, however, remained consistent with the predictions of the extreme-outcome rule, albeit with a noticeably diminished effect size. Interestingly, these human results were accompanied by a severe learning deficit in a large proportion of the sampled individuals. Speculating that this deficit may be the result of a different strategy to maximize rewards brought on by ambiguous discriminative control, Experiment 3 altered the instructions to be more evocative of the contingencies at play. This subsequently produced near ceiling levels of performance across all individuals tested while still preserving the basic predictions of the extreme-outcome rule at proportions nearly identical to those seen in Experiment 2's human subjects.

Taking the results at face value, it would seem that the risky choice of pigeons and humans, suggested by Ludvig, Madan, Pisklak, and Spetch (2014), is not so similar as had been supposed. Pigeons are largely indifferent to risk when expected values are equated and probabilistic discounting is not a confound. However, seen in the broader context of the pigeon risk-preference literature, such a conclusion would be rather ill-considered as, across different laboratories, findings of risk-aversion (Hamm & Shettleworth, 1987; Menlove, Inden, & Madden, 1979) and risk-seeking (Essock & Reese, 1974; Young, 1971; Lagorio & Hackenberg, 2012), along with risk-neutrality (Staddon & Innis, 1966) are not uncommon. This suggests that procedural details may be very important in determining how risk-preference is exhibited, a conclusion others have made as well (Hayden & Platt, 2009; Mazur, 2004).

In contrast to the large differences observed in the two pigeon experiments, the results of the human subjects showed the basic predictions of extreme-outcome rule to be fairly robust when participants were able to learn the catch trials. A rather unexpected finding, however, was the degree to which instructions played a role on the human behaviour and how this interacted with the presence and absence of the zero value outcome. Contrasting Experiments 1b and 2b, it can be seen that zero had a large influence on how participants reacted to the different amounts of reinforcement provided by the four options. In the former case, the presence of zero seemed to aid discrimination of the catch trial choices and in the latter case the absence of zero seemed to hinder discrimination for a large subset of the sample. The decreased levels of discrimination seen as a result of zero's removal could, however, be completely attenuated by simply altering the instructions to state what the "appropriate" response strategy should be, thereby placing further emphasis on the importance of procedural factors.

One complication raised by the observed influence of instructions concerns the validity of comparison between pigeons and humans. Much effort went into making the task for both species as similar as possible; however, the problem posed by the role of instructions in the human case revealed a glaring difference that had not been initially considered. Specifically, this difference pertained to the role of discriminative cues in the task. In the case of the pigeons, the behaviour of learning to peck the circles and search in the cups was shaped successively. Prior experience in the form of a controlling discriminative stimulus was not relied upon. Ideally one would want to take the same approach of gradually shaping the behaviour with the humans as well. But of course this tactic, while preferable, is not so feasible. Consequently, we have in the human case a scenario in which the presentation of instructions (i.e., discriminative stimuli) has to serve as a substitute for shaping the behaviour directly. The results of Experiment 2 and 3, as well as research on instructional control of human operant learning (for a review see Baron & Galizio, 1983), show that this is a far from foolproof method because it requires the assumption that the discriminative function of the instructions will not have any bearing on the independent and dependent variables of the task above and beyond what standard shaping would. Given that no control can be exerted on the learning histories that govern the discriminative function of the instructions (or the task more generally for that matter), this could be seen as a rather tenuous assumption to make.

One question that has yet to be considered concerns the general applied utility of the extreme-outcome rule. Even if we assume that the functional role of the instructions is of little consequence and the extreme-outcome rule represents a genuine population effect in humans, we are compelled to ask: to what extent is this population effect generalizable to individuals? Presumably it is the behaviour of actual individuals and not aggregates of individuals that we would like to explain. A common fallacy made of population statistics (t-tests, ANOVAs, etc.) is that they permit a downward generalization from populations to individuals. Only under rare circumstances is this ever the case (Branch & Pennypacker, 2013). It is therefore necessary that some consideration be made of the variability within the risk-preference choice types. Figure 3-1 illustrates the distributional information of the risk-preference trials across all three human experiments. Notably, the spread of the data in both the high- and low-value case is quite large, even in spite of the control imposed by laboratory conditions. Consequently, the likelihood of the extreme-outcome rule's predictions being applicable in any given individual case is not especially compelling. Since the rule's utility with respect to individuals appears limited, it is perhaps best to conceptualize it in terms of a population effect only.

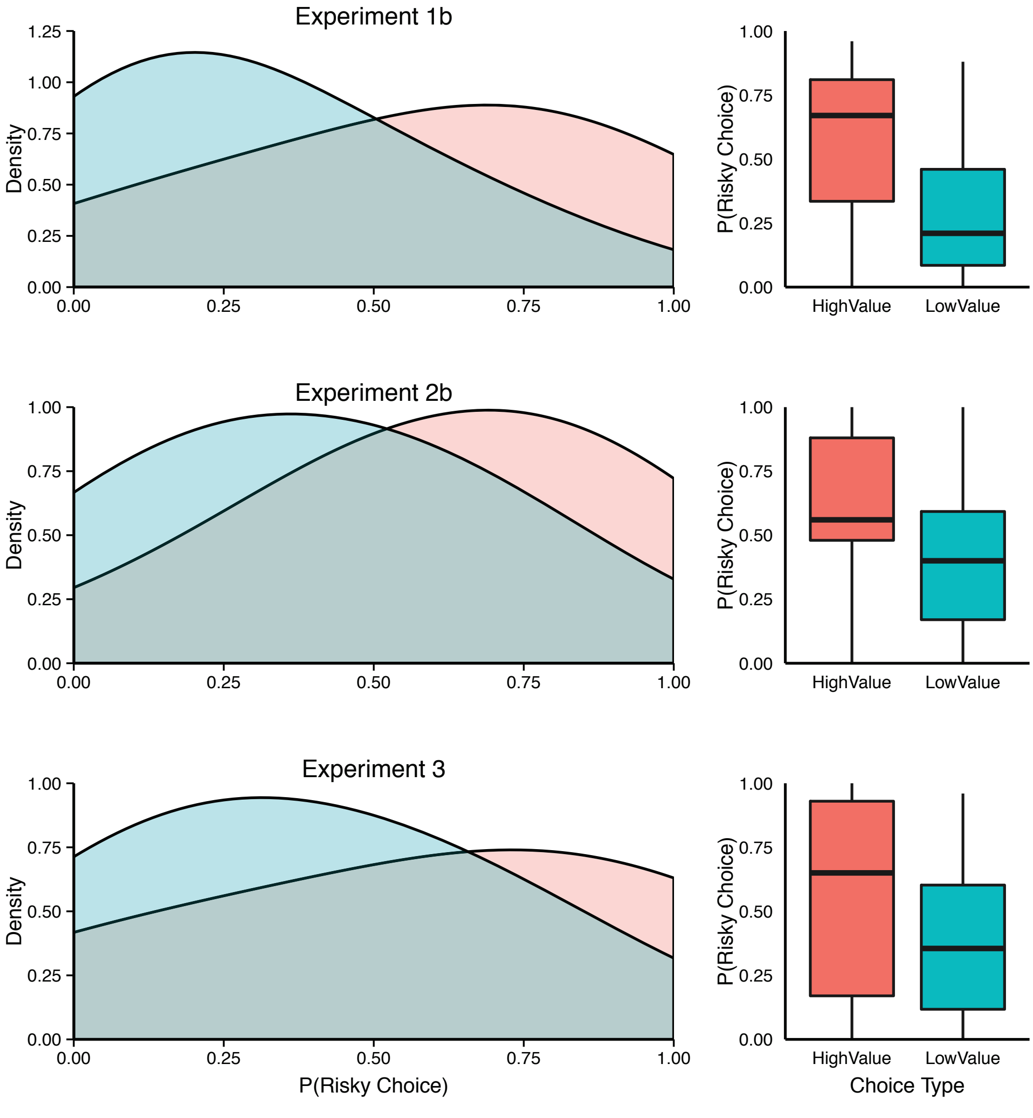

Figure 3-1. (Left) Density plots illustrating the distributions of the obtained human risk-preference scores across all three experiments using a Gaussian smoothing kernel. (Right) Corresponding boxplots of each experiment's risk-preference distributions. Each graph reflects only those subjects who satisfied the catch trial criteria.

All things considered, the extreme-outcome rule does seem to show some promise as a predictive tool in certain instances of human behaviour. This seems especially the case for larger aggregates of data that are more actuarial in nature. Its applicability to individuals and to other species, however, remains to be established. One possibility is that the true relevance of extreme outcomes on behaviour generally cannot be fully appreciated (in animals or humans) until the variances between the range of possible outcomes becomes quite drastic, as is the case in lottery or jackpot wins. Future appraisals of the rule will need to look at this possibility before any solid conclusion about its relevance to human and non-human affairs can be made.

References

Barnard, C. J., Brown, C. A. J., Houston, A., & McNamara, J. M., (1985). Risk-sensitive foraging in common shrews: An interruption model and the effects of mean and variance in reward rate. *Behavioral Ecology and Sociobiology, 18*, 139-146. doi:10.1007/BF00299042

Baron, A., & Galizio, M. (1983). Instructional control of human operant behavior. *Psychological Record,33*, 495-520.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215-233. doi:10.1002/bdm.443

Behar, I. (1961). Learned avoidance of non reward. *Psychological Report*, *9*, 43-52.

Biglan, A., & Kass, D. J. (1977). The empirical nature of behavior therapies. *Behaviorism*, *5*, 1-16.

Branch, M., & Pennypacker, H. (2013). Generality and generalization of research findings. In *APA Handbook of Behavior Analysis* (Vol. 1, pp. 151-167). Washington, DC: American Psychological Association.

Brown, S., & Heathcote, A. (2003). Averaging learning curves across and within participants. *Behavior Research Methods, Instruments, & Computers*, *35*, 11-21. doi:10.3758/BF03195493

Camilleri, A. R., & Newell, B. R. (2009). Within-subject preference reversals in description- and experience-based choice. *Cognitive Science Society*, *31*, 449-454.

Clements, K. C. (1990). Risk-aversion in the foraging blue jay, *Cyanocitta cristata. Animal Behavior, 40*, 182-183. doi:10.1016/S0003-3472(05)80678-5

Chung, S. H. (1965). Effects of delayed reinforcement in a concurrent situation. *Journal of the Experimental Analysis of Behavior. 8*, 439-444. doi:10.1901/jeab.1965.8-439

Cumming, G. (2011). Cohen's d. In *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis.* New York, NY: Routledge.

Essock, S. M., & Reese, E. P. (1974). Preference for and effects of variable — as opposed to fixed — reinforcer duration. *Journal of Experimental Analysis of Behavior*, *21*, 89-87. doi:10.1901/jeab.1974.21-89

Franz, V., & Loftus, G. (2012). Standard errors and confidence intervals in within-subjects designs: Generalizing Loftus and Masson (1994) and avoiding the biases of alternative accounts. *Psychonomic Bulletin & Review*, *19*, 395–404. doi:10.3758/s13423-012-0230-1

Green, L., Myerson, J., & Calvert, A. (2010). Pigeons' discounting of probabilistic and delayed reinforcers. *Journal of the Experimental Analysis of Behavior*, *94*, 113-123. doi:10.1901/jeab.2010.94-113

Hamm, S. L., & Shettleworth, S. J. (1987). Risk aversion in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *13*, 376-383. doi:10.1037/0097-7403.13.4.376

Hayden, B. Y., & Platt, M. L., (2009). Gambling for Gatorade: Risk-sensitive decision making for fluid rewards in humans. *Animal Cognition*, *12*, 201-207. doi:10.1007/s10071-008-0186-8

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517-523. doi:10.1016/j.tics.2009.09.004

Hertwig, R., G, B., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534-539. doi:10.1111/j.0956-7976.2004.00715.x

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291. doi:10.2307/1914185

Lagorio, C., & Hackenberg, T. (2012). Risky choice in pigeons: Preference for amount variability using a token-reinforcement system. *Journal of the Experimental Analysis of Behavior*, *98*, 139-154. doi:10.1901/jeab.2012.98-139

Ludvig, E., & Spetch, M. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events. *PLOS ONE*, *6.* doi:10.1371/journal.pone.0020262

Ludvig, E., Madan, C., & Spetch, M. (2014). Extreme outcomes sway risky decisions from experience. *Journal of Behavioral Decision Making*, *27*, 145-156. doi:10.1002/bdm.1792

Ludvig, E., Madan, C., Pisklak, J., & Spetch, M. (2014). Reward context determines risky choice in pigeons and humans. *Biology Letters*, *10.* doi:10.1098/rsbl.2014.0451

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin and Review*, *21*, 629-636. doi:10.3758/s13423-013-0542-9

Marshall, A. T., & Kirkpatrick, K. (2013). The effects of the previous outcome on probabilistic choice in rats. *Journal of Experimental Psychology-Animal Behavior Processes*, *39*, 24-38. doi:10.1037/a0030765

Mazur, J. E. (1984). Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 426-436. doi: 10.1037/0097-7403.10.4.426

Mazur, J. E. (1989). Theories of probabilistic reinforcement. *Journal of the Experimental Analysis of Behavior*, *51*, 87-99. doi:10.1901/jeab.1989.51-87

Mazur, J. E. (2004). Risky choice: Selecting between certain and uncertain outcomes. *The Behavior Analyst Today*, *5*, 190-203. doi:10.1037/h0100031

McDiarmid, C. G., & Rilling, M. E. (1965). Reinforcement delay and reinforcement rate as determinants of schedule preference. *Psychonomic Science*, *2*, 195-196. doi:10.3758/BF03343402

Menlove, R. L., Inden, H. M., & Madden, E. G. (1979). Preference for fixed over variable access to food. *Animal Learning and Behavior*, *7*, 499-503. doi:10.3758/BF03209709

Payne, J., Laughhunn, D., & Crum, R. (1980). Translation of gambles and aspiration level effects in risky choice behavior. *Management Science*, *26*, 1039-1060. doi:10.1287/mnsc.26.10.1039

Rachlin, H., Logue, A., Gibbon, J., & Frankel, M. (1986). Cognition and behavior in studies of choice. *Psychological Review*, *93*, 33-45. doi:10.1037/0033-295X.93.1.33

Ryle, G. (2009). Descartes' myth. In *The concept of mind* (60th Anniversary ed., pp. 1-13). New York, NY: Routledge. (Original work published 1949)

Schoemaker, P. (1990). Are risk-attitudes related across domains and response modes? *Management Science*, *36*, 1451-1463.

Skinner, B. F. (1969). An operant analysis of problem solving. In *Contingencies of reinforcement: A theoretical analysis.* New York: Appleton-Century-Crofts.

Staddon, J. R., & Innis, N. K. (1966). Preference for fixed vs. variable amounts of reward. *Psychonomic Science, 4*, 193-194. doi:10.3758/BF03342245

Stephens, D. W. (1981). The logic of risk-sensitive foraging preferences. *Animal Behavior. 29*, 628-629. doi:10.1016/S0003-3472(81)80128-5

Stevens, J. R., & Stephens, D. W. (2009). The adaptive nature of impulsivity. In G. J. Madden & W. K. Bickel (Ed.), *The Behavioral and Neurological Science of Discounting* (pp. 361-388). Washington, DC: American Psychological Association.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *The Journal of Business*, *59*, S251-S278. doi:10.1086/296365

Ungemach C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science. 20*. 473-479. doi:10.1111/j.1467-9280.2009.02319.x

Waddington, K. D. (1995). Bumblebees do not respond to variance in nectar concentration. *Ethology, 101*, 33-38. doi:10.1111/j.1439-0310.1995.tb00342.x

Wunderle, J. M., & O'Brien, T. G., (1985). Risk aversion in hand-reared bananaquits. *Behavioral Ecology and Sociobiology. 17*, 371-380. doi:10.1007/BF00293215

Young, J. S. (1981). Discrete-trial choice in pigeons: Effects of reinforcer magnitude. *Journal of the Experimental Analysis of Behavior. 35*, 23-29. doi:10.1901/jeab.1981.35-23