Gain-loss asymmetry in human experiential tasks

By

Rana Asgarova

A thesis

submitted to the Victoria University of Wellington

in fulfilment of the requirements for the degree of

Doctor of Philosophy in Psychology

Victoria University of Wellington

2019

Acknowledgments

I would like to thank my supervisors, Maree Hunt and Anne Macaskill, for their patience and enthusiasm in their guidance and encouragement at all stages of my time at the university. Should someone accuse me of being well educated, I will be proud and happy to name you as my principal aiders and abettors. I would like to extend this to my many accomplices from the Human Learning Lab, whose good company provided a happy place for sharing thoughts and snacks. I also want to acknowledge CORE and NZABA groups for the open sharing of ideas and data that provided numerous insights into my own research.

I want to thank my family and friends who have provided me with constant support from across the many continents, and I want to thank my husband, who good-naturedly crossed the seas with me and saw me through every struggle.

Finally, this thesis would not have been possible without the financial support of the Victoria University of Wellington Doctoral Scholarship.

Abstract

Prospect Theory models behaviour in one-off decisions where outcomes are described. Prospect Theory describes risk aversion when the choice is between gains and risk seeking when the choice is between losses. This asymmetry is known as the reflection effect. In choices about experienced outcomes, individuals show risk seeking for gains and risk aversion for losses. This change in the direction of gain-loss asymmetry is known as the description-experience gap. Across eight experiments, we examined gain-loss asymmetry in two experiential choice procedures. We compared the obtained results with predictions derived from Prospect Theory and the description-experience gap literature.

In Study 1, we evaluated the predictions of the reversed reflection effect in probability discounting. Probability discounting is loss in reinforcer value as a function of uncertainty. In typical tasks measuring discounting, participants choose between smaller, certain amounts and a larger amount at one of several probabilities. In choice from description, most participants show a gain-loss asymmetry consistent with the predictions of the reflection effect, discounting gains more steeply than losses. Across three experiments, we examined whether gain-loss asymmetry also occurred when participants experienced the outcomes they chose, when they chose between two uncertain options, and when these two contexts were combined. Across all of the above contexts, no consistent mean difference in discounting of gains and losses was observed. Rather, in most of the tasks that provided experienced outcomes, the participants showed steeper discounting in the first condition completed, whether it involved choices about gains or losses. Furthermore, subsequent conditions produced shallower discounting, but notably, not shallower than choice based on the expected value of the options. In Studies 2 and 3, we followed-up on this order effect by providing the participants with experience of probabilistic outcomes before the discounting tasks. Participants discounted losses more steeply than gains, consistent with the predictions of a reversed reflection effect.

In Study 2, we examined gain-loss asymmetry in a rapid-acquisition choice procedure using concurrent variable-interval schedules – the Auckland Card Task. Participants repeatedly chose between two decks of cards that varied in the frequency or magnitude of available gains or losses. Participants were more sensitive to changes in gain than loss frequency between the two decks, consistent with the predictions of a reversed reflection effect, while sensitivity to gain and loss magnitude did not show an asymmetry. We found a novel

asymmetry in the local effects of gains and losses. In the frequency tasks, gains disrupted the general pattern of responding more than losses. In the magnitude tasks, varying the magnitude of losses had a bigger effect on local-level patterns following outcomes than varying the magnitude of gains.

Across the two tasks we observed patterns of gain-loss asymmetry consistent with the predictions of a reversed reflection effect. We also observed several inconsistencies, particularly when comparing behaviour to choices that would maximize the expected returns. Our research suggested that sufficient exposure to chance outcomes and ensuring delivery of scheduled events are key challenges in further refinement of experiential choice in human operant tasks.

Table of Contents

List of Tables

List of Figures

General Introduction

Gain-loss asymmetry in choice from description

Rational decision making predicts that when an individual is faced with probabilistic outcomes, they ought to choose according to the expected return, or expected value, of the options. Expected value is calculated as amount multiplied by probability. For example, the expected value of a 0.50 probability of $100 is $50, and this is greater than the expected value of a 1.00 probability of $40, which is $40. Thus rational choice would be to choose the first option which has the highest expected value, yet human behaviour has consistently shown systematic violations of expected value predictions. Expected Utility Theory states that the subjective value of an outcome to the individual often varies from the expected value by assigning a numerical *utility* to these outcomes (Bernoulli, 1738; Von Neumann & Morgenstern, 1944). It further postulates that individuals ought to choose options that will maximize the expected utility of the outcomes, but this too has often been shown to be inconsistent with choice (e.g. Allais, 1953; Starmer, 2000; Bleichrodt, Abellan-Perpinan, Pinto-Prades, & Mendez-Martinez, 2007). One of the most popular theories that makes predictions about the systematic difference in behaviour when faced with risky gains or losses is Prospect Theory (Kahneman & Tversky, 1979) and its modified version, Cumulative Prospect Theory (Tversky & Kahneman, 1992). The central premise is that under some conditions, neither the expected value of a given prospect, nor its expected utility, predict behaviour. Here, we focus on three principles regarding decision under risk modelled by Prospect Theory that are key to this series of studies: loss aversion, probability weighting and the reflection effect. Together, these constitute a gain-loss asymmetry in how humans respond to probabilistic gains compared to probabilistic losses.

Loss aversion captures the tendency for people to be more affected by the prospect of a loss than the prospect of a gain of the same magnitude. The way in which Prospect Theory models loss aversion is through the value function illustrated in the left panel of Figure 1.1, by plotting a range of gain and loss outcomes on the x axis against the subjective value an individual might assign to each outcome on the y axis[1]. Prospect Theory predicts that change

[1]The value functions for outcomes zero and above was plotted using $v(x) = x^{\alpha}$ and for outcomes below zero was $v(x) = -\lambda(-x)^{\beta}$. Here, v is the value function applied to outcome x, α and β are scaling parameters describing sensitivity to value (note that we assumed α = β) and λ is the loss aversion scaling parameter. In our example, α and β were set to 0.88 and λ to 2.25 (Tversky & Kahneman, 1992).

in subjective value for negative outcomes, like losses, is steeper than change in subjective value for positive outcomes, like gains, that are equivalent in magnitude. This difference in rate captures the prediction of Prospect Theory that an outcome of -75 is not symmetrical in subjective value to an outcome of +75 (both indicated by dashed arrows); losing $75 is subjectively worse than gaining $75 is good.

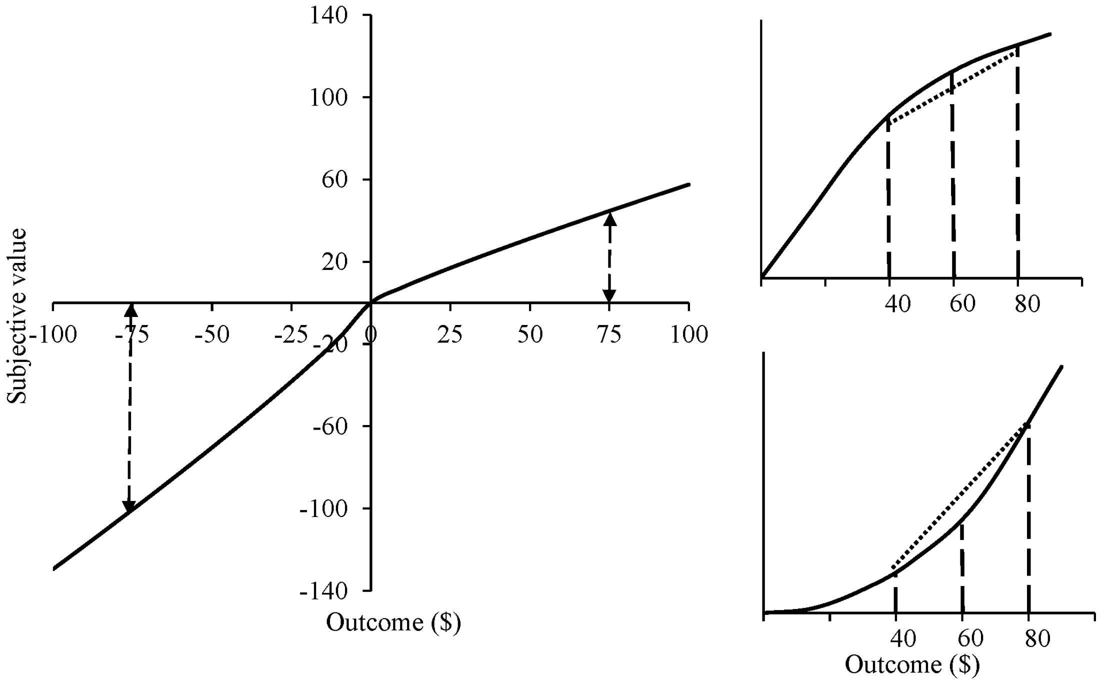

Figure 1.1. Left panel: value function that plots outcomes on the x axis and subjective value on the y axis. The function plots a hypothetical rate at which subjective value of positive and negative outcomes changes with change in magnitude as predicted by Prospect Theory. Right panel: concave (top) and convex (bottom) value functions. The solid line corresponds to the value function described in the left panel. The dotted line indicates change in subjective value of the plotted outcomes if the relationship was linear rather than concave/convex.

Furthermore, the shape of the function described by Prospect Theory for positive outcomes is concave, predicting risk aversion, while for negative outcomes it is convex, predicting risk seeking (Kahneman & Tversky, 1979). This is better illustrated on the right panel of Figure 1.1 which plots concave (top) and convex (bottom) functions, with monetary amounts on the x axis and their subjective value to the individual on the y axis. The dotted line indicates risk neutrality and choice in accordance to expected value, where an increase in outcome amount is accompanied by a linear increase in subjective value. Consider a choice between gaining $60 for certain and a risky prospect of a 50% chance of $40 and a 50% chance of $80. In terms of expected value, the two options are equivalent and the individual should show indifference. However, the top graph demonstrates that if the change in subjective value is concave rather than linear, then the individual would prefer to receive $60 for certain rather than the expected value of the risky prospect, given that at the $60 mark on the x-axis, the value function is higher than the dotted line on the y-axis. Conversely, if the change is convex (bottom graph), the value function predicts that the individual would prefer to risk the prospect rather than gain $60 for certain. The value function of Prospect Theory predicts a concave function for choices about gains and a convex function for choices about losses. Hence, a concave change in the value of a gain would predict risk averse behaviour and a convex change in the value of a loss would predict risk seeking behaviour.

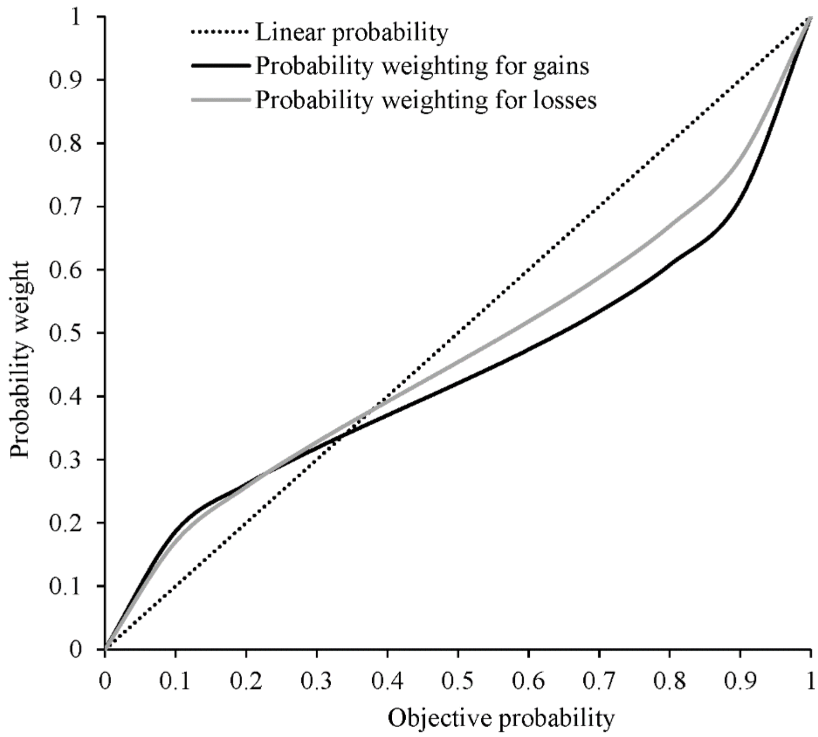

Figure 1.2. Probability weighting function for choice from description that plots objective probability on the x axis and weighted probability on the y axis. The dotted line corresponds to linear weighting of probabilities. The solid black and grey lines correspond to probability weighting predicted by Prospect Theory for gains and losses respectively.

Probability weighting is the application of decision weights to objective probabilities. Figure 1.2 demonstrates Prospect Theory’s prediction by plotting objective probability on the x axis and weighted (distorted) probability on the y axis. The dotted line shows linear weighting of probability and the solid black and grey lines show probabilities for gains and losses transformed by a weighting function[2]. Note that the weighting function assumes that small probabilities will be overweighted and medium to large probabilities underweighted relative to the dotted line, showing an inverse S-shaped function. The greatest deviation from linear weighting of probabilities can be observed at probabilities that are very high or very low. The separate weighting functions for gains and losses also predict a greater deviation from linear weighting of probabilities for gains rather than losses.

Together, the value and probability weighting functions demonstrate the reflection effect: we expect opposite choice behaviour for gains and losses. Consider a choice between a smaller gain for certain and a larger, probabilistic gain, which are equal in expected value. Underweighting at moderate to high probabilities implies that the risky option will be less attractive at these probabilities, which, in combination with risk averse behaviour suggested by the concave value function, predicts that people will be risk averse and choose the smaller gain. The reverse relationship is predicted for losses, where the choice is between a smaller loss for certain and a larger, probabilistic loss. That is, underweighting of moderate to large probabilities makes the risky choice more attractive and, combined with the risk seeking behaviour suggested by the convex value function, predicts that people will be risk seeking and choose the larger, chance loss. A similar reversal is expected at lower probabilities, although the predictions of Prospect Theory are less clear given that the probability weighting predicts changes in behaviour while the value function remains unchanged (see Hershey & Shoemaker, 1980 for a discussion). Overweighting at lower probabilities implies that risky gains will become more attractive and risky losses less, while the value function encourages risk aversion and risk seeking respectively. According to the four-fold pattern of risk attitudes elaborated on by Cumulative Prospect Theory, we would generally expect risk seeking behaviour for gains and risk averse behaviour for losses when probabilities are low (Tversky & Kahneman, 1992).

The exact shape of the value and probability weighting functions has been contested in the literature, but have largely supported the regularities we have plotted above (see Gonzalez & Wu, 1999 for a review). The reflection effect and the predictions it makes about

---

[2]The weighting functions were plotted using a one-parameter model $(p) = \frac{p^{\gamma}}{(p^{\gamma}+(1-p)^{\gamma})^{1/\gamma}}$, where ω is the probability weighing function applied to probability p, and γ is the parameter that specifies the inverse-S shape. In our example, the gain and loss curves were plotted with $\gamma = 0.61$ and γ = 0.69 respectively (Tversky & Kahneman, 1992).

a gain-loss asymmetry appears robust (e.g. Baucells & Villasis, 2010; Schoemaker, 1990). Baucells and Villasis (2010) found that most individuals showed the predicted reflection effect, although only around 50% of individuals displayed both risk averse behaviour for gains and risk seeking behaviour for losses. A significant minority of subjects showed risk averse behaviour for both gains and losses. The size of the reflection effect has also been contested, with Hershey and Schoemaker (1980) concluding that, if present, it was strongest with either small amounts, very large amounts or extreme probabilities.

Gain-loss asymmetry in choice from experience

Prospect Theory was developed based on a descriptive, non-experiential context where individuals are faced with a one-off decision (Kahneman & Tversky, 1979) and its predictions appear to largely hold in that context. Such non-experiential choice tasks where the individual relies on descriptions resemble some, but not all, decisions made in every-day life. An individual might make a one-off decision about whether to undergo surgery based on probabilistic information about the risks and benefits of that decision: a context that Prospect Theory was designed to model. Alternatively, they might make repeated decisions about whether to take a daily medication that reduces their risk of a medical problem, but also creates the risk of side effects, experiencing the outcome of their choices and adjusting which alternatives they subsequently choose accordingly. Prospect Theory was not designed to model this second type of context in which people make repeated decisions from experience.

Research using experiential choice procedures suggests that we should expect different behaviour from humans making decisions from experience as opposed to decisions from description. In one of the earliest demonstrations of this, Barron and Erev (2003) had participants choose between a 100% chance of 9 points (-9 in the loss condition) and a 90% chance of 10 points (-10 in the loss condition; Experiment 3b). Participants did not have descriptions of these probabilities but did receive feedback following each choice about the number of points they had received. The study showed the opposite pattern to what is typically seen in choice from description, with a greater proportion of risky choices for gains (indicative of greater risk seeking) than for losses. Hertwig, Barron, Weber and Erev (2004) found that responding in an experiential context reversed risk preference at both low and high probabilities relative to decisions from description. When the probability of a larger outcome was high and the participants were deciding from description, the majority showed the typical pattern of risk aversion for gains and risk seeking for losses. However, when the same choice was offered in an experiential context, the majority of participants showed risk seeking for gains and risk aversion for losses. Such a reversal was also observed when the probability of the larger outcome was small.

It appears that the relative preference for the risky option, and the direction of the gain-loss asymmetry, differ depending on whether participants experience the outcomes of their choices, which is known as the *description-experience gap* (Hertwig et al., 2004; Hertwig & Erev, 2009). The opposite pattern of behaviour suggested that the probability weighting function for decisions from experience might reverse from that described in Prospect Theory. Studies attempting to determine the shape of the probability weighting function in experiential choice have shown mixed results, an issue further compounded by a lack of consistency in the exact shape of the probability weighting function in decisions from description. As the recent meta-analysis by Wulff, Mergenthaler-Canseco, and Hertwig (2018) noted, the cause of this is largely methodological differences (also see Glöckner, Hilbig, Henninger, & Fiedler, 2016), but some consistencies existed: when participants experienced the outcomes in roughly the same frequency as intended, their choices were in the direction of less overweighing of small probabilities, which ranged from approximately linear to underweighting. We have demonstrated this in Figure 1.3, with arbitrarily chosen values[3] corresponding to an S-shaped function that is overall closer to a linear weighting of probabilities. We also included a distinction between the functions for gains and losses, which has been observed in some studies (e.g. Hau, Pleskac, Kiefer, & Hertwig, 2008; Abdellaoui, L'Haridon, & Paraschiv, 2011) and corresponds to Cumulative Prospect Theory's original prediction of greater deviation from linear probability for gains as compared to losses (Tversky & Kahneman, 1992).

Despite the lack of a universal prediction as to the shape of the probability weighting function, the differences between decisions based on description to that of experience remain fairly consistent. The opposite pattern of responding on the experiential task has been observed at aggregate (e.g. Barron & Erev, 2003) and individual (e.g. Kudryavtsev & Pavlodsky, 2012) levels of comparison, when the experience of number of outcomes vs. number of decisions is controlled for (Camilleri & Newell, 2013) and when the options included no small probabilities (i.e. 50% chance; Ludvig & Spetch, 2011). Wulff and colleagues (2018) concluded that the description-experience gap is a robust phenomenon,

[3] In our example, an S-shaped function was created with $\gamma = 1.39$ for gains and $\gamma = 1.31$ for losses.

with participants being more likely to choose the option in an experiential task that is opposite to what they have chosen in a purely descriptive task in a significant majority of the studies. The description-experience gap was more pronounced when participants chose between a safe and a risky option, rather than between two risky options, although both contexts confirmed its effect. A safe option is one in which an outcome is to occur at 100% probability, and a risky option is one in which an outcome has less than a 100% chance of occurrence.

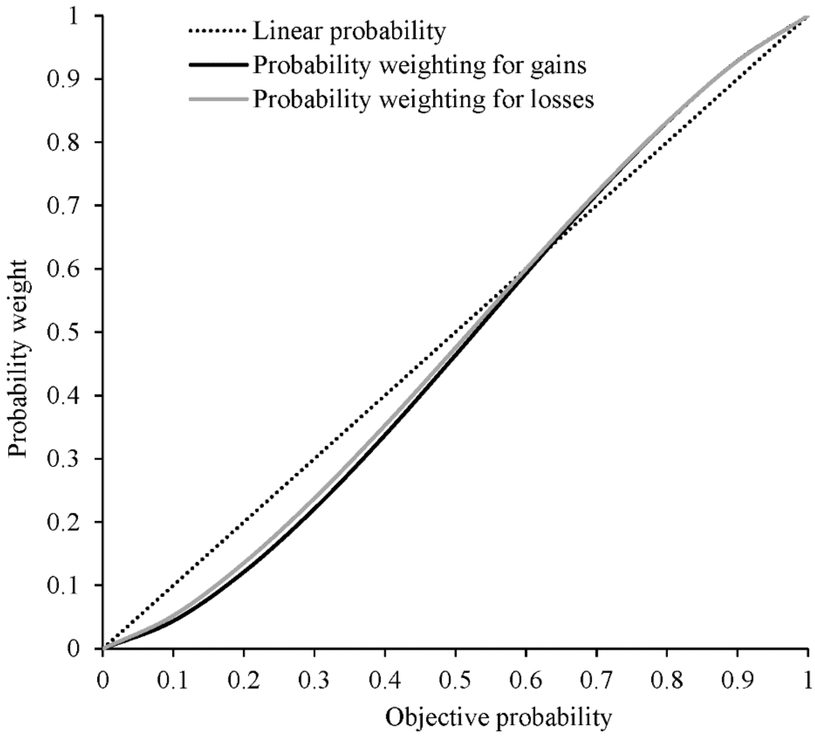

Figure 1.3. Probability weighting function for choice from experience that plots objective probability on the x axis and weighted probability on the y axis. The dotted line corresponds to linear weighting of probabilities. The solid black and grey lines correspond to probability weighting predicted by the description-experience gap literature for gains and losses respectively.

The consistent differences between behaviour based on description and experience emphasizes the necessity of further study into the extent of a gain-loss asymmetry. It would allow us to understand choice about gains and losses across the full variety of everyday decision contexts. Research using experiential tasks also facilitates comparisons to the non-human animal literature by addressing one common procedural difference: animal subjects are only able to experience the scheduled probabilities while human participants can rely on description, experience or both. Differences in preference for the risky option due to experience and its specific procedural components indicates how research from exclusively experiential animal procedures should be applied to human behaviour (see Kalenscher & van Wingerden, 2011 for a discussion). Animal literature has often demonstrated the same violations of expected utility observed with human participants (e.g. Real, 1996; Kacelnik & Bateson, 1996). However, inconsistencies have also been documented, where animal subjects behaved according to predictions of choice from description rather than experience (MacDonald, Kagel, & Battalio, 1991; Shafir, Reich, Tsur, Erev, & Lotem, 2008), suggesting that methodological refinements are necessary before effective translation can occur.

Overview of Current Studies

We aimed to compare the differences between making decisions about gains and losses in two experiential contexts where gain-loss asymmetry has not yet been demonstrated. In Study 1, we examined choice in probability discounting tasks combining description and experience in both safe-risky and risky-risky contexts. Probability discounting quantifies the effect of changing probability on the subjective value of an outcome. There is a lack of discounting studies looking at gain-loss asymmetry, particularly in experiential tasks, but the safe-risky procedures based on description have generally supported the reflection effect modelled by Prospect Theory. Whether the description-experience gap occurs and the reflection effect reverses in probability discounting procedures is unknown. In Study 2, we examined whether a gain-loss asymmetry occurs in a concurrent schedules task, which has also been used to examine the effect of changes in probability on preference among two alternatives. In Study 3, we addressed follow-up questions raised in Studies 1 and 2 on the effects of magnitude and cross-task experience.

Study 1: Gain-loss asymmetry in probability discounting

We often encounter choices among outcomes differing across multiple dimensions, such as probability and amount. For example, when deciding whether to purchase insurance, one alternative - buying insurance - requires a small, guaranteed payment while the other - foregoing insurance - might produce a larger loss in the future. Most individuals prefer guaranteed rather than chance gains, and chance rather than guaranteed losses. That is, the subjective value of an outcome is discounted with decreasing chance of its occurrence (probability discounting). We also prefer larger outcomes if positively valued (i.e. gains or rewards) and smaller outcomes if negatively valued (i.e. losses or punishers). However, in many decision-making situations, such as whether to purchase insurance, one alternative produces an outcome of a more preferred amount (e.g. a certain but smaller loss) and the other alternative produces an outcome of a more preferred probability (e.g. a larger loss that may not occur), and the person deciding must trade-off probability and amount.

Probability discounting

Probability-amount trade-offs can be studied within the probability discounting paradigm. A discounting perspective assumes that the subjective value of an outcome systematically decreases (i.e. is discounted) as the probability of that outcome occurring decreases. This approach to studying choice has been used to investigate risk-taking in gambling (e.g. Holt, Green, & Myerson, 2003), differences between smokers and non-smokers (e.g. Lawyer, Schoepflin, Green, & Jenks, 2011) and willingness to risk medical treatment (e.g. Weatherly & Terrell, 2014; Asgarova, Macaskill, Robinson, & Hunt, 2017).

Participants typically make a series of two-alternative choices and do not receive feedback about the outcomes of their choices (i.e. choice is made from description and is non-experiential). Table 2.1 presents possible alternatives within gains (top) and losses (bottom) tasks. Researchers often use the *titrating amount procedure* (Richards, Mitchell, de Wit, & Seiden, 1997). In this procedure, a participant makes a series of choices between a smaller, certain amount and a larger, uncertain amount. After each choice, the smaller, certain amount is either decreased or increased in response to what was chosen on the previous trial, and this adjustment is intended to make the previously selected alternative less appealing. Such adjustments are done over several trials in order to identify when the two alternatives are roughly equal in subjective value: this is termed the *indifference point.* For example, an individual might be indifferent between a 50% chance of gaining $100 and gaining $40 for certain. In this example, when the gain of $100 only has a 50% chance of occurring, it lost $60 of its value with decreased probability. In a losses condition, an individual might be indifferent between losing $60 dollars for certain and a 50% probability of losing $100, indicating that $100 has lost $40 of its subjective value with decreased probability. Note that in this example, a probabilistic gain has lost more of its subjective value than a probabilistic loss of the same probability.

Table 2.1

Examples of probability-amount trade-offs in probability discounting tasks

	Smaller certain	Larger uncertain
Gains		
E.g. of alternatives	Gain $50 for certain	50% chance of gaining $100
Trade-off	Better on **certainty** (guaranteed gain)	Better on **amount** (larger gain)
Losses		
E.g. of alternatives	Lose $50 for certain	50% chance of losing $100
Trade-off	Better on **amount** (smaller loss)	Better on **certainty** (uncertain loss)

Quantifying probability discounting

Using titrating amount procedures, researchers can derive indifference points for several probabilities of gaining or losing a given larger, uncertain amount in order to quantify how the subjective value of the larger outcome changes as a function of changes in its probability. The rate at which the subjective value of the larger outcome changes with decreasing probability can be described using one of several mathematical functions. Here we will focus on two methods, the hyperbolic model and calculating area under the curve.

The hyperbolic model. The hyperbolic model (Equation 2.1) assumes the subjective value of an outcome will decrease more steeply with change at higher probabilities and more shallowly with change at lower probabilities (Mazur, 1987; Rachlin, Raineri, & Cross, 1991).

$$V = \frac{A}{(1+h\Theta)} \qquad \text{(Equation 2.1)}$$

In this model, the subjective value of the larger, uncertain outcome (*V*) is derived from its undiscounted size (*A*), odds against its successful occurrence (*Θ*), and a free parameter *h*. The *h* parameter best describes change in value of the larger, uncertain outcome as a function of changes to its probability for a given individual. Odds against is computed as

(1-probability)/probability and indicates the expected mean number of non-occurrences of an event before it occurs. For example, a 30% chance of losing $100 is equivalent to an odds against losing $100 of 2.3, which is, on average, one loss for every 2.3 non-occurrences of a loss. This conversion of probabilities to odds against allows for a hyperbolic shape of the discounting function. The left panel of Figure 2.1 demonstrates how *h* describes different rates at which an outcome can lose its value. Two sets of hypothetical indifference points are plotted, where black triangles correspond to indifference points for losses and white triangles correspond to indifference points for gains. Fitting Equation 2.1 to the data will produce the solid and dashed curves, respectively. In this example, the losses curve has a larger *h* parameter than the gain curve. The larger the *h* parameter, the steeper the rate at which the larger, uncertain outcome loses its value with decreasing probability. Conversely, the smaller the *h* parameter the better the larger, uncertain outcome holds its value with decreasing probability of its occurrence.

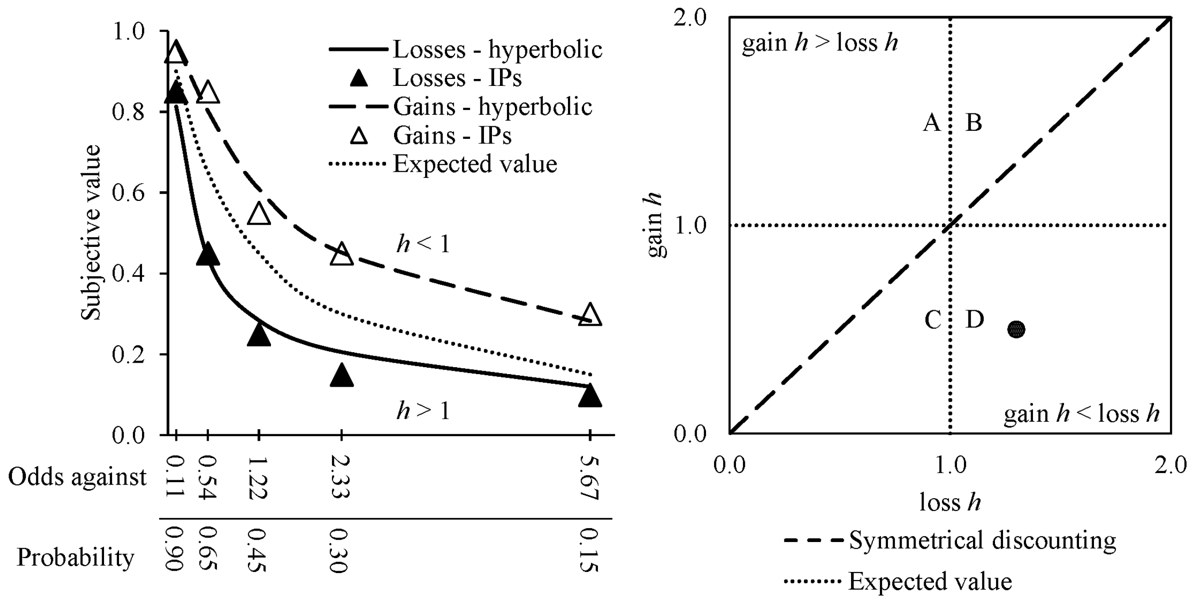

Figure 2.1. Left panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) using hypothetical data. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Right panel: The *h* parameter for gains plotted against the *h* parameters for losses for the hypothetical data in the left panel. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate *h* value (1) when decisions are made based on expected value.

The relationship between the h parameters can also be shown on a modified Brinley plot (Blampied, 2017) to examine the consistency of the discounting of gains and losses across individuals, where loss h is plotted on the x axis and gain h on the y axis. The data in the left panel of Figure 2.1 is represented as a single data point in the right panel, describing h parameters for both gains and losses. The dashed diagonal reference line indicates identical discounting of gains and losses, where loss h = gain h. Points above the line indicate steeper discounting of gains than losses, whereas points below the line indicate steeper discounting of losses than gains. The steeper discounting of losses than gains in the left panel is represented by a data point below the diagonal line on the right panel.

Discounting can also be compared to behaviour predicted solely by the expected value of options and corresponds to choices that would maximize gains or minimize losses. This is also of interest as it allows us to determine whether choice was overall risk seeking or risk averse, and thus subject to probability weighting predicted by the Prospect Theory. In a probability discounting task, choices based on the expected value would result in $h = 1$ and is indicated by the dotted line in the left panel of Figure 2.1. When $h = 1$, Equation 2.1 determines subjective value (V) solely from odds against (Θ) and its undiscounted size (A). A decision based on expected value alone can be considered risk neutral, where the decrease in probability corresponds to a proportionate decrease in subjective value (see Shead & Hodgins, 2009 for a discussion). Responses that deviate from this would indicate change to subjective value that is higher or lower than the expected value alone would predict (thus requiring a scaling parameter, h, not equal to 1). When $h < 1$ and the discounting curve is showing a shallower decrease in value, this is indicative of risk seeking for gains and risk aversion for losses, where the larger outcome holds its value better than predicted by the expected value. When $h > 1$ and the discounting curve is showing a steeper decrease in value, this is indicative of risk aversion for gains and risk seeking for losses, where the larger outcome holds its value worse than predicted by the expected value.

The relationship to expected value can also be represented in the right panel, where $h = 1$ is indicated by the dotted vertical and horizontal lines, allowing the data to be further split into four quadrants. Data points in quadrant B would indicate discounting of both gains and losses that is steeper than predicted by expected value ($h > 1$), whereas in quadrant C it is shallower ($h < 1$). Our sample data on the left graph shows discounting of losses that is steeper than predicted by expected value and discounting of gains that is shallower, hence the data point in the right panel is in the D quadrant.

The hyperbolic model generally describes data well, with the *h* parameter capturing differences across individuals in discounting of different commodities, such as real money (Lawyer et al., 2011) and hypothetical or real cigarettes (Green & Lawyer, 2014). However, the hyperbolic model is one of many mathematical functions that have been used in literature, with each using a different combination of parameters to predict the rate at which an outcome ought to lose its value. Determining the model with the best fit to the data allows for more accurate predictions on how an outcome will be impacted by changing its probability. The single parameter hyperbolic model generally describes data better than a single parameter exponential model, which assumes a constant rate at which an outcome loses its value (see Green & Myerson, 2004 for a review). However, the comparison of the hyperbolic one-parameter to hyperboloid two-parameter models is more nuanced. Two such models have gained prominence in literature, the Green and Myerson (2004; Equation 2.2) and the Rachlin (2006; Equation 2.3) equations.

$$V = \frac{A}{(1+h\Theta)^s} \qquad \text{(Equation 2.2; Green \& Myerson, 2004)}$$

$$V = \frac{A}{(1+h\Theta^s)} \qquad \text{(Equation 2.3; Rachlin, 2006)}$$

The hyperboloid models use two scaling parameters, *h* and *s*, where *s* is a non-linear scaling parameter proposed to indicate non-linear sensitivity of subjective value to odds against (Rachlin, 2006). This greater number of parameters improves the models' flexibility in fitting to different data sets and generally produces higher indicators of goodness of fit, but also raises issues of coherence in interpreting variation in the two parameters (Green & Myerson, 2004). Both *h* and *s* impact the shape of the discounting curve, are not independent and are often correlated, more so in Equation 2.2 where *h* is directly raised to the power of *s* (see Mitchell, Wilson & Karalunas, 2015, and Young, 2017 for a discussion in a delay discounting context). This non-independence presents a challenge to the models as it introduces issues of multicollinearity between the predictors. Furthermore, a comparison based on the *h* parameter across the two groups of interest (e.g. different magnitude of the larger, uncertain option) would only be meaningful if *s* is equivalent across the two groups and independent of *h*, which has not been consistently demonstrated in probability discounting (e.g. Estle, Green, Myerson, & Holt, 2006). Furthermore, there is no current consensus on interpreting the difference between *h* and *s* parameters, and whether restrictions should be placed on any variation in them when fitting the models in order to permit comparisons across groups (McKerchar, Green, & Myerson, 2010). We therefore opted to use the hyperbolic model in our subsequent analysis.

Area under the curve. A method that avoids the model fitting issues described above is the area under the curve (AUC) measure of discounting (Myerson, Green, & Warusawitharana, 2001). Here, the rate of discounting is calculated by dividing the area under the plotted indifference points into trapezoids as shown by the dashed lines in the left panel of Figure 2.2, applying Equation 2.4 to each trapezoid and summing them to produce a single value.

$$Trapezoid\ area = (x_2 - x_1) * ((y_1 + y_2)/2)) \qquad \text{(Equation 2.4)}$$

x_1 and x_2 are successive odds against and y_1 and y_2 are successive subjective values, all expressed as a proportion of their maximum possible value (i.e. normalized to range from 0 to 1). The resultant value ranges from 0 to 1 and is the area under the plotted data relative to the maximum possible area, which is indicated by the square ending at the maximum possible x and y axis values. The AUC for losses would be derived from the total grey area and for gains from the total dotted area. Higher AUC values correspond to a greater proportion of the shaded area relative to total area and indicates shallower discounting, while lower AUC values correspond to a lesser proportion of the shaded area relative to total area and indicates steeper discounting. Note that this is in the reverse direction to the *h* parameter, where a higher *h* corresponds to a lower AUC and vice versa.

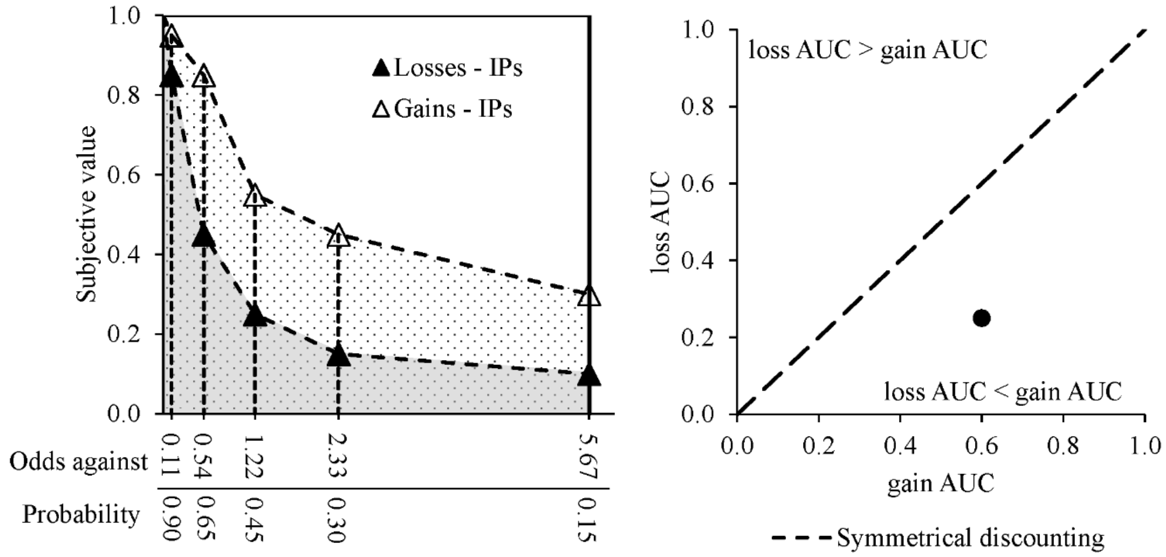

Figure 2.2. Left panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white

triangles) or loss (black triangles) using hypothetical data. Dashed lines indicate the trapezoid areas used to calculate the area under the curve value. The shading corresponds to the total area used for gains (dotted) and losses (grey). Right panel: The AUC value for losses plotted against the AUC value for gains for the hypothetical data in the left panel. The diagonal dashed line represents symmetrical discounting of gains and losses.

As with the *h* parameter, individual data can be plotted on a modified Brinley plot on the right panel of Figure 2.2. Note that for the AUC graph, we have chosen to swap the axis data relative to the *h* graph, such that gains are plotted on the x axis and losses on the y axis, to facilitate comparison relative to the dashed diagonal line. As with the modified Brinley plot of *h* values described above, points above the line would indicate steeper discounting of gains than losses and points below the line indicate steeper discounting of losses than gains. The steeper discounting of losses than gains in the left panel is represented by a data point below the diagonal line on the right panel.

We opted to supplement our analysis using the *h* parameter with AUC because each approach has some advantages and some disadvantages. While AUC avoids model fitting issues, its values do not indicate whether systematic discounting has occurred (i.e. a monotonic decrease in value with increasing odds against). Furthermore, AUC values cannot be compared to an AUC value derived from choice based on expected value (for example, the dotted line in the left panel of Figure 2.1 produces an AUC of approximately 0.33), since the AUC value corresponding to discounting based on expected value can correspond to different distributions of indifference points on the graph. On the other hand, AUC values provide a more accurate description in instances where discounting data are poorly described by the hyperbolic model. Furthermore, some studies have found different patterns in the data when using AUC and *h* (e.g. Weatherly & Derenne, 2013). In all of our analysis, we used both AUC and *h*, and applied the following logic. Where examinations of any systematic patterns in residuals and/or low R^2 values (proportion of the variance in the data explained by the fitted model) indicate that the hyperbolic model does not describe the data well, then any effect captured by AUC is more likely to be reliable. Any discrepancies between measures derived from a well-fitting model and AUC would most likely suggest that AUC is capturing noise in the distribution of indifference points around an otherwise monotonic trend, and the *h* results become more reliable.

Discounting in a non-experiential, safe-risky context

In studies using non-experiential probability discounting tasks, a common finding is that gains are discounted more steeply than losses (Estle et al., 2006; Mitchell & Wilson, 2010; Weatherly & Derenne, 2013; Shead & Hodgins, 2009). That is, the rate at which a probabilistic gain loses its subjective value in response to decreasing probability is steeper than that of a probabilistic loss, thus demonstrating an asymmetrical effect of probability on discounting of gains versus losses. The gain-loss asymmetry observed in probability discounting is thus consistent with the reflection effect described by Prospect Theory (Kahneman & Tversky, 1979), showing that people do not value equal sized losses and gains symmetrically.

While this asymmetry is commonly reported, the extent of it appears to be dependent on the specific context. The rate of discounting is affected by the magnitude of the discounted outcome and this effect is not symmetrical for gains and losses. Research suggests that larger gains are discounted more steeply than smaller gains (e.g. Green, Myerson & Ostaszewski, 1999; Myerson, Green, Hanson, Holt & Estle, 2003; see also Weatherly & Derenne, 2013 for non-significant results), but the magnitude of losses does not affect discounting rate to the same extent (Estle et al., 2006; Mitchell & Wilson, 2010; Green, Myerson, Oliveira & Chang, 2014).

This difference in the effect of magnitude results in a more pronounced gain-loss asymmetry at higher magnitudes, but the range of lower magnitude at which no gain-loss asymmetry is observed is inconsistent across studies. An additional complication is that statistical tests conducted on parameters derived from model fits versus AUC are not always consistent. Weatherly and Derenne (2013) reported no main effect of condition using log-transformed *h* from the hyperbolic model (with generally poor fits ranging from R^2 of 0.70-0.72 for gains and 0.58-0.59 for losses), but a significant main effect using AUC, where gains were discounted more than losses when data were averaged across the $1000 and $100000 magnitudes. The effect of magnitude itself was not significant using either measure of discounting. The average AUC reported by Mitchell and Wilson (2010) showed that gains were discounted more steeply than losses for both $10 and $100 magnitudes. Estle et al. (2006) reported analysis using AUC, with no significant main effect of condition when participants discounted $200 and $40000 (Experiment 2) and a significant main effect of condition when participants discounted $100, $20000 and $60000 (Experiment 4).

In non-experiential probability discounting tasks, gain-loss asymmetry has sometimes also been observed in studies that assessed whether choice is risk seeking or risk averse relative to the expected value of the outcomes. The four-fold pattern of risk preference described by Prospect Theory predicts that people are risk averse for gains and risk seeking for losses for most probabilities (Tversky & Kahneman, 1992). However, at lower probabilities individuals are expected to become more risk seeking for gains and risk averse for losses. In probability discounting, we are not able to use *h* values to examine responses to specific probabilities, as it quantifies rate of change across a range of probabilities. AUC is similarly a summary statistic about all of the data points, rather than individual probabilities. We can, however, determine whether overall the subjective value is lost at a steeper or shallower rate than expected value would predict. In probability discounting, the predicted pattern of risk aversion for gains and risk seeking for losses at most probabilities would correspond to generally steeper discounting of both gains and losses relative to the expected value rate.

Studies comparing observed discounting curves for gains and losses to those based on the expected value have produced mixed results. Shead and Hodgins (2009) found that the rate at which an outcome lost its value was, in general, risk averse for gains and risk seeking for losses relative to the expected value. Estle et al. (2006) did not report an analysis comparing discounting to the expected value. My examination of their data showed that discounting of gains depended on the magnitude of the larger outcome. That is, discounting was closer to behaviour predicted by the expected value with lower amounts and more risk averse with higher amounts. Conversely, discounting of losses was unaffected by magnitude and was closer to discounting predicted by the expected value. Mitchell and Wilson (2010) similarly observed that discounting rate for gains depended on the magnitude of the larger outcome, showing a similar pattern to Estle et al. (2006), but the loss condition showed an overall pattern of risk aversion. Weatherly and Derenne (2013) did not report a comparison to the expected value, but my examination of the data suggested that choice was, in general, risk seeking for gains and risk averse for losses.

The most consistent pattern across the studies we examined is that for both gains and losses, the indifference points were shallower than the expected value at probabilities lower than approximately 25% (higher odds against; e.g. Mitchell & Wilson, 2010), corresponding to Prospect Theory's probability weighting function predictions. The variability observed above appears to be mostly due to data at higher probabilities, both across studies and between gains and losses. At probabilities higher than approximately 75% (lower odds against), discounting tended to be steeper than the expected value and match Prospect Theory predictions, but more consistently for gains rather than losses (e.g. Shead & Hodgins, 2009). Discounting at the middle range of probabilities (25-75%) showed a less consistent pattern of responding for both gains and losses. This result may be due to encompassing the probability range at which Prospect Theory predicts a change from underweighting to overweighting of probabilities according to the probability weighting function. No clear cut-offs exist for where underweighting of low probabilities begins, with the most common observation is that it lies between 0.30 and 0.40 probabilities (Tversky & Kahneman, 1992; Wu & Gonzalez, 1996; Gonzalez & Wu, 1999).

Overall, if a gain-loss asymmetry was observed, it tended to be in the direction of steeper discounting of gains than losses, consistent with the reflection effect in Prospect Theory. It also does not appear to be a strong effect, which is consistent with some of the reflection effect literature that did not find a robust difference (e.g. Hershey & Schoemaker, 1980). The distribution of indifference points relative to the expected value tended to match Prospect Theory predictions, more so for gains than losses.

Discounting in experiential and risky-risky contexts

According to the description-experience gap literature, choice from experience results in a different preference for the risky option compared to choice from description. Furthermore, the differences to choice from description are more pronounced when the participant is deciding between a safe and a risky option, rather than two risky options.

Experience of outcomes. In probability discounting tasks combining experience and description, participants receive feedback about each choice before making subsequent choices. Participants on these tasks have shown systematic discounting of both hypothetical (Greenhow, Hunt, Macaskill & Harper, 2015) and real monetary gains (Scheres et al., 2006; Hinvest & Anderson, 2010). Greenhow et al. (2015) developed a mixed description-experience game-style task where participants navigated a skier down a slope. Participants also regularly encountered a ski jump scenario with a choice between a smaller number of points to be gained for certain and a larger number of points to be gained at one of five probabilities. The key aspect was that participants were exposed to the outcomes of each trial based on their choices, with immediate feedback on their score. The procedure demonstrated systematic discounting of hypothetical game points, suggesting that experiential computerized tasks are a viable procedure for comparing probability discounting of gains and losses when outcomes of choices are experienced. No studies have examined gain-loss asymmetry in such mixed description-experience tasks to our knowledge.

We examined gain-loss asymmetry in the mixed description-experience ski task by contrasting it to a non-experiential money task as part of my honours thesis (unpublished). In a first experiment, we observed that the gain-loss asymmetry reversed as predicted by the description-experience gap, with steeper discounting of gains than losses in the money task and steeper discounting of losses than gains in the ski task (graphs reprinted in Appendix A). In a second experiment, we replicated this comparison of discounting of gains and losses from the first experiment, and introduced experiential features to the money task in order to test whether a similar reversal of the gain-loss asymmetry would be observed. However, we did not replicate this reversed pattern in either the mixed description-experience money or ski game contexts; group data showed similar discounting of gains and losses and individual data suggested this was due to a large range of variability in whether gains or losses were discounted steeper (graphs reprinted in Appendix A). While the data from the mixed description-experience tasks did not resemble that of discounting in a non-experiential context, neither did it show discounting consistent with the predictions of a reversed reflection effect. Therefore, Study 1 continued the examination of the effect of introducing experiential outcomes on the gain-loss asymmetry in discounting, as well as included an examination of discounting relative to choice predicted by the expected value.

Notably, our probability discounting procedures combined both description and experience, which is unlike most of the procedures that have examined the description-experience gap. In a probability discounting task, given the number of trials required for a titrating amount procedure, it is impractical to provide sufficient trials for participants to fully experience the probabilities examined. We adopted a combined description-experience condition, where the probabilities were described, but the participants also received experience with most of the outcomes. Research that has provided participants with both descriptions of and experience with the probabilities of choice outcomes suggests that experience predominates, and decisions resemble those made from experience alone (Jessup, Bishara, & Busemeyer, 2008; Lejarraga & Gonzalez, 2011). This combination of description and experience is also analogous to our example of taking daily medication in the general introduction. An individual might make repeated decisions about whether to take a daily medication that they have been told reduces their risk of a medical problem while also describing a risk of side effects, experiencing the outcome of their choices and adjusting subsequent choices accordingly. For simplicity, we will refer to the mixed description-experience task as experiential in subsequent discussions, but we will return to this distinction in the General Discussion.

Based on the description-experience gap literature and the reversed reflection effect, we expected losses to be discounted more steeply than gains in an experiential task (Hertwig & Erev, 2009). Furthermore, as we noted in the general introduction, the probability weighting function for experiential choice tends to range from almost linear weighting to an S-shape. Wulff et al. (2018) observed that behaviour in the experiential context showed a higher proportion of choices maximizing the mean returns than in the descriptive context: experience was associated with more optimal choice. Therefore, when participants experience the outcomes of their choices, we would expect either shallower discounting of both gains and losses relative to choice based on the expected value of the outcomes, or at least choice patterns that are in general closer to the expected value.

Risky-risky context. Typical discounting tasks establish the subjective value of the larger, uncertain alternative by contrasting it to a smaller amount to receive for certain. These form a useful analogue of the common premise of having to trade-off two valued dimensions, certainty and amount. Under such conditions, the predictions of Prospect Theory in terms of a reflection effect have largely been supported. We can modify a typical safe-risky task by changing the probability of the smaller, certain alternative to 0.99. Mathematically, this alternative is almost certain to happen and practically the participants are unlikely to experience the 0.01 rare outcome. The trade-off largely remains unchanged, with one option offering (more) certainty and the other a higher amount, but based on the behaviour economic literature we might expect a disproportionate change in behaviour.

According to the probability weighting function of Prospect Theory, a numerical decrease in probability is not always accompanied by a proportionate decrease in preference (refer to Figure 1.2; Tversky & Kahneman, 1992). Rather, individuals treat the transition from certainty to near certainty as categorical rather than continuous, whereas a similar change among moderate probabilities (e.g. from 45% to 44%) is more linear in nature. However, under conditions where both outcomes are risky, behaviour more closely approximates that predicted by the expected value of the outcomes (see Andreoni & Sprenger, 2011 and Bleichrodt et al., 2007 for discussions). For example, Bleichrodt et al.

(2007) observed no overweighting of small probabilities between 0.1 and 0.2 in a risky-risky context using health outcomes. Additionally, McCord and Neufville (1987) observed that choice between two risky gains did not result in overvaluing of outcomes characterized by the concave shape of the Prospect Theory value function. Therefore, we might expect a risky-risky non-experiential task to show discounting closer to expected value, consistent with some studies that have shown that Prospect Theory predictions do not hold well for choice between risky options, although reflection effects have also been documented in a risky-risky context (e.g. Budescu & Weiss, 1987). Probability discounting using two risky options has not been investigated to our knowledge.

Risky-risky and experience combined. Wulff et al. (2018) concluded that the description-experience gap was reduced, but not eliminated, when both alternatives were risky. This difference appeared to be partially due to the experiential choice in the safe-risky contexts resulting in linear to S-shaped weighting of probability, different from that of descriptive choice, while risky-risky procedures resulted in an inverse S-shaped function akin to descriptive choice. Therefore, we also examined experiential discounting in a risky-risky context for patterns that might deviate from experiential choice in a safe-risky context.

Correlations

In addition to comparing absolute levels of discounting of gains and losses, we also compared whether the tendency to choose the risky option was maintained within individuals across the tested contexts. Studies that have correlated discounting of gains and losses have observed negative correlations (Shead, Callan, & Hodgins, 2008; Shead & Hodgins, 2009), although this was not consistently significant (Mitchell & Wilson, 2010). Studies that have compared discounting of gains across tasks have generally observed positive correlations when the conditions varied in magnitude of the larger outcome (Greenhow et al., 2015; but see Yi, de la Piedad, & Bickel, 2006 for non-significant results), when the rewards were real or hypothetical (Hinvest & Anderson 2010; Matusiewicz, Carter, Landes, & Yi, 2013) and when participants were tested after a delay between condition (Peters & Buchel, 2009; Ohmura, Takahashi, Kitamura, & Wehr, 2006; but see Matusiewicz et al., 2013 for non-significant results). Similar positive correlations have been observed for losses when the conditions varied in magnitude of the larger outcome (Green et al., 2014). Overall, the correlations observed in the literature suggested that greater risk seeking in one condition was associated with greater risk seeking in the other, and it is of interest whether such a relationship can also be observed in the atypical contexts described above.

Table 2.2 summarizes our specific predictions based on Prospect Theory and description-experience gap literature.

Table 2.2

Predictions of gain-loss asymmetry for the four possible context combinations in a probability discounting task

	Description	Experience
Safe-risky	Discounting of gains steeper than of losses, both steeper than expected value-based curve	Discounting of losses steeper than of gains, both shallower than/close to expected value-based curve
Risky-risky	Discounting of gains steeper than of losses, both steeper than/close to expected value-based curve	Discounting of losses steeper than of gains, both steeper than/close to expected value-based curve

Experiment 1: Discounting in a risky-risky, non-experiential money task

Study 1, Experiment 1 aimed to pilot an online risky-risky discounting procedure, where the probability of the smaller, certain option was changed from 1.00 to 0.99, and to contrast it to discounting on a safe-risky money task. Our first aim was to determine the feasibility of this contextual change to an online format in terms of the rate of unsystematic data and the fit of the hyperbolic model to the data. These indicators were compared to the literature and my honours thesis data set. We piloted online data collection because this approach would facilitate more efficient and convenient data collection from a more diverse participant group. If this experiment suggested that online data collection produced data of a quality comparable to in-lab data collection, then we planned to use this approach for subsequent studies in this thesis. We chose a non-experiential task to pilot online because more is known about typical rates of systematic data and hyperbolic model fits in such contexts.

Our second aim was to examine gain-loss asymmetry in the safe-risky task; we expected steeper discounting of gains than losses (e.g. Estle et al., 2006) and discounting curves to be steeper than curves produced by expected value. Our third aim was to examine any consistent patterns in discounting of gains relative to losses, as well discounting relative to the expected value, which might indicate a gain-loss asymmetry in the risky-risky context. Based on the literature, we expected to observe steeper discounting of gains than losses and choice patterns that were closer to decisions based on the expected value. Our analysis was also supplemented by examining correlations within and across tasks.

Method

Participants

Forty-four participants attending Victoria University of Wellington were recruited through the School of Psychology Research Programme online tool and participated in partial fulfilment of a course requirement. The Victoria University School of Psychology Human Ethics committee reviewed and approved all aspects of all studies presented in this thesis (approval number: 0000024336).

Materials

Non-experiential safe-risky money task. The task was coded using JavaScript in Qualtrics. Participants completed a total of 30 questions for each condition. The following instructions were shown with each set of questions:

> *In this task, you will see pairs of monetary outcomes. Please click on the outcome you would choose if you were given this choice in real life. Please read each pair carefully as each will be slightly different from the previous pair.*

Below the instructions, participants were presented with five questions to a screen and had to scroll down to complete them. Each question presented them with two alternatives with the smaller amount at 1.00 probability on the left and the larger amount at one of five probabilities (0.90, 0.65, 0.45, 0.30, and 0.15) on the right. Each set of five questions contained one question for each probability of the larger amount. The order in which these questions were listed on the screen was the same in each set and was randomized once at the start of each condition for each participant. The smaller, certain amount in the first set of five questions in each condition was set to $50, half of the larger, uncertain amount ($100). For each subsequent choice at the same probability, the certain amount was adjusted by 10% of the larger amount ($10) based on the participant’s previous choice at that probability. For gains, choosing the larger, uncertain gain resulted in an increase in the smaller amount for that probability in the next set while choosing the smaller gain resulted in a decrease. For losses, choosing the uncertain loss resulted in a decrease in the size of the smaller loss in the next set, while choosing the smaller loss resulted in an increase. Thus the size of the smaller alternative was always adjusted in the direction that would make the participants' previous choice at that probability less appealing. Participants clicked a button to advance to the next set of five questions and, once all 30 questions for the current condition were completed, to the next condition or task.

Non-experiential risky-risky money task. The task was coded as described above, with the exception of the smaller amount set to 0.99 probability in each trial. The titrating procedure was as described above.

Procedure

The experiment was completed online via the survey engine Qualtrics. Participants were shown a screen with a brief summary of the experiment, indicating that they had the option to withdraw at any stage before completion of the survey, and an informed consent statement. Each participant completed both tasks and both conditions within each task. The order of the tasks and conditions was randomized. Upon completion, participants were shown a screen debriefing them on the purpose of the study.

Data Analysis and Exclusions

In the safe-risky task, we calculated the indifference point at each probability of the larger outcome, which indicated a smaller, certain outcome that was equal in subjective value to the larger, uncertain outcome. We took the mean of the smaller, certain amount at the last trial and the subsequent smaller, certain amount that would have been presented following one additional adjustment. For example, if the participant chose the smaller, certain option of $60 at the last gain trial, the subsequent smaller, certain amount would have been $50 and so the indifference point for this probability would be $55. The only exception was that the indifference points could not be below $0 or above $100. Indifference points for the risky-risky task were calculated in the same manner, and indicated a smaller outcome to occur with 0.99 probability that was equal in subjective value to the larger outcome to occur at a lower probability. Indifference points were expressed as a proportion of the larger, uncertain amount.

We examined participants' indifference points for non-systematic discounting using the widely used criteria developed by Johnson and Bickel (2008; see Smith, Lawyer & Swift, 2018 for a review on non-systematic data in discounting). According to the first criterion, data were systematic if an indifference point was not larger than the indifference point at the previous odds against by more than 20% of the uncertain amount ($20). According to the second criterion, the indifference point for the largest odds against had to be smaller than the first indifference point by at least 10% ($10) of the uncertain amount. This was based on the assumption that an outcome should lose at least some of its value (quantified as 10%) with decreasing probability. Table 2.3 specifies the number of participants who failed the criteria within each condition. Non-systematic data were assumed to reflect inattention or lack of understanding of the options presented and we opted to present the analysis with unsystematic data sets removed. However, including all participants did not affect conclusions. When the analysis concerned only the money or only the ski task data, we used the fully systematic responses in each task (row 3 in Table 2.3). When the analysis concerned comparisons across tasks, we used the fully systematic responses across tasks (row 4 in Table 2.3).

Table 2.3

Number of participants (n = 44) in Study 1, Experiment 1 who had unsystematic data by criterion, and total participants with systematic data.

	Safe-risky		Risky-risky	
	Gain	Loss	Gain	Loss
Criterion 1	1	2	1	4
Criterion 2	4	5	1	6
Total systematic in each task	36 (81.81)		35 (79.55)	
Total systematic across tasks	30 (68.18)			

Note. Percentages are in parentheses.

The hyperbolic equation was fitted to each individual's indifference points and to the group median indifference points using nonlinear, least squares regression. R^2 values were calculated in order to estimate the goodness of fit for the hyperbolic equation. The h parameters were not normally distributed (Shapiro-Wilk's test for normality all $p < .001$) and were log transformed to normalize the distribution for analysis. Log transforming the data resulted in Shapiro-Wilk's test p values > .05 for all but the losses condition on the safe-risky task ($W = 0.93$, $p = .027$). Examination of histograms showed a slight positive skew for losses, but the shape of the distribution was otherwise normal. AUC values were normally distributed (Shapiro-Wilk's test for normality all $p > .05$) with the exception of gain condition on risky-risky task ($W = 0.92$, $p = .016$). Examination of histograms showed a slight negative kurtosis for gains, but the shape of the distribution was otherwise normal. Thus, statistical tests that require normality assumptions were deemed appropriate. Effect sizes for *t* tests (Cohen's d_z) were calculated using the spreadsheet provided by Lakens (2013) and for *F* tests (partial eta squared, η_p^2) were derived from SPSS output.

Results and Discussion

Did the online format produce reliable estimates of discounting?

Application of Johnson and Bickel (2008) criteria produced a greater number of unsystematic data sets than expected, suggesting poor data quality. The percentage of data sets that were not fully systematic (31.82%) was higher than the overall frequency seen in probability discounting studies (19%; Smith, Lawyer, & Swift, 2018) and my honours thesis data. Notably, the percentage of unsystematic data sets for the safe-risky (18.18%) and risky-risky (20.45%) tasks was comparable to literature, showing that the risky-risky manipulation itself resulted in largely sensible data. Examination of Table 2.3 showed that the majority of criteria violations were on criterion two, where the participants showed no sensitivity to changes in probability from highest to lowest odds against. The resultant systematic data sets produced reasonable R^2 values for the group data and the majority of the individuals' data (see Table 2.4). Although we observed a greater range of individual R^2 values in the current sample, the median R^2 were comparable to my honours thesis experiments and the literature (e.g. Lawyer et al., 2011).

Table 2.4

Study 1, Experiment 1 money task h, R^2 and AUC values

	Group median indifference points			Individual participants		
	h	R^2	*AUC*	*h* median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Safe-risky gain	2.79	0.98	0.21	2.61 (1.46, 4.97)	0.88 (0.69, 0.95)	0.23 (0.02)
Safe-risky loss	1.56	0.96	0.30	1.59 (0.87, 4.22)	0.80 (0.36, 0.90)	0.30 (0.03)
Risky-risky gain	2.84	0.97	0.19	2.39 (1.29, 4.19)	0.87 (0.80, 0.93)	0.22 (0.02)
Risky-risky loss	2.38	0.97	0.23	1.87 (1.01, 5.47)	0.79 (0.48, 0.89)	0.27 (0.03)

Note. Q_1 = first quartile, Q_3 = third quartile.

In order to check whether task progression resulted in poorer data quality due to fatigue or boredom, we calculated the percentage of unsystematic responses and median R^2 values from 1st to 4th condition completed by the participants. The percentage of unsystematic responses was lowest in the 1st condition (2.27%) and highest in the second (18.18%), subsequently decreasing for the 3rd (15.91%) and 4th (9.09%) conditions, suggesting that the relatively high percentage of unsystematic data observed did not reflect participants becoming fatigued as a result of the relatively large number of choices they were presented with. This also showed that the majority of unsystematic data were not from the early conditions, suggesting that participants understood the instructions. Median R^2 values for systematic data were stable across conditions: 1st 0.86, 2nd 0.78, 3rd 0.89, and 4th 0.86.

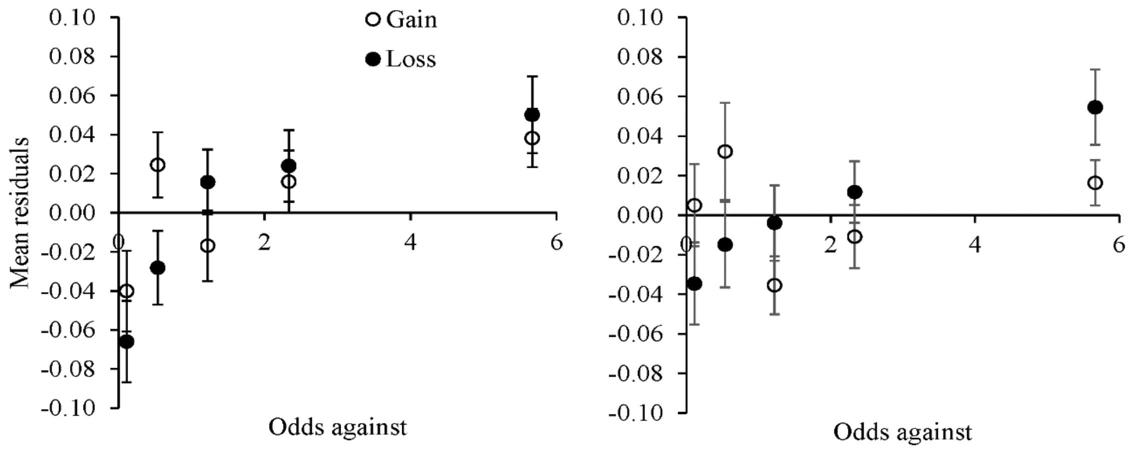

Figure 2.3. Mean residuals calculated from individual data on the safe-risky (left) and risky-risky (right) money tasks as a function of odds against in the gain (white circles) and loss (black circles) conditions. Error bars are standard error of the mean.

Lastly, examination of residuals showed systematic effects of odds against on residuals in both tasks, with a linear increase in mean residuals from lowest to highest odds against that seemed more pronounced for losses. Figure 2.3 shows that at lower odds against, mean residuals were generally negative and at higher odds against they are generally positive. This suggested that the hyperbolic model underestimated discounting rate at lower odds against (i.e. indifference points were steeper) and overestimated discounting rate at higher odds against (i.e. indifference points were shallower) for most of the participants. However, a repeated-measures ANOVA confirmed a significant linear trend for the effect of odds against on residuals for losses on the safe-risky task only (see Table 2.5 for sphericity and ANOVA statistics). Mauchly's test indicated that the assumption of sphericity had been violated for all four of the conditions, therefore degrees of freedom were corrected using the Greenhouse-Geisser and Huynh-Feldt estimates of sphericity where appropriate. For the safe-risky task, there was a significant effect of odds against on residuals in both conditions, but subsequent tests for whether the trends were linear showed a non-significant result for gains ($F(1, 35) = 3.58$, $p = .067$) and a significant linear trend for losses ($F(1, 35) = 11.53$, $p = .002$). For the risky-risky task, there were no significant effects of odds against in either condition.

These analyses suggested that the online format resulted in systematic discounting for most participants within the majority of the conditions, but data quality was overall poorer than in our in-lab data collections. The residual analysis showed a systematic linear pattern for the losses condition on the safe-risky task, suggesting a poor fit of the hyperbolic model to the data in one of the four conditions.

Table 2.5

Study 1, Experiment 1 money task sphericity and ANOVA statistics

	Sphericity				ANOVA			
	χ^2	df	p	ε	F	df	p	η_p^2
Safe-risky								
Gains	27.02	9	.001	.70[a]	2.80	2.82, 98.53	.047	.074
Losses	27.60		.001	.85[b]	5.04	3.38, 118.26	.002	.126
Risky-risky								
Gains	55.43	9	<.001	.63[a]	1.70	2.52, 85.53	.181	.048
Losses	27.25		.001	.84[b]	2.51	3.35, 113.84	.056	.069

Note. a = Greenhouse-Geisser correction; b = Huynh-Feldt correction.

Did the participants display a gain-loss asymmetry on the safe-risky money task?

The top panel of Figure 2.4 shows the individual and group data for the safe-risky task. No order effects on discounting rate were observed. The group median indifference points and their discounting curves for gains (dashed line) and losses (solid line) for the two tasks (*h* and AUC values are reported in Table 2.5) are shown on the right. Discounting on the safe-risky money task was steeper for gains than losses, and discounting of both gains and losses was steeper than the expected value (dotted line).

Individual participants' data are presented in the modified Brinley plots on the left for log(*h*) and on the right for AUC (means represented by a data point with error bars). For the safe-risky money task, 24 of the 36 data points (66.67%) on the log(*h*) graph indicated individuals with steeper discounting of gains than losses, as seen in the median data curves. However, this difference was not significant ($t(35) = -0.98$, $p = .334$, $d_z = 0.16$). Repeating the same analysis with AUC showed a gain-loss asymmetry where AUC for losses was significantly higher than for gains ($t(35) = 2.10$, $p = .043$, $d_z = 0.35$). The dotted reference lines on all graphs indicate choice based on the expected value ($h = 1$, $\log(h) = 0$). Most individuals discounted gains and losses more steeply than predicted by the expected value (points mostly in quadrant B), and the group mean log(*h*) value was significantly above zero for gains ($t(35) = 6.96$, $p < .001$, $d_z = 1.16$) and for losses ($t(35) = 3.28$, $p = .002$, $d_z = 0.55$).

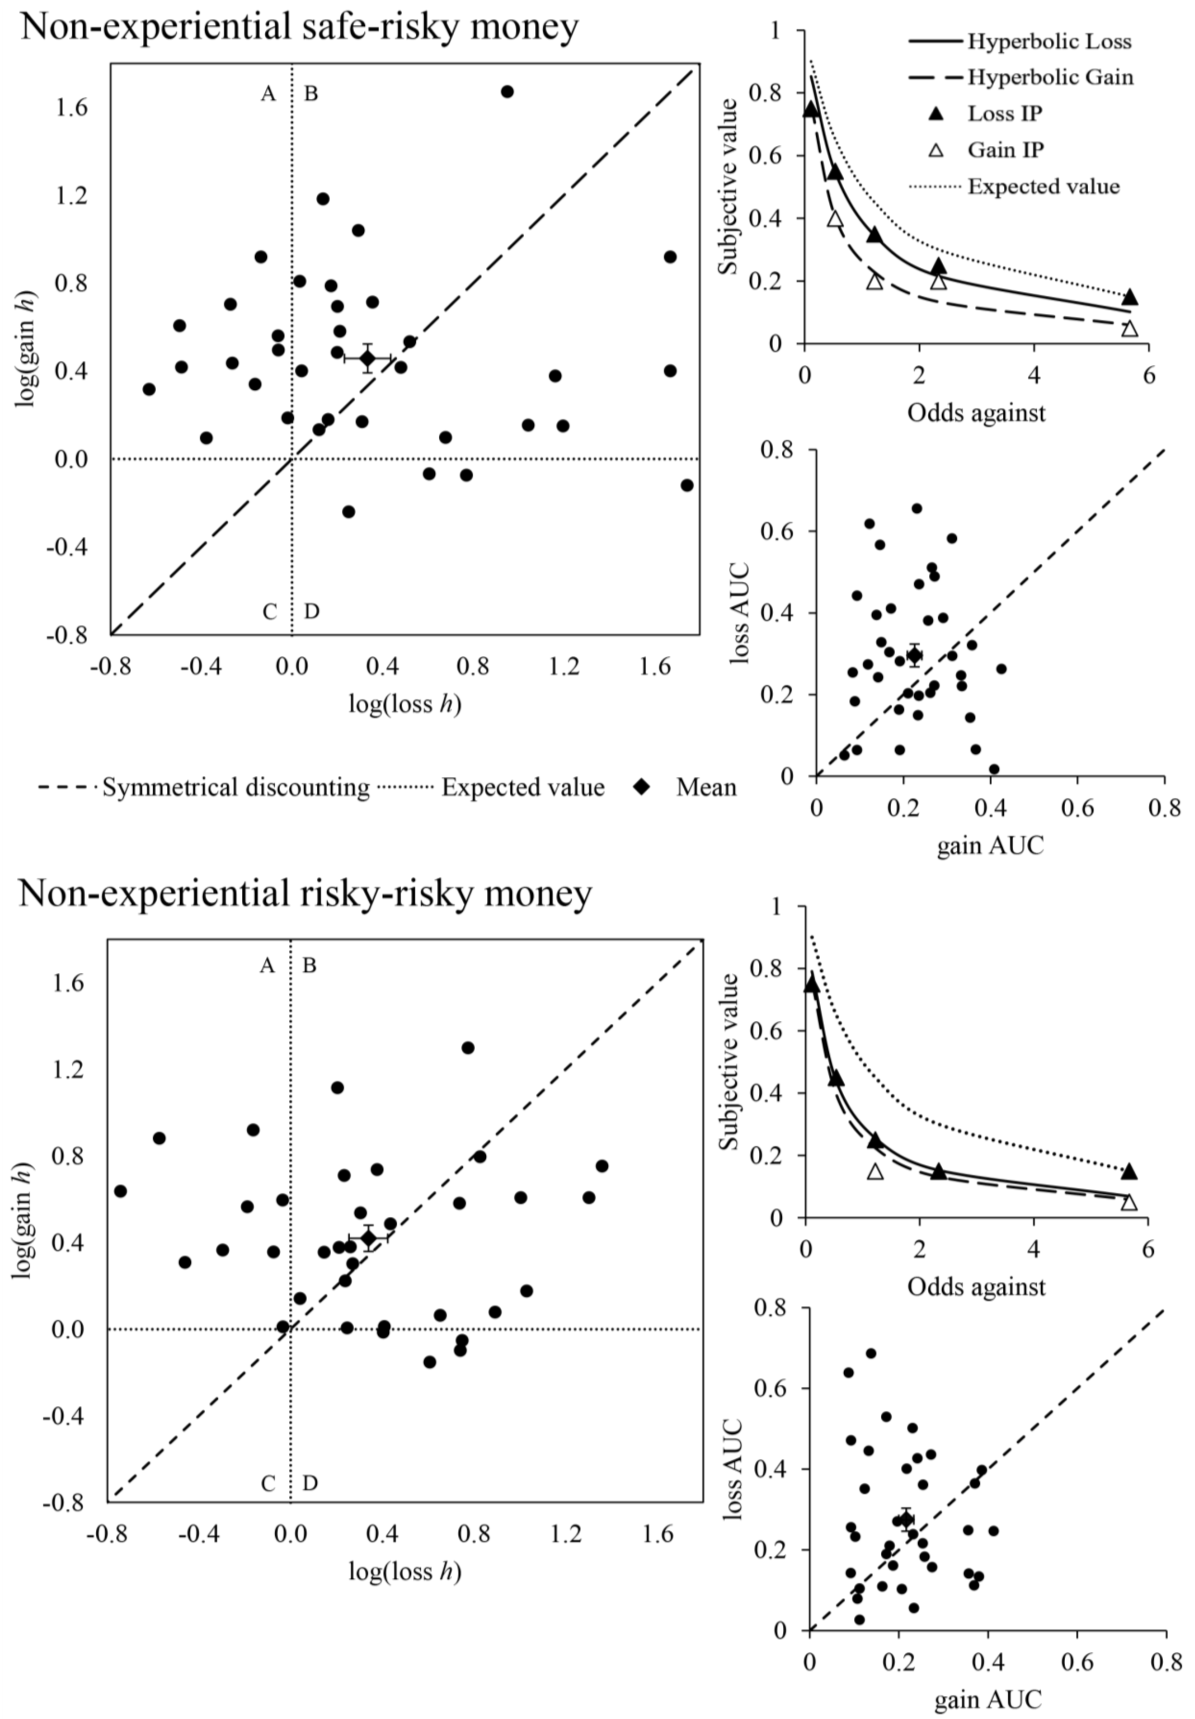

Figure 2.4. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual in the safe-risky (top half) and risky-risky (bottom half) money

tasks. The black diamond represents the mean and the error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged *h* value (0) when decisions are made based on expected value. Right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) for safe-risky (top half) and risky-risky (bottom half) money tasks. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Below the curves, AUC for losses is plotted against AUC for gains for each individual in the safe-risky (top half) and risky-risky (bottom half) money tasks. Group means and symmetrical discounting are as described for the left panel graph.

We expected discounting on the non-experiential safe-risky money task to be consistent with the gain-loss asymmetry seen in my honours thesis first experiment and the predictions of Prospect Theory. Discounting for both gains and losses was steeper than choice based on the expected value, consistent with our predictions. Furthermore, both measures of discounting showed a difference in median discounting rates consistent with literature (e.g. Estle et al., 2006). However, analysis of individual parameters using different discounting measures showed some inconsistencies. For one measure, log(h), this difference was not significant, but given the varied R^2 and some systematic deviations in residuals, the validity of using this parameter is questionable. For the AUC measure, the steeper discounting of gains than losses was significant and consistent with the median data. On balance, there was evidence for a gain-loss asymmetry consistent with the reflection effect in choice from description.

Did the participants display a gain-loss asymmetry on the risky-risky money task?

The bottom panel of Figure 2.4 shows the individual and group data for the risky-risky task. The hyperbolic model fitted to the median indifference points on the right showed that discounting on the risky-risky money task was similar for gains and losses, where both were steeper than expected value. For the individual data, 20 of the 35 data points (57.14%) on the log(h) graph were above the diagonal reference line and the discounting of gains and losses was not significantly different ($t(34) = -0.74$, $p = .467$, $d_z = 0.12$). As with log(h), AUC analysis did not show a gain-loss asymmetry ($t(35) = 1.69$, $p = .101$, $d_z = 0.29$). Discounting of gains ($t(34) = 6.94$, $p < .001$, $d_z = 1.17$) and losses ($t(34) = 4.01$, $p < .001$, $d_z = 0.68$) was steeper than predicted by the expected value.

We observed no consistent gain-loss asymmetry in the risky-risky condition, with both the median data and individual log(*h*) and AUC parameters showing no consistent difference between the gain and loss conditions. However, discounting for both gains and losses was steeper than choice based on expected value, similar to the safe-risky condition, suggesting that the risky-risky context affects the consistency of discounting of gains relative to losses, but not relative to the expected value.

Correlations between tasks and conditions

Correlations are presented in Table 2.6. Discounting based on log(*h*) across the two tasks was significantly, positively correlated for gains and for losses, suggesting that there was consistency in how the participants' discounting varied across the safe-risky and risky-risky contexts. The same pattern was observed across the two tasks based on AUC. Within each task, gains and losses were negatively, but not significantly, correlated based on both log(*h*) and AUC data.

Table 2.6

Study 1, Experiment 1 money task correlations for log(h) and AUC parameters

	Log(*h*)	AUC
E2.1		
Gains safe-risky - Losses safe-risky	-0.10	-0.14
Gains risky-risky - Losses risky-risky	-0.19	-0.10
Gains safe-risky - Gains risky-risky	0.58*	0.57*
Losses safe-risky - Losses risky-risky	0.76*	0.72*

Note. **p* < .01

Experiment 2: Discounting with no safe option and in experiential tasks

Study 1, Experiment 1 showed that the online format produced fewer systematic data sets than we expected, and therefore, for subsequent experiments we changed to in-lab data collections. Study 1, Experiment 1 also established that when participants discounted in a risky-risky context, they did not show the typical pattern of discounting gains steeper than losses seen in choice from description. In Study 1, Experiment 2, we combined a series of three data collections (sub-samples) that were aimed at examining the effects of choosing between two risky options, experience of outcomes and a combination of the two on the gain-loss asymmetry in ski and money task contexts.

Across all three sub-samples, we compared the discounting of gains and losses using two tasks: an experiential computer ski game adapted from Greenhow et al. (2015) and a money task. The money task was either non-experiential, or it resembled hypothetical tasks typically used but with experienced outcomes (i.e. combined description and experience). Including two distinct discounting tasks using different commodities and contexts allowed us to systematically replicate our measurement of gain-loss asymmetry. Any consistencies in gain-loss asymmetry across the tasks would speak to the consistency of the effects of experience and the risky-risky context on choice.

Table 2.7

Summary of the three sub-samples in Study 1, Experiment 2

Sub-sample	Money task	Ski task	*n*
E2.1	Experiential safe-risky	Experiential safe-risky	41 (21 recruited combined with 20 from honours thesis)
E2.2	Non-experiential risky-risky	Experiential risky-risky	52
E2.3	Experiential risky-risky	Experiential risky-risky	53

Table 2.7 summarizes the tasks and sample sizes in the three sub-samples. In E2.1, we expanded on the experiential data derived from my honours thesis to investigate the effects of experience on discounting in ski and money tasks, and combined our data sets for analysis. We expected steeper discounting of losses than gains in both tasks, in line with the reversed reflection effect in the description-experience gap literature (Hertwig & Erev, 2009). We also expected discounting to be either shallower than or close to choice based on the expected value of the outcomes. In E2.2 and E2.3, we combined the experiential component and the risky-risky context in money and ski game tasks. In E2.2, we compared discounting in a non-experiential money and an experiential ski tasks. In E2.3, we compared discounting in experiential money and experiential ski tasks. We also examined whether choice was risk seeking or risk averse in each condition and task, and determined whether there were any correlations within the sub-samples. We expected steeper discounting of losses than gains with inclusion of the experiential component. We also expected the combination of experiential component with risky-risky context to show generally steeper discounting than

choice based on the expected value, unlike that in experiential safe-risky tasks (Wulff et al., 2018).

Method

Participants

We recruited a total of 126 participants attending Victoria University of Wellington through the School of Psychology Research Programme online tool over several trimesters. The analyses below included data from 20 additional participants from the honours thesis data set for a total of 146 participants. First-year psychology students participated in partial fulfilment of a course requirement. The three parts of Study 1, Experiment 2 were run as three separate experiments, but are combined here for simplicity of explanation. Participants were first recruited for E2.1, and participants for E2.2 and E2.3 were recruited simultaneously.

Materials

All tasks were coded using Visual Basic. The titrating procedure was unchanged from Study 1, Experiment 1.

Ski Task. In the ski game the participants navigated a skier through a ski slope. The goal was to avoid certain features which deducted points and target certain features that added points. The ski task began with the following instructions:

> *You are a 'ski boarder' competing for points. The object of the task is to gain as many points as possible. You gain points for each jump you make over 'moguls' which look like: [image of a mogul]. You lose points for running into trees or rocks. Every so often you have to make a 'free run' at a jump platform. Before making such a jump you must choose ONE of TWO possible jump scenarios. Read the jump scenarios carefully as they will change from trial to trial. Use the mouse to click on the option you wish to choose. You move the player around using the DOWN, LEFT and RIGHT arrows. You can only move left or right, straight down, or at an angle downwards (make sure you spend some time at the start trying out the keys in order to become familiar with movement). Please wear the headphones attached to the computer during this task.*

Participants were given three minutes to familiarize themselves with the game rules and earn points before the experimental trials began (Figure 2.5, A). It was important for the participants to accumulate points so that points could subsequently be lost in the loss condition. Ten points were gained by navigating the player character over moguls (small ski jumps made from packed snow) and three points were lost for running into trees or rocks. The accumulated total score was presented on the top right of the screen. If the participant had fewer than fifty points after the first sixty seconds of the game, a prompt was generated by the program that read "Don't forget to ski over moguls and get points!". On average, participants accumulated 91 points before the experimental trials. After three minutes, the experimental trials began. Participants skied for 15 seconds between discounting choices.

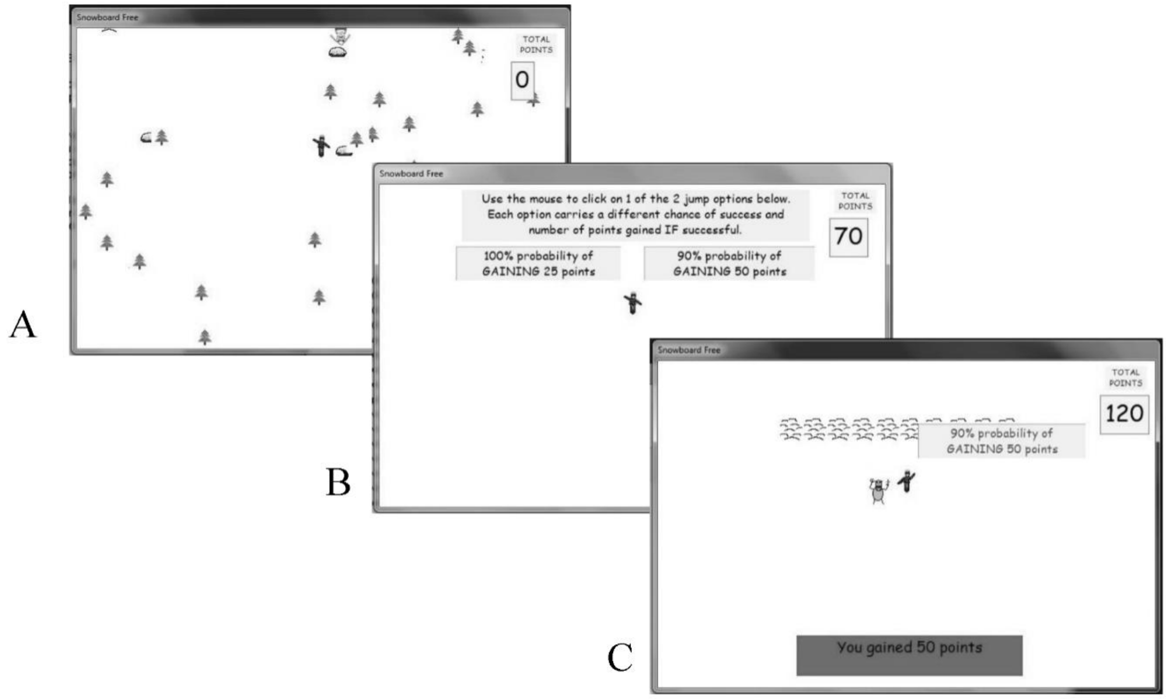

Figure 2.5. Screen capture of the ski task, consisting of the layout of the game in-between discounting options (A), the instructions and discounting options (B), and the outcome after clicking one of the options (C).

The instructions for the gains condition read: *Use the mouse to click on 1 of the 2 jump options below. Each option carries a different chance of success and number of points gained IF successful.*

The instructions for the losses condition read: *Use the mouse to click on 1 of the 2 jump options below. Each option carries a different chance of failure and number of points lost IF failed.*

The options were presented below the instructions (Figure 2.5, B). Once a choice was made, a simulation of the skier attempting the jump was played, and the outcome appeared at the bottom of the screen (Figure 2.5, C). If the jump resulted in points gained or no points lost (both being successful jumps), the outcome message box was coloured green and the skier remained upright. If the jump resulted in no points gained or points lost (both being failed jumps), the message box was coloured red and the skier was shown to have fallen into the snow. The accumulated total score was updated with the result of the jump attempt. The larger, uncertain amount was set at 50 points, while the smaller, uncertain amount started at half of the larger, uncertain amount (25 points), and was adjusted by 10% (5 points) of the larger, uncertain in the same manner as in Study 1, Experiment 1.

Non-Experiential Money Task. Instructions were as described in Study 1, Experiment 1. The alternatives were presented below the instructions, with the certain, smaller alternative on the left and the larger, uncertain on the right. The questions were arranged in six sets as described in Study 1, Experiment 1, but the questions in this in-lab version were shown on the screen one at a time (Figure 2.6). Choosing an option advanced the participants to the next question after a 500 millisecond delay until all the questions for a given condition were completed.

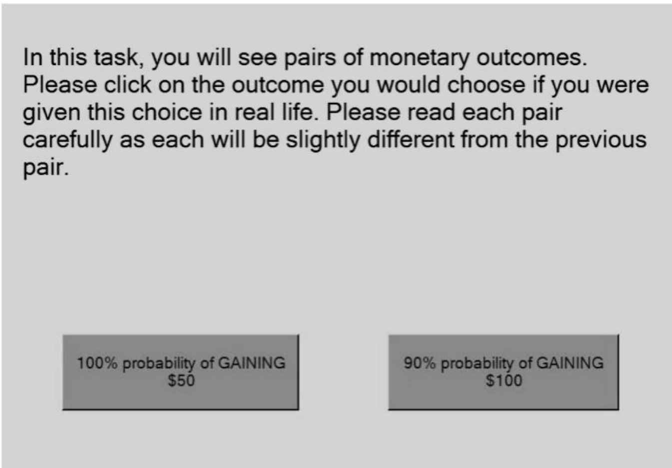

Figure 2.6. Screen capture of the money task instructions and options. Participants selected one of the options by clicking on them

Experiential Money Task. The instructions were as in Study 1, Experiment 1. Once the participants indicated their choice, an animation of a turning hourglass was played for two seconds above the options (Figure 2.7, B). The outcome appeared below the hourglass for two seconds after the animation ended (Figure 2.7, C). A counter indicating accumulated balance was added below the instructions and updated with each outcome (Figure 2.7, A).

The initial balance was set to $3000 in order to be comparable to the points gained from the non-discounting portion of the ski game. In the ski game, the first experiment of my honours thesis indicated that participants gained, on average, 1438 points during the non-discounting portion of the task, which is approximately 30 times the amount of the larger, uncertain. Hence, in the second experiment of my thesis the proportional balance in the money task was set to $3000 (30 times $100). This accumulated balance was subsequently deducted from or added to once the discounting portion commenced and avoided the possibility of a negative balance if the first condition scheduled was losses. Given the thirty trials (six sets of five probabilities), the starting value of $3000 ensured that if the first scheduled condition was losses, the participant chose to risk a loss of $100 in each trial and by chance failed at each trial (i.e. outcome was loss), then the total balance would not decrease below $0. This balance was also implemented in this set of experiments.

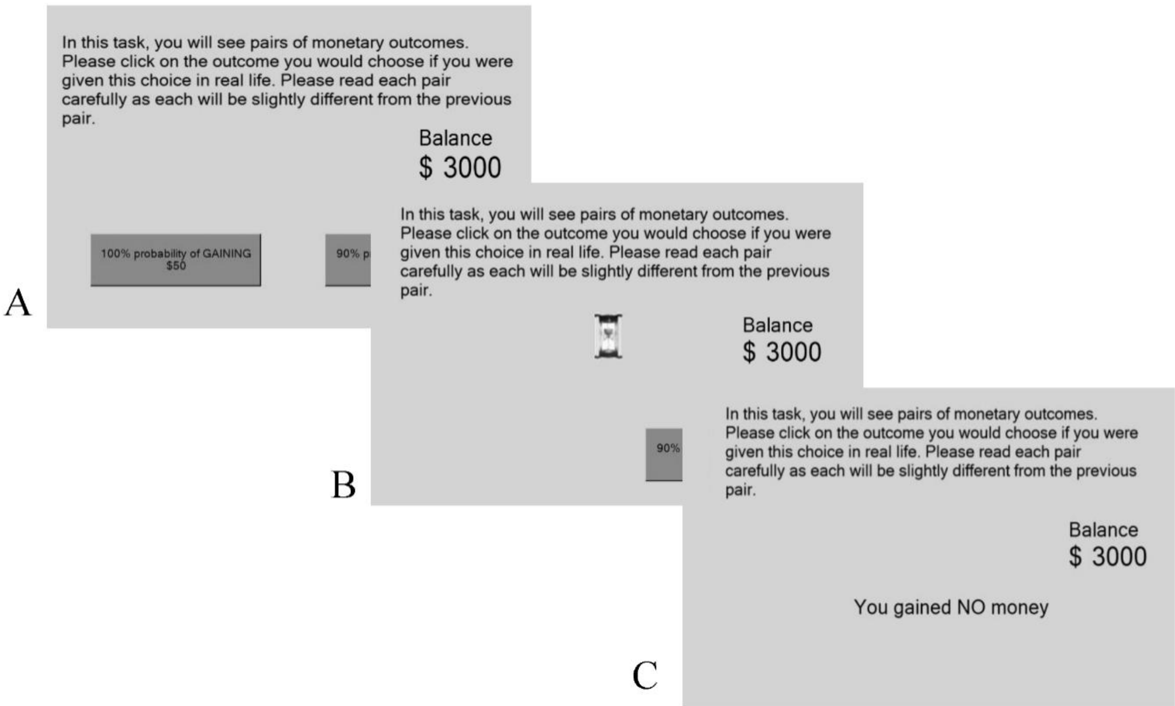

Figure 2.7. Screen capture of the experiential money task, consisting of the layout of the discounting options balance and instructions (A), the animation after choosing one of the options (B), and the outcome (C).

The Risky-Risky Manipulation. The risky-risky versions of the ski and money tasks were as described above, with the only exception that the probability of the smaller, certain option was displayed as 99%. For tasks that were also experiential, the probability of experiencing the event remained at 100% despite the displayed 99%.

Procedure

The participants were assigned to a cubicle with a computer terminal and headphones and tested in groups of four or fewer. Participants completed an informed consent procedure before starting the tasks. Within each sub-sample, each participant completed both tasks and both conditions within each task. The order of the money and ski tasks was counterbalanced. The order of the gain and loss conditions was also counterbalanced within the tasks.

Data Analysis and Exclusions

Data were analysed as described in Study 1, Experiment 1. Table 2.8 specifies the number of participants who failed the Johnson and Bickel (2008) criteria within each condition. We opted to present the analysis with unsystematic data sets removed. However, including all participants did not affect conclusions.

Table 2.8

Number of participants (n = 146) in Study 1, Experiment 2 who had unsystematic data by criterion, and total participants with systematic data.

	Money		Ski	
E2.1 (*n* = 41)	Experiential safe-risky		Experiential safe-risky	
	Gain	Loss	Gain	Loss
Criterion 1	3	1	1	1
Criterion 2	0	1	0	2
Total systematic in each task	36 (87.80)		37 (90.24)	
Total systematic across tasks	33 (80.49)			
E2.2 (*n* = 52)	Non-experiential risky-risky		Experiential risky-risky	
	Gain	Loss	Gain	Loss
Criterion 1	1	3	4	0
Criterion 2	1	5	1	2
Total systematic in each task	46 (88.46)		48 (92.31)	
Total systematic across tasks	43 (82.69)			
E2.3 (*n* = 53)	Experiential risky-risky		Experiential risky-risky	
	Gain	Loss	Gain	Loss
Criterion 1	1	1	0	2
Criterion 2	0	1	0	1
Total systematic in each task	51 (96.23)		51 (96.23)	
Total systematic across tasks	50 (94.34)			

Note. Percentages are in parentheses.

Individual *h* parameters were not normally distributed (Shapiro-Wilk's test for normality all $p < .001$). Log transforming the data resulted in *p* values > .05 for all conditions, except ski loss in E2.3 ($W = 0.95$, $p = .021$). The distribution in the ski loss condition showed slight positive kurtosis, but the shape of the distribution was otherwise normal. Therefore, log(*h*) values were used for analysis for all conditions. For AUC, Shapiro-Wilk's test for normality showed significant results for ski gains in 2.1 ($W = 0.91$, $p = .006$), money losses in 2.2 ($W = 0.92$, $p = .003$) and ski losses in 2.3 ($W = 0.91$, $p = .001$), all of which were showing slight positive skew and kurtosis, but the shape of the distributions was otherwise normal.

Results and Discussion

The hyperbolic model provided good fits to the individual and group median data (see Table 2.9 for R^2 values). Notably, the three sub-samples provided a narrower range of R^2 values and a lower percentage of overall unsystematic data sets than in Study 1, Experiment 1. Furthermore, no systematic trends were observed in mean residuals (see Appendix B), suggesting that the poorer quality of data in Study 1, Experiment 1 was most likely due to the features of the online format.

Table 2.9

Study 1, Experiment 2 ski and money task h, R^2 and AUC values

	Group median indifference points			Individual participants		
E2.1	*h*	R^2	*AUC*	*h* median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Money gain	1.74	0.97	0.27	1.68 (1.32, 3.05)	0.83 (0.77, 0.90)	0.27 (0.02)
Money loss	1.83	0.96	0.23	1.87 (1.06, 2.90)	0.85 (0.73, 0.91)	0.26 (0.02)
Ski gain	1.25	0.98	0.31	1.25 (0.79, 2.27)	0.91 (0.86, 0.95)	0.32 (0.02)
Ski loss	1.09	0.99	0.33	1.22 (0.73, 1.60)	0.87 (0.77, 0.93)	0.33 (0.02)
E2.2						
Money gain	2.06	0.99	0.21	1.83 (1.15, 3.34)	0.86 (0.78, 0.92)	0.24 (0.02)
Money loss	1.60	0.98	0.26	1.81 (0.75, 2.85)	0.86 (0.73, 0.92)	0.30 (0.02)
Ski gain	1.30	0.96	0.28	1.30 (0.79, 2.15)	0.89 (0.82, 0.93)	0.30 (0.01)
Ski loss	1.60	0.98	0.26	1.50 (0.90, 2.25)	0.90 (0.82, 0.94)	0.30 (0.02)
E2.3						
Money gain	1.66	0.94	0.23	1.60 (1.00, 2.87)	0.85 (0.79, 0.89)	0.25 (0.02)
Money loss	1.79	0.91	0.22	1.65 (1.12, 2.85)	0.87 (0.79, 0.92)	0.24 (0.01)
Ski gain	1.02	0.93	0.30	1.02 (0.69, 1.60)	0.86 (0.81, 0.92)	0.32 (0.02)
Ski loss	1.33	0.92	0.31	1.32 (0.91, 2.00)	0.86 (0.72, 0.92)	0.30 (0.02)

Note. Q_1 = first quartile, Q_3 = third quartile.

Examination of the group median and individual data indicated an order effect in some sub-samples, where discounting depended on the order in which participants completed the tasks and conditions. Therefore the data in each sub-sample were examined for both condition- (gain-loss vs. loss-gain) and task- (money first vs. money second) order effects. The right panels of Figures 2.8-2.10 show the hyperbolic equation fitted to the group median indifference points as described in Study 1, Experiment 1 and are split by order of condition. The individual log(*h*) and AUC data are shown in the left and right panels respectively, where black data points denote participants who experienced gains first and white data points denote participants who experienced losses first. Means are represented by data points with errors bars, with the circles denoting data split by order and the square data point denoting the overall mean not split by order.

E2.1 Gain-loss asymmetry in experiential money and ski tasks

The experiential money task showed that the participants did not discount losses more steeply than gains. Rather, participants who completed the gains condition first discounted gains more steeply than losses and participants who completed the losses condition first discounted losses more steeply than gains. Group median curves show a distinct reversal in which condition is discounted more steeply based on order. Both the individual log(*h*) (left panel of Figure 2.8) and AUC parameters (right panel) show more black circles above the diagonal line, indicating steeper discounting of gains in the gains first order, and more white circles below the line, indicating steeper discounting of losses in the losses first order. A mixed measures ANOVA, with condition (gains vs. losses) as a within-subjects factor, condition order (gain-loss vs. loss-gain) as a between-subjects factor, and task order (money first vs. money second) as a between-subjects factor confirmed that log(*h*) values were determined by a significant interaction between condition and condition order (see Table 2.10). Simple main effects analysis using a Bonferroni correction showed that the difference between gains (M = 0.39, SE = 0.08) and losses (M = 0.17, SE = 0.07) when gains were first was significant (p = .027), but the difference between gains and losses when losses were first was not significant (p = .126).

Notably, while the same analysis using AUCs confirmed the log(*h*) ANOVA results (see Appendix C), the follow-up tests indicated that the difference between gains and losses when gains were first was not significant (p = .057), while the difference between gains (M = 0.31, SE = 0.03) and losses (M = 0.22, SE = 0.03) when losses were first was significant (p =

.016). The group and individual data showed that the log(h) parameters for both gains ($t(35) = 5.18$, $p < .001$, $d_z = 0.78$) and losses ($t(35) = 4.71$, $p < .001$, $d_z = 0.86$) were greater than choice based on the expected value ($h = 1$, $\log(h) = 0$).

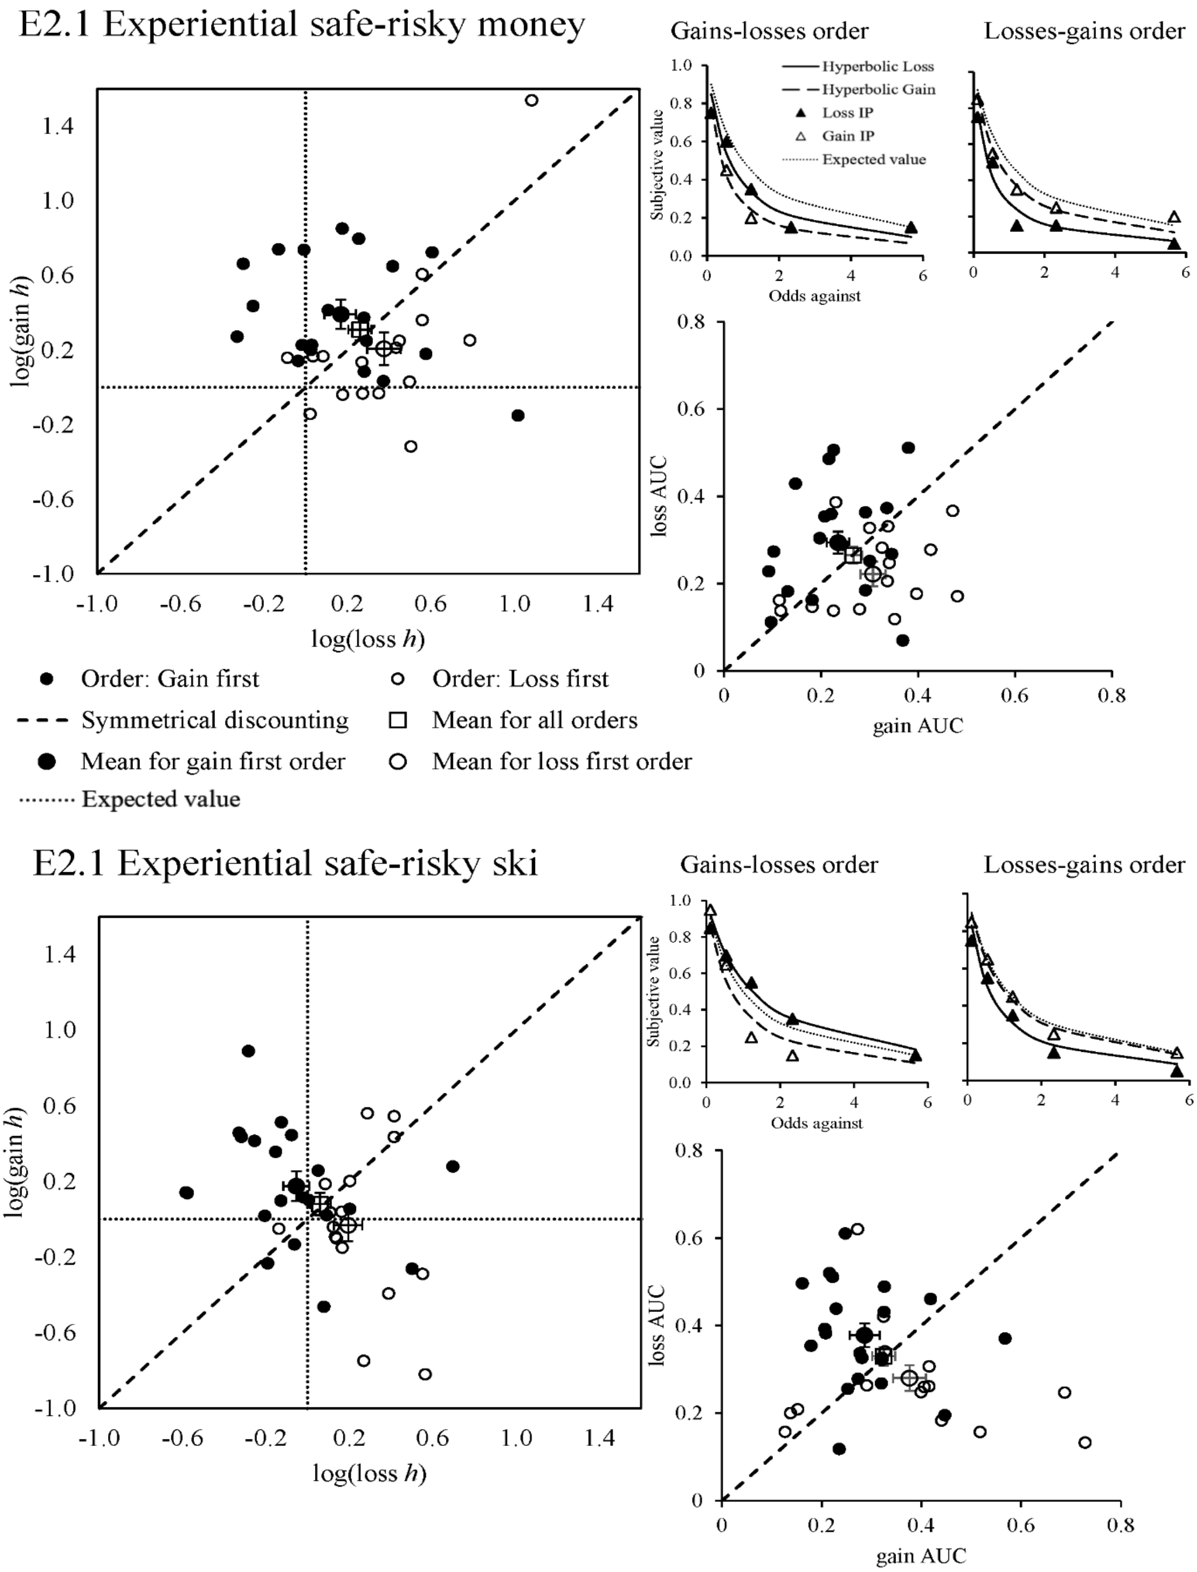

Figure 2.8. Left panel: Logged h parameters for gains plotted against logged h parameters for losses for each individual in the money (top half) and ski (bottom half) tasks, split by order;

black circles are gain-loss and white circles are loss-gain order. The larger black and white circles represent the means split by order, the white square the overall mean and the error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged h value (0) when decisions are made based on expected value. Right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) for money (top half) and ski (bottom half) tasks. For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Below the curves, AUC for losses is plotted against AUC for gains for each individual in the money (top half) and ski (bottom half) tasks. Group means and symmetrical discounting are as described for the left panel graph.

Data for the experiential ski task showed a similar effect of order on the difference between gains and losses. A mixed measures ANOVA showed no significant main effects and the only significant interaction was between condition order and condition. Simple main effects analysis using a Bonferroni correction showed no significant effect of condition on log(h) values either when gains were presented first (p = .051), or when losses were presented first (p = .073). For AUC, while the ANOVA results were consistent with log(h), simple main effects analysis showed a significant difference between gains (M = 0.29, SE = 0.03) and losses (M = 0.38, SE = 0.03) for the gains first order (p = .047) and no significant difference between gains and losses for the losses first order (p = .060). Contrary to the money task, discounting in both conditions approximated discounting based on the expected value. Neither log(h) for gains ($t(36)$ = 1.35, p = .187, d_z = 0.22) nor losses ($t(36)$ = 1.21, p = .235, d_z = 0.20) were significantly different from choice based on the expected value.

E2.2 Gain-loss asymmetry in risky-risky non-experiential money and experiential ski tasks

Discounting in the non-experiential risky-risky money task showed the same pattern as in Study 1, Experiment 1: discounting of gains and losses were not significantly different, and both were discounted more steeply than choice based on the expected value (Figure 2.9). A mixed measures ANOVA showed no significant main-effects of order or condition, and non-significant interactions with the exception of task order and condition order (Table 2.10).

Table 2.10

Study 1, Experiment 2 ANOVA results using log(h)

	F	*df*	*p*	η_p^2	*F*	*df*	*p*	η_p^2
	Money				Ski			
E2.1	Safe-risky experiential				Safe-risky experiential			
Condition	0.06	1, 32	.814	.002	<0.001	1, 33	.999	<.001
Condition order	0.03		.877	.001	0.01		.924	<.001
Task order	0.14		.709	.004	1.20		.281	.035
Condition by Condition order	**6.99**		**.013**	**.179**	**7.10**		**.012**	**.177**
Condition by Task order	0.87		.359	.026	0.16		.694	.005
Condition order by Task order	0.21		.653	.006	2.52		.122	.071
Condition by Condition order by Task order	1.75		.195	.052	0.33		.569	.010
E2.2	Risky-risky non-experiential				Risky-risky experiential			
Condition	1.68	1, 42	.202	.038	0.21	1, 44	.648	.005
Condition order	2.44		.125^	.055	0.93		.339	.021
Task order	0.04		.834	.001	0.91		.346	.020
Condition by Condition order	0.01		.930	<.001	0.39		.534	.009
Condition by Task order	0.01		.906	<.001	0.65		.424	.015
Condition order by Task order	**4.12**		**.049^**	**.089**	0.01		.916	<.001
Condition by Condition order by Task order	0.01		.914	<.001	0.03		.865	.001
E2.3	Risky-risky experiential				Risky-risky experiential			
Condition	0.13	1, 47	.723	.003	4.00	1, 47	.051	.078
Condition order	0.73		.396	.015	1.15		.289	.024
Task order	**12.88**		**.001**	**.215**	1.96		.168	.040
Condition by Condition order	**23.60**		**<.001**	**.334**	0.37		.546	.008
Condition by Task order	0.49		.488	.010	2.01		.163	.041
Condition order by Task order	0.23		.637	.005	0.02		.891	<.001
Condition by Condition order by Task order	0.86		.358	.018	0.11		.746	.002

Note. Bold emphasis added to significant results. ^ indicates results that differed when using AUC (table in Appendix C).

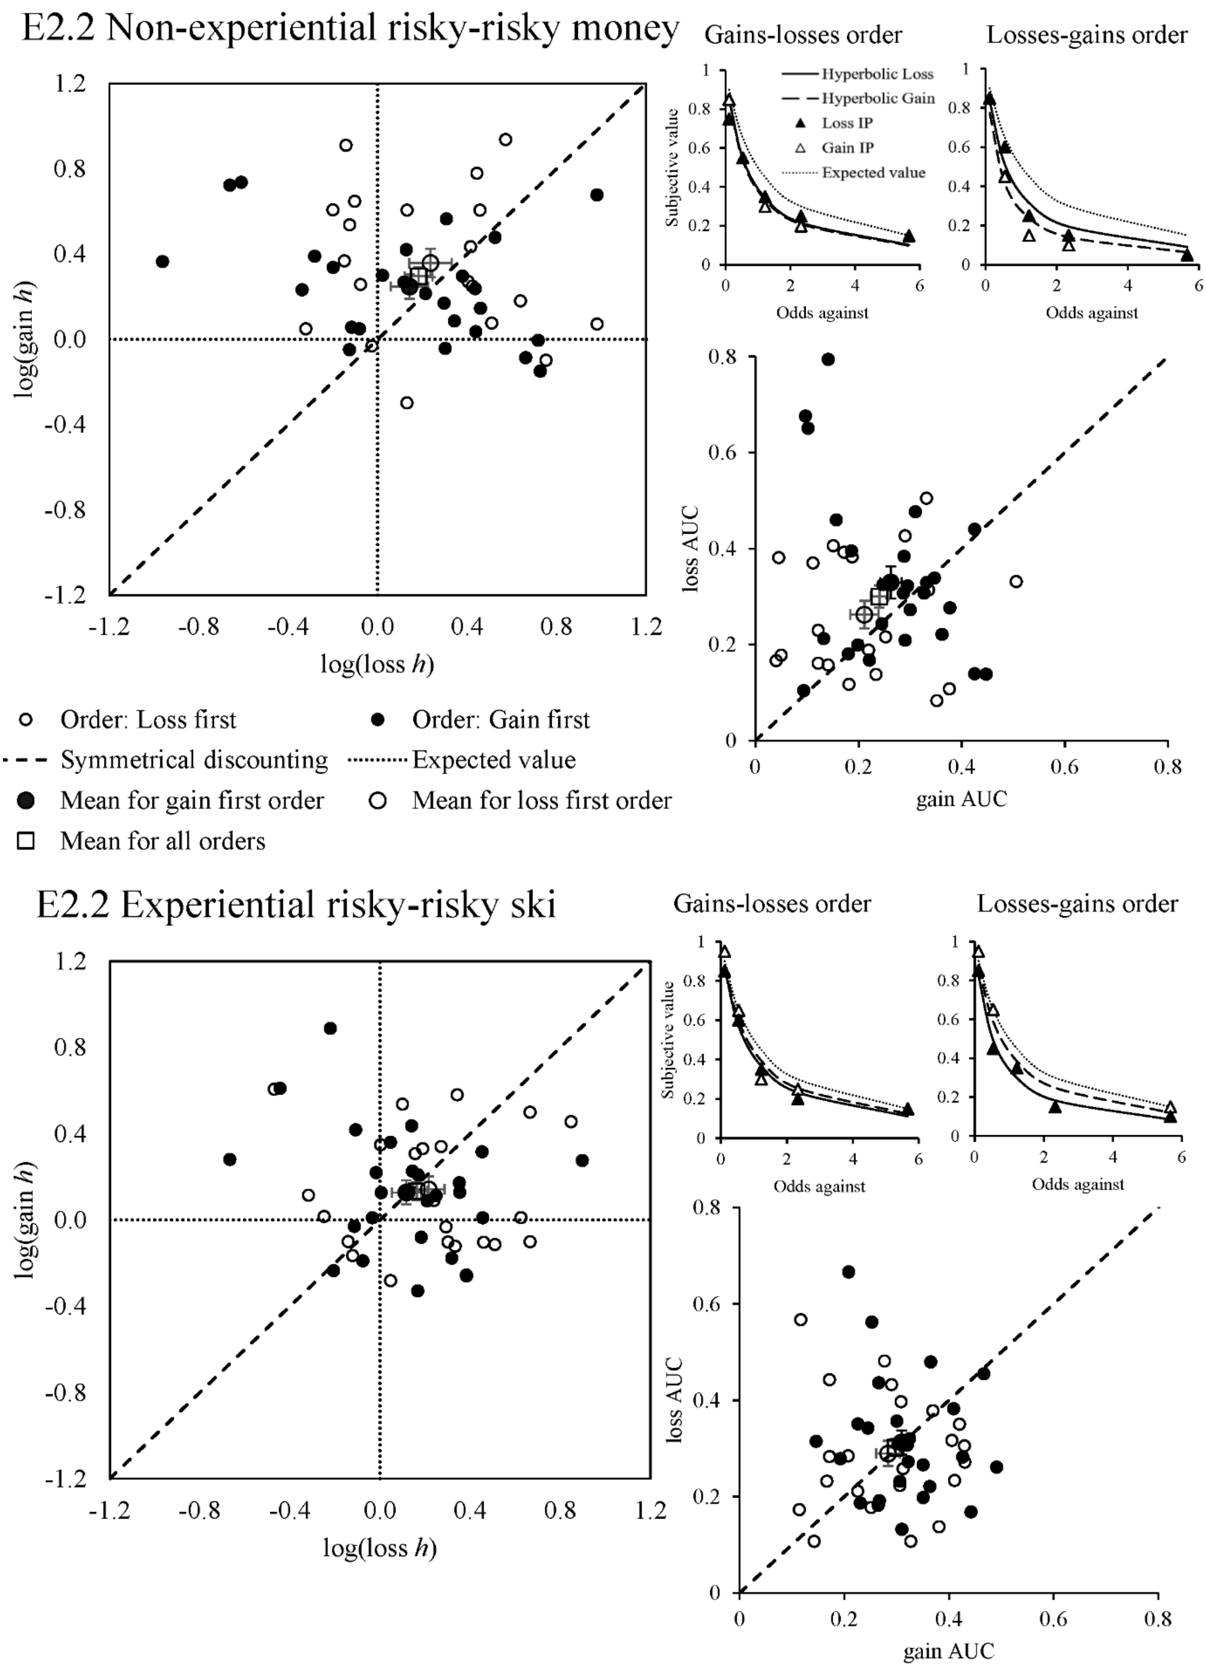

Figure 2.9. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual in the money (top half) and ski (bottom half) tasks, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and white

circles represent the means split by order, the white square the overall mean and the error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged *h* value (0) when decisions are made based on expected value. Right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) for money (top half) and ski (bottom half) tasks. For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value ($h = 1$). Below the curves, AUC for losses is plotted against AUC for gains for each individual in the money (top half) and ski (bottom half) tasks. Group means and symmetrical discounting are as described for the left panel graph.

Simple main effects analysis using a Bonferroni correction showed that when the money task was first, the average log(*h*) across gains and losses in the loss first order (M = 0.37, SD = 0.07) was steeper than in the gains first order (M = 0.14, SD = 0.06; p = .015). When the money task was second, there was no significant difference in the average log(*h*) between the two condition orders (p = .743). These significant effects of order were in contrast to the AUC results, which found no significant interaction of condition order and task order, but a significant main effect of condition order. Mean AUC parameter for gains and losses was significantly lower when losses were first (M = 0.24, SD = 0.02) compared to when gains were first (M = 0.30, SD = 0.02). Logged *h* was significantly higher than zero for both gains ($t(45) = 6.80$, $p < .001$, $d_z = 1.00$) and losses ($t(45) = 2.93$, $p = .005$, $d_z = 0.43$).

Discounting in the ski task was similar to that of E2.1, in that discounting on one condition was not consistently steeper than in the other. However, a mixed measures ANOVA showed no significant main effects or interaction when using log(*h*) and when using AUC, suggesting that discounting was unaffected by condition or task order. Unlike the ski task in E2.1, both gains and losses conditions were discounted more steeply than expected value; log(*h*) values were significantly higher than zero for gains ($t(47) = 3.30$, $p = .002$, d_z = 0.48) and losses ($t(47) = 3.41$, $p = .001$, $d_z = 0.49$).

E2.3 Gain-loss asymmetry in risky-risky experiential money and ski tasks

A combination of experiential and risky-risky features in the money task produced order effects akin to E2.1 (see Figure 2.10). A mixed measures ANOVA showed no significant main effects of condition or condition order, a significant main effect of task order and one significant interaction between condition and condition order (see Table 2.10). The main effect of task order indicated that the average log(*h*) for gains and losses was higher when the money task was first (M = 0.37, SE = 0.04) compared to when the money task was second (M = 0.15, SE = 0.05; p = .001). Simple main effects analysis using a Bonferroni correction showed that the differences between gains (M = 0.40, SE = 0.06) and losses (M = 0.18, SE = 0.06) when gains were first (p = .002) and between gains (M = 0.10, SE = 0.06) and losses (M = 0.36, SE = 0.06) when losses were first (p = .001) were both significant. Analysis of AUC parameters confirmed the ANOVA results and the follow-up tests. Log(*h*) values were significantly higher than zero in the gain ($t(50)$ = 5.50, p < .001, d_z = 0.77) and loss ($t(50)$ = 5.94, p < .001, d_z = 0.83) conditions.

Discounting in the ski task largely replicated the patterns in E2.2. A mixed measures ANOVA showed no significant main effects or interactions, showing no consistent gain-loss asymmetry or effect of order on discounting. Analysis using AUC confirmed these results. Log(*h*) values were significantly higher than zero in the loss condition ($t(50)$ = 3.35, p = .002, d_z = 0.47), but not in the gain condition ($t(50)$ = 0.89, p = .377, d_z = 0.12).

Correlations between conditions and tasks

In addition to examining gain-loss asymmetry within tasks, we also considered if there was consistency in how discounting rates varied across the tasks. In all three subsamples, and for both log(*h*) and AUC, gains and losses were significantly positively correlated across tasks, while gains and losses were not significantly correlated within the tasks (see Table 2.11). A positive correlation suggested that even while discounting rates were sometimes affected by order of presentation, there was also consistency in how the participants' discounting varied across the two tasks.

Table 2.11

Study 1, Experiment 2 correlations for log(h) and AUC values.

	Log(*h*)	AUC
E2.1		
Gains money - Losses money	-0.22	0.06
Gains ski - Losses ski	-0.25	-0.22
Gains money - Gains ski	0.69*	0.61*
Losses money - Losses ski	0.58*	0.55*

Table 2.11 (continued)

	Log(h)	AUC
E2.2		
Gains money - Losses money	-0.26	-0.16
Gains ski - Losses ski	-0.17	-0.10
Gains money - Gains ski	0.44*	0.48*
Losses money - Losses ski	0.59*	0.62*
E2.3		
Gains money - Losses money	0.15	0.16
Gains ski - Losses ski	0.09	0.08
Gains money - Gains ski	0.67*	0.59*
Losses money - Losses ski	0.47*	0.37*

Note. $^*p < .01$

Do the risky-risky and experience contexts produce a consistent gain-loss asymmetry?

Study 1, Experiment 2 aimed to further investigate the gain-loss asymmetry in risky-risky and experiential contexts. E2.1 resolved the discrepancy in the two data sets from my honours thesis: experience of outcomes resulted in an order effect consistent across the money and ski contexts, where participants discounted outcomes more steeply in the first condition they completed. This effect of order was also apparent in the experiential risky-risky money task in E2.3 and absent in the non-experiential risky-risky money task in E2.2, suggesting that this order effect depends primarily on the experiential component. The effect of order also appears to be more likely in the money rather than ski tasks. Unlike the ski task in E2.1, neither of the ski tasks in sub-samples E2.2 and E2.3 showed an effect of order on discounting, despite the inclusion of the experiential component. The addition of the risky-risky manipulation to the existing experiential features in these ski tasks does not seem to sufficiently explain this, as such a combination in the money task still produced order effects in E2.3. Additionally, the money tasks in E2.2 and E2.3 both showed further influences of task order, which was not observed in any of the ski tasks. The task order effect in E2.3, in particular, indicated that discounting became shallower with task progression, which is in the same direction as the condition order effects.

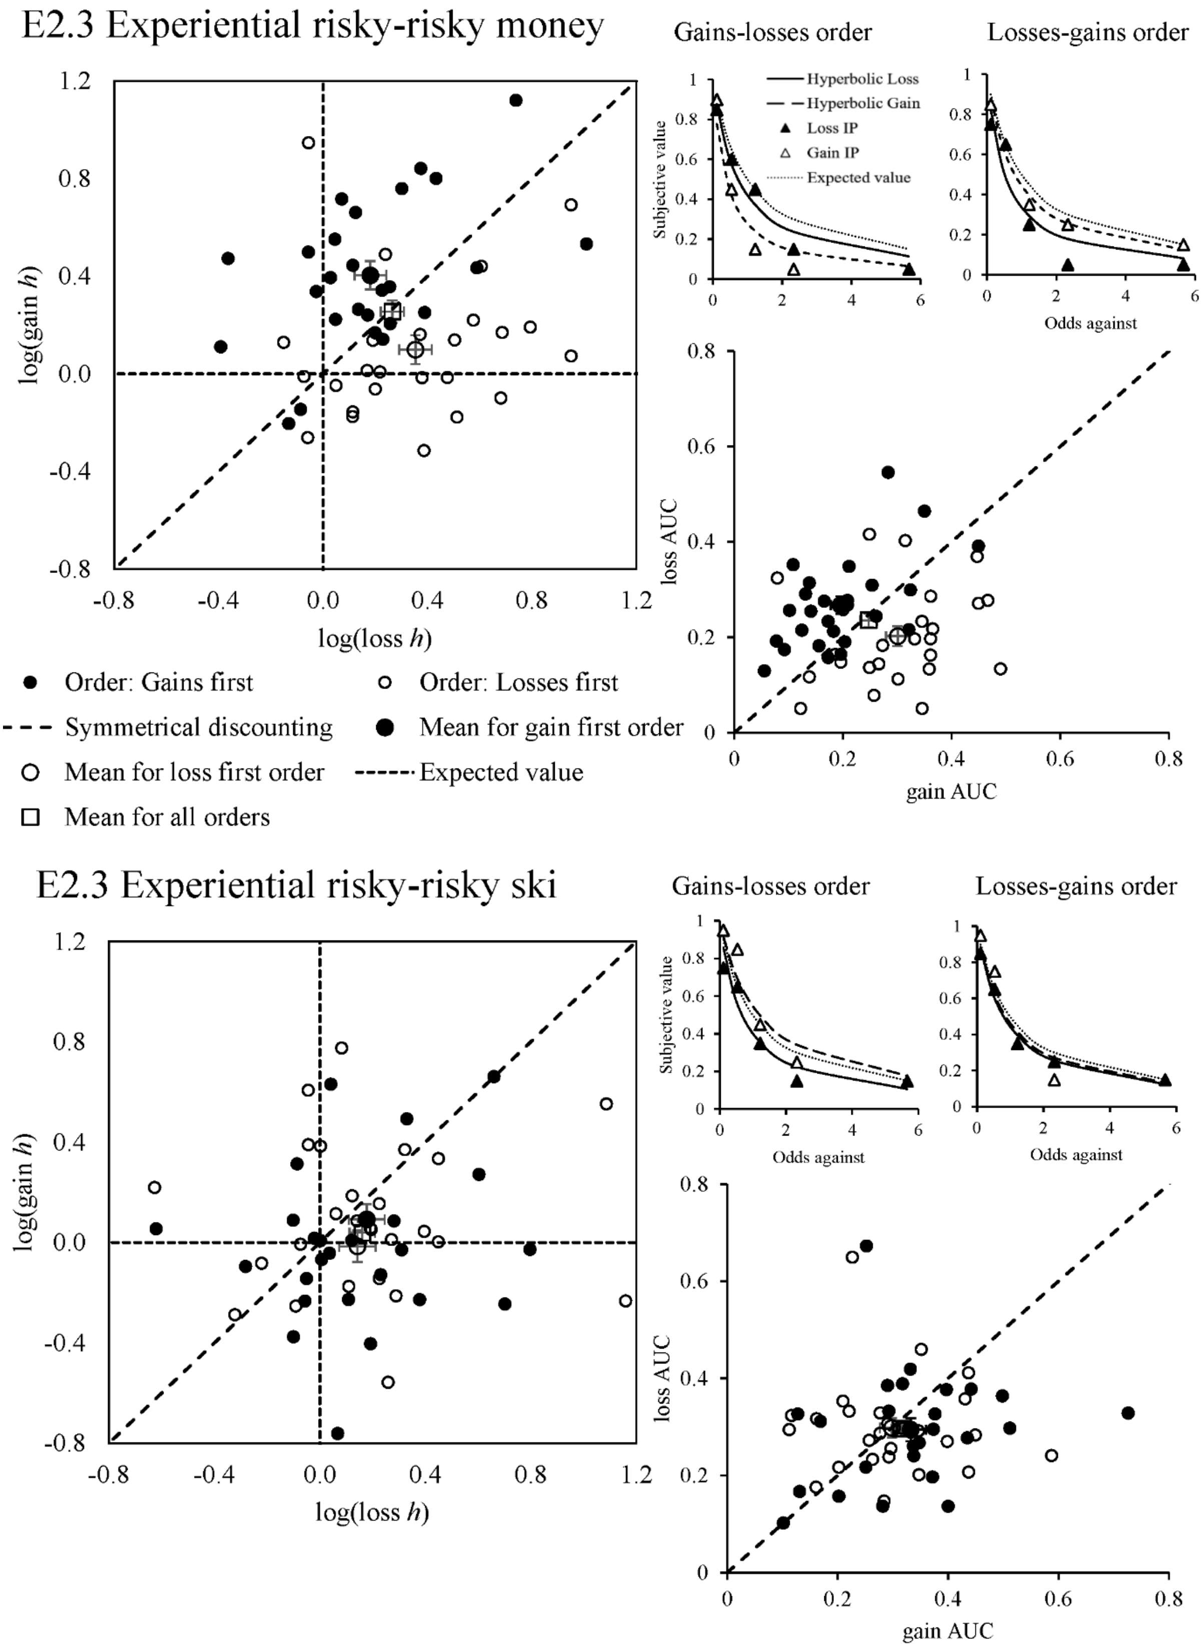

Figure 2.10. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual in the money (top half) and ski (bottom half) tasks, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and

white circles represent the means split by order, the white square the overall mean and the error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged *h* value (0) when decisions are made based on expected value. Right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) for money (top half) and ski (bottom half) tasks. For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Below the curves, AUC for losses is plotted against AUC for gains for each individual in the money (top half) and ski (bottom half) tasks. Group means and symmetrical discounting are as described for the left panel graph.

In all of the money tasks, participants discounted gains and losses more steeply than choice based on expected value. While this was consistent with our predictions for the risky-risky contexts, both experiential and non-experiential, it was inconsistent with the predictions for the experiential safe-risky context, where we expected shallower discounting. The predicted pattern was not altogether absent, as the safe-risky experiential ski task in E2.1 showed discounting that was not significantly different from expected value, echoing the differences between the money and ski tasks in relation to the consistency of order effects.

Was the observed order effect due to experience of outcomes?

Overall, discounting in atypical contexts, (i.e. risky-risky and experience), does not produce the consistent gain-loss asymmetry that is seen in choice from description in either the ski or the money task contexts. Rather, discounting depends on the recent experience of the participants, more so for the money context rather than the ski context. Similar order effects have been noted in studies where the participants experienced two different probability discounting conditions. Matusiewicz et al. (2013) conducted a study where participants completed a real-rewards lottery-style task and a hypothetical money task over two sessions held one week apart. The real-rewards task did not show outcomes for each option, but at the end of each session the participants experienced one of their previously chosen options randomly selected by the experimenter. Participants assigned to the real-rewards condition in the first session showed shallower discounting in the real-rewards condition in the second session if they experienced one of their probabilistic choices at the end of the first session. This effect was strongest in the participants that won as opposed to lost on their probabilistic choice in the first session. A similar order effect has been noted by Hinvest and Anderson (2010) who compared two experiential tasks, but only one condition had real lottery-style rewards at the end of the session. They found an effect of order of task presentation, where completing the experiential task with no real rewards in session one resulted in shallower discounting in session two where the participants played for potentially real lottery-style rewards. Hinvest and Anderson proposed that the experience of probabilistic outcomes resulted in subsequent greater risk taking for the task offering potentially real rewards. No difference in discounting was observed for the reverse order, which seems to be inconsistent with Matusiewicz and colleagues (see Yi, Stuppy-Sullivan, Pickover, & Landes, 2017 for a discussion on the effect of switching between control and experimental conditions on discounting).

Experiment 3: Experiential money test-retest task

In Study 1, Experiment 2, discounting of gains relative to losses was largely inconsistent across participants in both experiential and risky-risky contexts. The most consistently observed pattern was the effect of order on discounting. In some of the tasks, whether gains or losses were discounted more steeply depended on the order they were experienced in. That is, participants discounted outcomes more steeply in the first condition than in the second condition regardless of whether those outcomes were gains or losses. This effect of order was consistently seen in experiential money tasks, but less so in experiential ski tasks. Furthermore, discounting in the money task in E2.3 also showed susceptibility to task order, with participants discounting more shallowly after the ski task.

However, follow-up tests that subdivided each condition by order did not consistently show a significant difference in discounting rate. In E2.1, there was only a significant difference in discounting rate between gains and losses in one of four groups: the gains first order in the money task, which showed higher mean *h* values for gains than losses. In E2.3, the experiential money task showed a significant difference in discounting rate for both groups: the gains first order showed steeper discounting of gains than losses, and the loss first order showed steeper discounting of losses than gains. Notably, in both of the sub-samples, this was a post-hoc exploratory analysis as we did not expect to find order effects and Study 1, Experiment 2 was underpowered to detect this effect. The interaction effect sizes for the money task in E2.1 (η_p^2 = .18) and E2.3 (η_p^2 = .33) are considered large (Cohen, 1988). A power analysis with alpha at .05 and power at .80, recommends at least 54 participants to detect a large effect size and up to 128 to detect a moderate effect size in the follow-up tests to the interaction (calculation done in GPower 3.1.9.2). A larger sample would also allow us to determine whether the significantly steeper discounting of both losses and gains relative to expected value in the money tasks was driven by the first condition alone, or whether *h* remained below 1 in the second condition completed.

Therefore, Study 1, Experiment 3 replicated the safe-risky experiential money task from E2.1 with a larger sample in order to determine whether gain-loss asymmetry occurs in probability discounting of experienced outcomes after taking order into account. We expected that discounting of gains and losses would depend on condition order. We also administered a second experiential money task in order to determine whether discounting rates stabilized following initial exposure to the task or whether discounting rates became shallower with continued exposure to experienced outcomes. The latter result could explain the task order effect in the money task found in E2.3. We also examined whether discounting rates were correlated across the two repetitions of the gain condition and the two repetitions of the loss condition. Finally, we examined whether the relationship between gains and losses at the individual level was correlated across the two tasks.

Method

Participants

One hundred and twenty-seven undergraduates participated in partial fulfilment of course requirements. Informed consent was collected.

Apparatus

The safe-risky experiential money task was as described in Study 1, Experiment 2.

Procedure

The participants were tested as described in Study 1, Experiment 2. Each participant completed two experiential money discounting tasks with a gain and a loss condition within each. The two tasks were separated by a 30 second interval stating the following on a full-screen blue background:

> *You have completed Part 1. Please take a break before proceeding to Part 2. Part 2 will start automatically in 30 seconds.*

The order of presentation of the two conditions (gains and losses) was counterbalanced in the first task. In the second task, half of the participants experienced the same order as in Task 1 and half the reverse order.

Data Analysis and Participant Exclusions

Data were analysed as described in Study 1, Experiment 1. Individual *h* parameters were not normally distributed (Shapiro-Wilk's test for normality all $p < .001$). Log transforming the data resulted in *p* values > .05 for gains in both tasks, but not losses in Task 1 ($W = 0.97$, $p = .008$) or Task 2 ($W = 0.98$, $p = .032$). Examination of histograms showed slight positive skew for losses in both tasks and some outliers at the higher end of the distribution, but the shape of the distribution was otherwise normal. We also determined the degree and direction of each participants' gain-loss asymmetry across the two task repetitions by calculating a *gain-loss asymmetry score*: log (loss *h*/gain *h*). Positive logged ratios indicate higher *h* parameter for losses than gains and negative logged ratios indicate the reverse. The gain-loss asymmetry scores for Task 1 and Task 2 were normally distributed (Shapiro-Wilk's test for normality all $p > .05$).

Table 2.12 specifies the number of participants excluded based on Johnson and Bickel's (2008) criteria. The analysis below uses participants who had fully systematic data across tasks, with 59 experiencing the two tasks in the same order (31 with gains first) and 57 in the swapped order (27 with gains first in the first task).

Table 2.12

Number of participants (n = 127) in Study 1, Experiment 3 who had unsystematic data by criterion, and total participants with systematic data.

E2.1 (*n* = 41)	Task 1		Task 2	
	Gain	Loss	Gain	Loss
Criterion 1	4	3	2	1
Criterion 2	1	3	1	1
Total systematic in each task	118 (92.91)		123 (96.85)	
Total systematic across tasks	116 (91.34)			

Note. Percentages in parenthesis.

Results

The hyperbolic model provided good fits to the individual and group median indifference points (see Table 2.13). No systematic trends were observed in the mean residuals (see Appendix B).

Table 2.13

Study 1, Experiment 3 money task h, R^2 and AUC values in Task 1 and 2

	Group median indifference points			Individual participants		
	h	R^2	*AUC*	*h* median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Task 1 Gain	1.38	0.96	0.27	1.55 (0.91, 2.70)	0.88 (0.80, 0.92)	0.28 (0.01)
Task 1 Loss	1.83	0.96	0.23	1.80 (1.03, 2.84)	0.89 (0.78, 0.94)	0.26 (0.01)
Task 2 Gain	1.30	0.96	0.28	1.32 (0.83, 2.23)	0.87 (0.79, 0.93)	0.30 (0.01)
Task 2 Loss	1.46	0.96	0.24	1.49 (0.91, 2.73)	0.86 (0.79, 0.93)	0.27 (0.01)

Note. Q_1 = first quartile, Q_3 = third quartile.

Was discounting in Task 1 affected by condition order in Task 1?

In Task 1, we replicated the steeper discounting in the first condition seen in Study 1, Experiment 2 and showed that discounting was steeper than that based on expected value in both the first and second conditions (individual and group median data are presented in Figure 2.11). A mixed measures ANOVA showed no significant main effects of order or condition, but a significant interaction between order and condition (see Table 2.14). The inset in the left panel of Figure 2.11 shows this interaction; when losses were presented first, gains (M = 0.14, SE = 0.05) were discounted significantly less steeply than losses (M = 0.32, SE = 0.05; p = 0.014), but when gains were presented first, there was no significant difference between discounting of gains and discounting of losses (p = .177; Bonferroni correction applied). The log(h) parameters were significantly higher than those describing choice based on expected value for losses in the loss-first order ($t(57) = 7.31$, $p < .001$, d_z = 0.96), gains in the loss-first order ($t(57) = 2.82$, $p = .007$, $d_z = 0.37$), losses in the gain-first order ($t(57) = 4.10$, $p < .001$, $d_z = 0.54$) and gains in the gain-first order ($t(57) = 6.53$, $p < .001$, $d_z = 0.86$; Bonferroni adjusted alpha of .0125).

The same analysis conducted with AUC confirmed the above ANOVA results, with the exception that the follow-up tests detected a significant difference between gains (M = 0.33, SE = 0.02) and losses (M = 0.22, SE = 0.01) when losses were first ($p < .001$), as well as between gains (M = 0.24, SE = 0.02) and losses (M = 0.29, SE = 0.01) when gains were first (p = .049; see Appendix D for AUC ANOVA results).

Table 2.14

Study 1, Experiment 3 ANOVA results using log(h)

	F	df	p	η_p^2	F	df	p	η_p^2
	Task 1				Task 2 (Task 2 order)			
Condition	0.63	1, 114	.428	.006	2.14	1, 114	.147	.018
Condition order	0.10		.754	.001	0.72		.400	.006
Condition by Condition order	**7.39**		**.008**	**.061**	0.28		.599	.002
					Task 2 (Task 1 order)			
Condition					2.11	1, 114	.149	.018
Condition order					0.001		.980	<.001
Condition by Condition order					3.05		.084	.026

Note. Bold emphasis added to significant results.

Was discounting in Task 2 affected by condition order in Task 2?

Discounting in Task 2 did not demonstrate the effect of order observed in Study 1, Experiment 2 and Task 1, but did show the same pattern of steeper discounting than that based on expected value (bottom panel, Figure 2.11). A mixed-measures ANOVA showed no significant main effects of order or condition, and no significant interaction when using log(h) or AUC. The inset in the bottom left panel and the overlapping curves in the right panel of Figure 2.11 demonstrate the non-significant interaction of order and condition. Log(h) parameters were significantly higher than those describing choice based on expected value for losses in the loss-first order ($t(54)$ = 4.31, p < .001, d_z = 0.58), losses in the gain-first order ($t(60)$ = 4.69, p < .001, d_z = 0.60) and gains in the gain-first order ($t(60)$ = 3.93, p < .001, d_z = 0.50), but not gains in the loss-first order ($t(54)$ = 2.38, p = .021, d_z = 0.32; Bonferroni adjusted alpha of .0125).

Was discounting in Task 2 affected by condition order in Task 1?

We examined whether discounting in Task 2 was related to condition order in Task 1 by running a mixed measures ANOVA, with condition in Task 2 as the within-subjects factor and order in Task 1 as the between-subjects factor. There was no significant main effect of condition, order, or interaction when using log(h) or AUC.

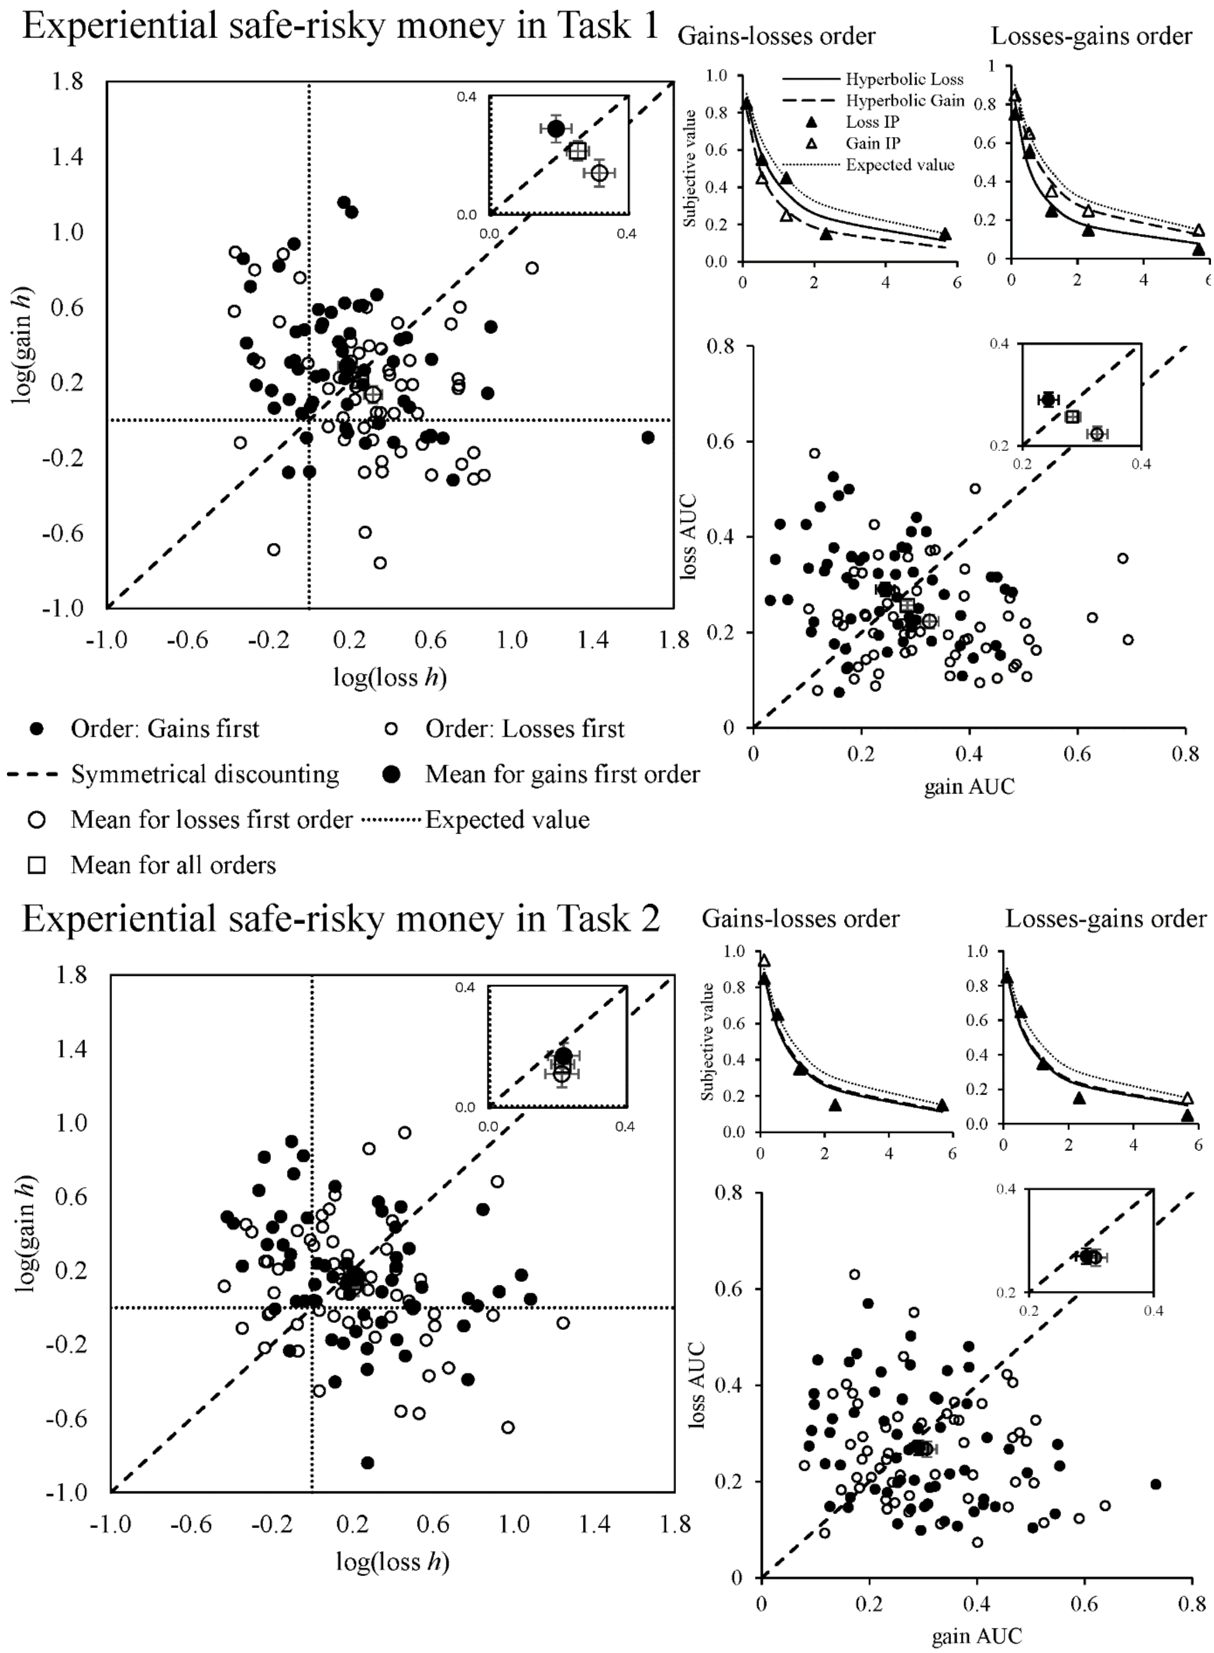

Figure 2.11. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual in the money (top half) and ski (bottom half) tasks, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and

white circles represent the means split by order, the white square the overall mean and the error bars are standard error of the mean. This is also shown in the inset. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged h value (0) when decisions are made based on expected value. Right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles) for money (top half) and ski (bottom half) tasks. For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Below the curves, AUC for losses is plotted against AUC for gains for each individual in the money (top half) and ski (bottom half) tasks. Group means and symmetrical discounting are as described for the left panel graph.

Did discounting of gains and losses change from Task 1 to Task 2?

Discounting of gains was more affected by order in the first task and task repetition, while discounting of losses was less affected by these factors. Examination of the top panel of Figure 2.12 showed that the majority of gains data points were below the diagonal line for log(h) on the left and above the diagonal line for AUC on the right, showing that participants generally had shallower discounting in the second gains condition they completed. A mixed measures ANOVA, with log(h) for gains as the dependent variable, task (first vs. second) as the within-subjects factor and order at Task 1 (gain-loss vs. loss-gain) as the between subjects factor, showed significant main effects of task (F(1, 114) = 9.47, p = .003, η_p^2 = .077) and order (F(1, 114) = 4.15, p = .044, η_p^2 = .035), and a non-significant interaction (F(1, 114) = 1.65, p = .202, η_p^2 = .014). Participants had higher logged h values for gains on the first task (M = 0.21, SE = 0.03) than on the second task (M = 0.14, SE = 0.03), regardless of the order of condition presentation on the first task. Participants also had higher logged h values for gains averaged across Task 1 and Task 2 when they completed the gains conditions first (M = 0.24, SE = 0.04) as compared to second (M = 0.12, SE = 0.04) in Task 1. The same analysis conducted based on AUC values showed a significant main effect of order (F(1, 114) = 5.28, p = .023, η_p^2 = .044) as above, but also a non-significant main effect of task (F(1, 114) = 3.50, p = .064, η_p^2 = .030) and a significant interaction (F(1, 114) = 14.31, p < .001, η_p^2 = .112). When gains were discounted first in Task 1, the mean gain AUC in Task 1 (M = 0.24, SE = 0.02) was significantly lower than in Task 2 (M = 0.29, SE = 0.02; p < .001). When losses were discounted first in Task 1, the mean gain AUC in Task 1 (M = 0.33, SE = 0.02) was not significantly different from Task 2 (M = 0.31, SE = 0.02; p = .179).

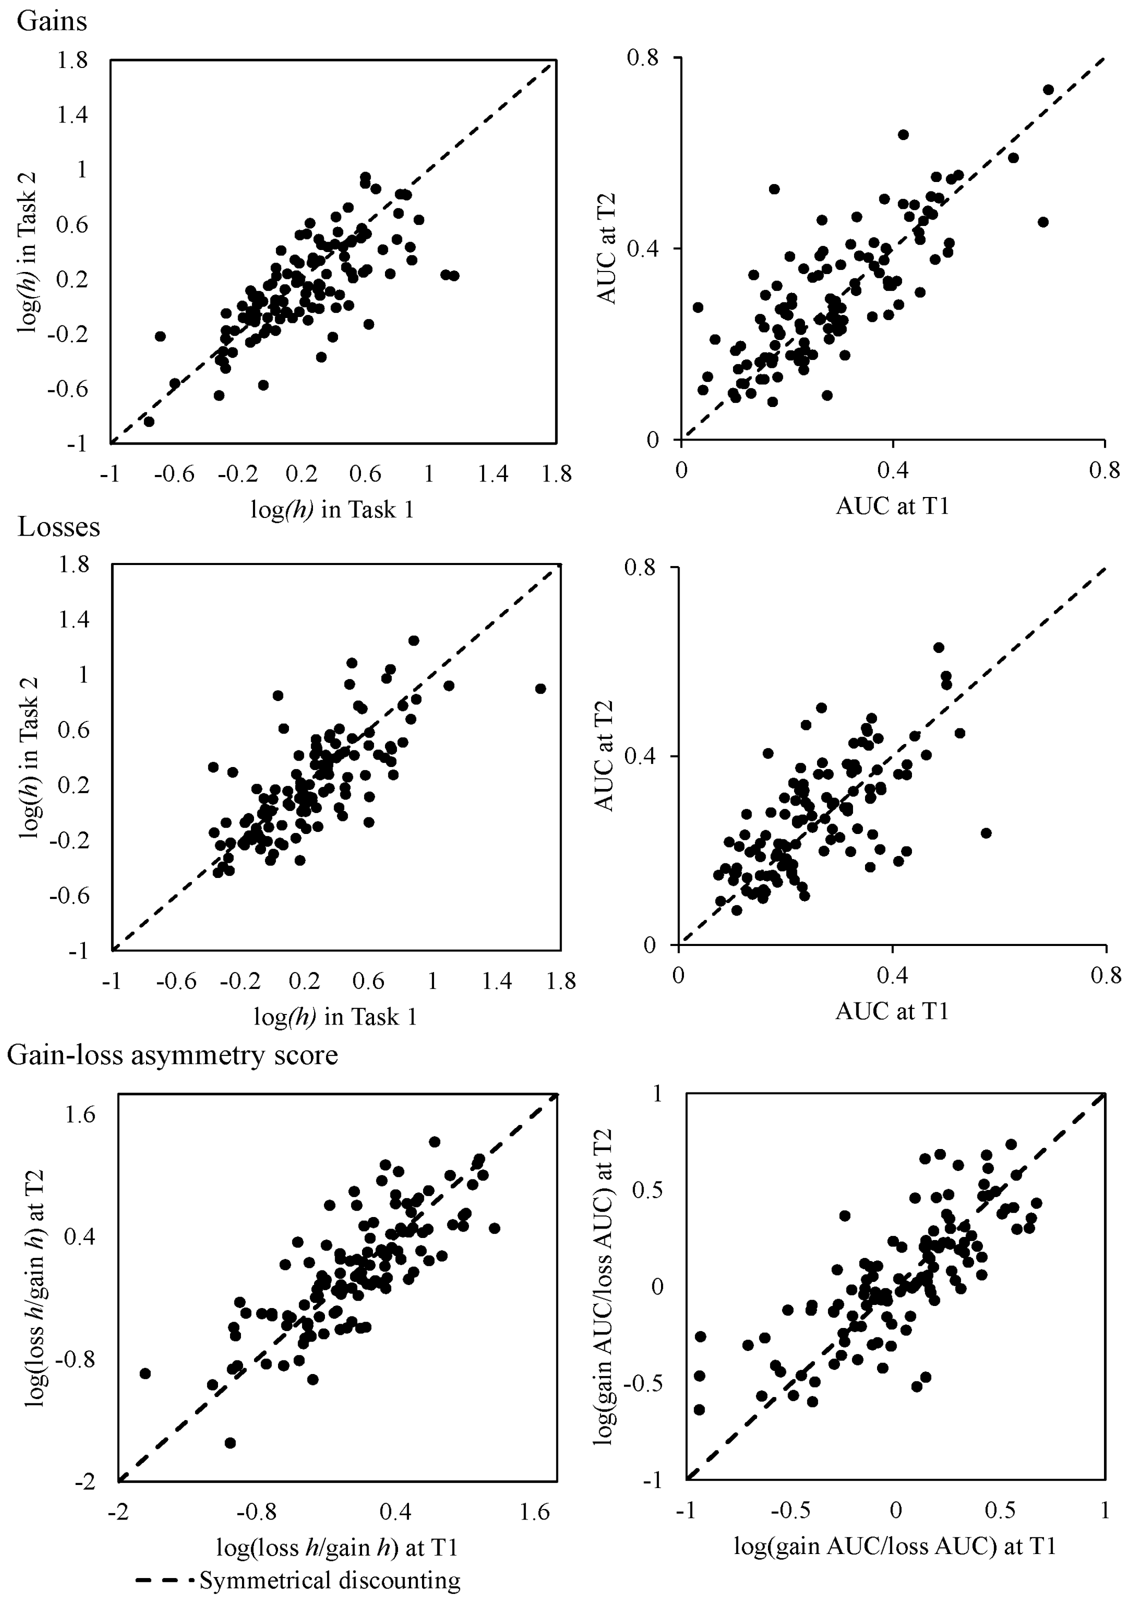

Figure 2.12. Top two rows: Log(h) (left) and AUC (right) parameters in Task 1 plotted against Task 2 for gains (top row) and losses (second row). The diagonal dashed line

represents identical *h* or AUC parameters across tasks. Third row: Gain-loss asymmetry score using log(*h*) on the left (log(loss *h*/gain *h*)) and AUC on the right (log(gain AUC/loss AUC) in Task 1 plotted against gain-loss asymmetry score in Task 2. The diagonal dashed line represents identical scores across tasks.

Examination of the middle panel of Figure 2.12 did not show any apparent trends in the discounting of losses based on log(*h*). There were no significant main effects of task ($F(1, 114) = 2.71$, $p = .103$, $\eta_p^2 = .023$), order ($F(1, 114) = 3.14$, $p = .079$, $\eta_p^2 = .027$) or interaction ($F(1, 114) = 0.63$, $p = .428$, $\eta_p^2 = .006$) for losses using log(*h*). The same analysis conducted on AUC values showed a non-significant main effect of task ($F(1, 114) = 2.14$, $p = .146$, $\eta_p^2 = .018$) and a non-significant interaction ($F(1, 114) = 3.89$, $p = .051$, $\eta_p^2 = .033$). There was, however, a significant effect of order ($F(1, 114) = 7.67$, $p = .007$, $\eta_p^2 = .063$); participants had lower AUCs for losses averaged across Task 1 and Task 2 when they completed losses first ($M = 0.24$, $SE = 0.01$) as compared to second ($M = 0.29$, $SE = 0.01$), in Task 1.

Correlations between tasks and conditions

Gain-loss asymmetry scores were strongly significantly correlated across the two repetitions of the task for log(*h*) ($r(114) = .79$, $p < .001$) and AUC ($r(114) = .77$, $p < .001$). The bottom panel of Figure 2.12 shows each participant's gain-loss asymmetry score for Task 1 plotted against their gain-loss asymmetry score for Task 2. Most participants' data points fall along the diagonal line, indicating that they demonstrated similar gain-loss asymmetry in Task 1 and Task 2. Log(*h*) values were significantly, positively correlated across tasks for losses and gains (see Table 2.15). Notably, discounting of gains and losses was significantly, negatively correlated in Task 1 and in Task 2. The same significant correlations were observed with AUC values.

Table 2.15

Study 1, Experiment 3 correlations for log(h) and AUC

	Log(*h*)	AUC
E2.1		
Gains Task 1 - Losses Task 1	-0.28*	-0.26*
Gains Task 2 - Losses Task 2	-0.27*	-0.23*
Gains Task 1 - Gains Task 2	0.73*	0.78*
Losses Task 1 - Losses Task 2	0.74*	0.68*

Note. *$p < .01$

Study 1, Experiment 3 aimed to establish whether the order effect observed in experiential discounting in Study 1, Experiment 2 was replicable in a larger sample. Discounting of gains and losses in Task 1 was determined by an interaction of condition order and type of outcome discounted as in Study 1, Experiment 2, with steeper discounting in the first condition of Task 1. However, discounting patterns in Task 2 were less clear. Discounting in Task 1 was associated with discounting in Task 2, yet condition order for either Task 1 or Task 2 did not significantly affect discounting rates in Task 2. Notably, the negative correlation between gains and losses within each task suggested that the degree to which individuals were risk averse or risk seeking (i.e. tendency to choose the risky option) was largely maintained across gains and losses. Lastly, group mean discounting of losses was more consistent across tasks, whereas discounting of gains was more affected by task context. Discounting of gains became shallower as the task progressed, while discounting of losses was not similarly affected.

Discussion

We examined gain-loss asymmetry in probability discounting across several contexts by contrasting safe-risky with risky-risky, as well as descriptive with experiential components. We have summarized the findings below in Table 2.16 by updating Table 2.2 with whether the predicted pattern was observed or not.

We expected a gain-loss asymmetry in probability discounting rate and that the direction of this asymmetry would be reversed in choice from description versus choice from experience. While we replicated the gain-loss asymmetry in choice from description in the safe-risky context (e.g. Estle et al., 2006), discounting of gains and losses was not significantly different in the risky-risky context, inconsistent with our predictions (see Budescu & Weiss, 1987 for similar results). In almost every experiential money task and one of the experiential ski tasks, discounting of experienced gains differed significantly from discounting of experienced losses. However, the direction of the difference between discounting curves was primarily determined by whether discounting was measured earlier or later in the session and not by whether choices were about gains or losses (see Hinvest & Anderson, 2010 and Matusiewicz et al., 2013 for other examples of order effects in probability discounting). Hence, although choice based on a combination of experience and description largely did not resemble choice based on description alone, it was also not consistent with our predictions based on the reversed reflection reported in the description-experience gap literature.

Table 2.16

Predictions and observed results of gain-loss asymmetry for the four possible context combinations in a probability discounting task

	Description	Tested in	Supported?	Experience	Tested in	Supported?
Safe-risky	Discounting of gains steeper than of losses	E1 money	Yes	Discounting of losses steeper than of gains	E2.1 ski E2.1 money E3 T1 money E3 T2 money	No, order No, order No, order No, null
	Both steeper than expected value-based curve	E1 money	Yes	Both shallower than/close to expected value-based curve	E2.1 ski E2.1 money E3 T1 money E3 T2 money	Yes No, steeper No, steeper No, mostly steeper
Risky-risky	Discounting of gains steeper than of losses	E1 money E2.2 money	No, null No, null	Discounting of losses steeper than of gains	E2.2 ski E2.3 ski E2.3 money	No, null No, null No, order
	Both steeper than/close to expected value-based curve	E1 money E2.2 money	Yes Yes	Both steeper than/close to expected value-based curve	E2.2 ski E2.3 ski E2.3 money	Yes Yes Yes

Note. Order = interaction between order and condition significant; Null = no significant results observed, null hypothesis not rejected. Steeper = discounting significantly steeper than choice based on expected value.

We expected discounting of gains and losses to be generally steeper than discounting based on expected value in contexts, except the safe-risky choice from experience context, where we expected generally shallower discounting. Our predictions based on the inverse S-shaped probability weighting function were supported in the safe-risky choice from description (Tversky & Kahneman, 1992), as well as in all of the risky-risky conditions (Wulff et al., 2018). Choice from experience in the safe-risky context largely did not support our predictions based on an S-shaped probability weighting function. Rather, we found the same patterns seen in other contexts of generally steeper discounting than choice based on the expected value, with the exception of data from one ski task.

Order effect as transition from description to experience

Studies comparing decisions from description with decisions from experience have found that decisions from experience are more risk averse for losses and more risk seeking for gains than decisions from description (Hertwig & Erev, 2009). In our experiential choice tasks, participants received information about the probabilities of different outcomes through both description *and* experience, accruing experience as the session progressed. It is possible that the shallower discounting we observed later in the session reflected a transition from choice controlled primarily by the descriptions of the probabilities of each outcome to choice controlled primarily by experience.

Research that has provided participants with both descriptions of and experience with the probabilities of chosen outcomes suggested that experience predominates, and decisions resemble those made from experience alone (Jessup et al., 2008; Lejarraga & Gonzalez, 2011). These findings suggested that including descriptions of the outcomes in our experiential task should not have significantly impacted participants' decisions. In our experiments we observed that participants' behaviour changed after they experienced some number of probabilistic outcomes, consistent with the research showing dominating effect of experience. However, the observed order effects also meant that our participants have not had sufficient trials to experience probability and the impact of the initial description likely remained. Therefore, the inconsistent direction of gain-loss asymmetry indicated that we captured their preferences while they were transitioning from decision from description to decision from experience.

Jessup and colleagues (2008) observed similar behaviour in description-only and description with experience conditions in the first block of ten trials, but by the second block participants started showing a relatively consistent description-experience-gap. Their participants also had several practice trials before the actual experiment. Although our participants experienced thirty trials in each condition, these trials were divided among the five probabilities and the experience of each probability for each condition would at most reach six by completion. Study 1, Experiment 3 repeated the conditions, so that by the end of Task 2 the participants would have experienced sixty trials of gains and losses each. Repeating the task twice also did not guarantee sufficient exposure to probabilistic outcomes, as the participants had, at most twelve trials for each probability within these conditions. Consequently, we may have measured discounting rate earlier in the transition while comparable research has sampled later (Jessup et al., 2008; Lejarraga & Gonzalez, 2011). Consistent with this idea, Wulff et al. (2018) found that the number of probabilistic outcomes participants sampled when able to try out the chance option was predictive of the size of the description-experience gap.

Discounting of gains more susceptible to context than discounting of losses

Discounting of gains was less stable across the sessions than discounting of losses. When comparing discounting across task administrations in Study 1, Experiment 3 and controlling for order of task presentation, there was no significant difference between discounting of losses in Task 1 and Task 2 based on log(*h*) values. However, a similar comparison showed that discounting of gains in the first task was steeper than discounting of gains on the second task, regardless of task order in the first task. Analysis using AUC showed some inconsistencies with the log(*h*) results. However, given the good fits of the hyperbolic model in Study 1, Experiment 3, the validity of the log(*h*) results appears to be high and AUC results were likely capturing noise in the data. Furthermore, this bears similarities to the differential effect of magnitude in literature, where discounting of gains was more susceptible to changes in magnitude than discounting of losses (e.g. Estle et al., 2006; Green et al., 2014). Our findings suggested that certain features of the experiential context, which were more salient in the money than the ski tasks, affected the discounting of gains to a greater extent than discounting of losses.

Consistency across task administration

While the order effect resulted in an inconsistent direction of discounting, we can comment on some properties of experiential discounting that appeared to be consistent across repetitions and task contexts. We observed significant, negative correlations between gains and losses in Study 1, Experiment 3 based on both log(*h*) and AUC values. This negative relationship supported the work of Shead and Hodgins (2009), who also observed that an individual’s tendency to choose the larger, uncertain option was relatively consistent across gains and losses. A positive correlation between the logged ratios of losses over gains in Task 1 and Task 2 of Experiment 3 suggested a degree of consistency in each individual’s level of gain-loss asymmetry. For most individuals, if gains were discounted more steeply than losses in Task 1, they were also discounted more steeply than losses in Task 2 and vice versa. Furthermore, across all of the three experiments, gains and losses on one task were positively correlated with gains and losses respectively on the other task, affirming a cross-context similarity in discounting rate as seen in the literature (e.g. Green et al., 2014; Greenhow et al., 2015). This within-individual stability may reflect the impact of each participant's idiosyncratic pre-experimental history with probabilistic outcomes.

Discounting relative to choice based on the expected value

We also compared participants' discounting rates to those based on the expected value. Prospect Theory predicts risk aversion for gains and risk seeking for losses in most participants for descriptive choice, which corresponded to the pattern we observed despite the inclusion of experience. The generally steeper discounting relative to choice based on the expected value in the risky-risky conditions also supported the findings of a generally inverse S-shaped probability weighting function in risky-risky contexts (Wulff et al. 2018). Wulff et al. also observed that experience was associated with more optimal choice, that is, a higher proportion of choices maximizing the mean returns. Although participants demonstrated shallower discounting in the second condition they experienced, bringing discounting closer to that based on the expected value, discounting remained significantly steeper than that predicted by expected value in most tasks.

Notably, in selecting a range of probabilities to observe the rate of change in subjective value, the studies of gain-loss asymmetry in probability discounting have generally chosen more probabilities in the lower and higher range, rather than the middle (approximately 25-75%). While this constituted a distribution that was a roughly equal sampling of lower and higher probabilities within each study, the conversion of probabilities to odds against does not result in an equal distribution of points along the x axis. Figure 2.1 shows that probabilities at 0.50 and above correspond to a smaller portion of the curve, while probabilities below 0.50 correspond to a larger portion of the curve. Hence, lower probability events were likely to have a higher influence on the parameter estimates. We chose our probabilities based on the range we have piloted in my honours thesis, but in retrospect this resulted in a heavier sampling of the middle range relative to the tail-ends of the probabilities, opposite to that in the literature (Shead & Hodgins, 2009; Estle et al., 2006; Weatherly & Derenne, 2013; Mitchell & Wilson, 2010). Based on the probability weighting functions we have plotted in Figures 1.2 and 1.3, the middle range of probabilities remains roughly unchanged with the reversal of the S-shaped function, with underweighting in both choice from description and choice from experience. We speculate that the relative consistency of steeper curves, corresponding to more risk seeking for losses and risk averse behaviour for gains than expected value would predict, was due to the heavier sampling of the middle range of probabilities. This is speculative because there is no definitive probability weighting function in either the descriptive or the experiential literature (see Wulff et al. 2018 for a review), but does suggest further refinements to discounting tasks that should include careful mapping of a broader range of probabilities.

The effect of magnitude on the gain-loss asymmetry

We considered whether our choice of magnitude of the larger outcome diminished a possible gain-loss asymmetry. Studies that have manipulated the magnitude of the probabilistic option observed that smaller probabilistic gains were discounted less steeply than larger probabilistic gains, while discounting of losses was largely unaffected by change in magnitude of the probabilistic option (Estle et al., 2006; Mitchell & Wilson, 2010; but see also Weatherly & Derenne, 2013). Therefore, a more pronounced gain-loss asymmetry is expected at higher magnitudes, but the magnitude at which no gain-loss asymmetry is observed was inconsistent across studies (e.g. Mitchell & Wilson, 2010). However, our choice of $100 was within the range in which no gain-loss asymmetry has been observed in some previous studies (e.g. Estle et al., 2006).

We also considered the magnitude of the on-screen balance as an additional aspect of our procedure that could have affected discounting. The balance varied from the first to the second condition as participants lost or gained money. Using Task 1 in Study 1, Experiment 3 as an example, the first condition had a starting balance of $3000. In the losses-gains order the balance had decreased to $1739, on average, by the start of the second condition but in the gains-losses order it the balance had increased to $4477, on average, by the start of the second condition. The balance might affect the subjective, undiscounted value of the larger, uncertain alternative, and therefore affect discounting via the magnitude effect. Research suggests that larger gains are discounted more steeply but the magnitude of losses does not affect discounting rate (e.g. Estle et al., 2006), and that a similar effect occurs when the relative value of gains is manipulated by altering other rewards in the context (Dai, Grace & Kemp, 2009).

The possible effect of the on-screen balance on the subjective magnitude of the outcomes likely does not account for the order effect observed, for several reasons. Firstly, the balance changed in the opposite direction for each order, yet the direction of change in discounting was the same for each. Secondly, the magnitude effect would predict steeper discounting of gains in the second condition in the loss-gain order, and no effect of order on losses, which is inconsistent with the patterns we observed. Thirdly, we did not observe an order effect in the second task in Study 1, Experiment 3, even while the balance fluctuated in the same manner as in conditions with consistent order effects, and Hinvest and Anderson (2010) observed an order effect in the absence of an on-screen balance. These aspects of our results suggest that balance is an unlikely cause of order effects, although worthy of further investigation.

To conclude, we would like to return to the trade-off between amount and probability when considering medical insurance. The context for such decisions is usually descriptive rather than experiential; the individual may not have feedback on their initial decision for many years. Previous research suggested that, in this type of decision context, the subjective value of gains and losses for an individual are likely to be asymmetric. However, decisions made from a combination of description and experience - such as about whether to take a daily medication producing side effects - are more likely to change across time as the individual accumulates experience with the outcome of their choices. Specifically, the current study suggests that with more experience, the probabilistic option would become subjectively more valuable to the decision maker while, and at the same time, the direction of asymmetry between gains and losses is preserved for each individual. This effect was also more likely in the safe-risky, rather than risky-risky contexts. Lastly, decisions involving losses might be less susceptible to contextual changes than decisions involving gains. The set of experiments in Study 1 point to the importance of methodological diversity to develop a full picture of the drivers of choice.

Study 2: Gain-loss asymmetry in a concurrent schedules task

Study 1 explored the extent to which a gain-loss asymmetry was observed in probability discounting when the choice scenario was experiential, risky-risky and a combination of the two. Contrary to our initial predictions of a reverse gain-loss asymmetry when choice was made experiential, we did not observe a consistent difference in discounting of gains versus losses in safe-risky or risky-risky tasks. Instead, order of the conditions interacted with the type of outcome being discounted. This significant interaction was consistent across most of the experiential money tasks, as seen in Experiments 2.1, 2.3 and 3 Task 1, and less consistent in the ski tasks, only seen in Experiment 2.1. In Experiment 3, Task 1, when the participants discounted losses first, losses were discounted significantly more steeply than gains, a finding consistent across both log(h) and AUC analysis. When the participants discounted gains first, the difference between gains and losses did not appear to be large, with a significant difference when using AUC but not log(h). This pattern suggested that both discounting of losses and discounting earlier in the task produced steeper curves; the weaker effect in the gains first order was likely due to these factors working in opposition.

The gain-loss asymmetry appeared to be task order dependent. However, it is not clear which features of the task being experienced by the participants contributed to this change in discounting. First, progression through each condition involved not only increasing experience of probabilistic events, but also the magnitude and probability of each option varied, the separate contributions of which is not clear in a titrating-amount discounting procedure. Second, Experiment 3, Task 2 further demonstrated that discounting of gains was more susceptible to task repetition, generally becoming shallower, but the discounting of losses was not affected to the same extent. This suggested different rates of learning about probabilistic choices involving gains as opposed to losses.

In Study 2, we used a theoretically and procedurally related concurrent-schedules task to continue our examination of gain-loss asymmetry in experiential choice and to further build upon the above findings. In a concurrent schedules task participants choose between two alternatives. One alternative could be associated with lower and the other with higher frequency of reinforcement, but participants typically do not know the exact schedules (i.e. no description is provided) and are encouraged to sample the alternatives and experience the outcomes. For example, in a concurrent schedules task using variable-intervals, the left alternative could be scheduled to provide a reinforcer on average every seven seconds (VI7s) and the right on average every fourteen seconds (VI14s). After sampling each option and in order to maximize the number of reinforcers, most participants would be expected to respond more on the left relative to the right. The participants experience a range of such schedules to track how their response changes when we change some aspect of reinforcer delivery. Fundamentally, discounting and concurrent schedules tasks quantify similar decision-making behaviour (Rachlin, 2006). Both demonstrate that discounting of the subjective value of an outcome occurs as a function of changes to its probability. In probability discounting, the value of an outcome is discounted by its probability. In concurrent schedules tasks using variable intervals, the value of reinforcers is discounted due to the lower probability of receiving them, with greater loss in value on the leaner (e.g. VI14s) than the richer (e.g. VI7s) side. This difference in value is demonstrated by the fact that participants typically allocate more responses to the richer side.

In our experiential discounting task, we observed an effect of condition order, presumably due to a transition in behaviour informed by description to behaviour that is also informed by experience, but the nature of the titrating procedure meant we could not isolate choice on later trials from choice on earlier trials. Concurrent schedule tasks allow the participant to learn the associated contingencies through repeated sampling and preferences are usually estimated from the final few trials; unlike in a discounting procedure learning/early trials do not directly contribute to this final estimate. Concurrent schedule tasks also allow for gain-loss asymmetry to be examined at both an extended and a local level of analysis (Baum, 2002, 2003). An extended level of analysis describes aggregated patterns of behaviour over extended exposure to changes in reinforcers, while a local level of analysis describes behaviour immediately after reinforcer delivery, which tends to be transient. While the former is possible in both tasks, the latter is not possible in our discounting task procedure.

The probability discounting procedure we used in Study 1 and concurrent schedule tasks both allow for an extended level of analysis. The titrating amount procedure used thirty trials to calculate indifference points and subsequently the rate at which the outcome lost its subjective value with changes in its probability. We could not derive a measure of discounting rate from fewer than thirty trials as it took at least six trials for the participant to have the option to reach either the lowest ($0) or the highest ($100) possible smaller certain value. Furthermore, taking their choices on any given trial in isolation would not provide us with a measure of rate: change in subjective value as a function of increasing odds against. In a concurrent schedules task, participant's choices are aggregated across blocks of choices and can also be expressed as a rate of change in behaviour allocation in response to changes in reinforcement rate. Both procedures can thus provide a measure of gain-loss asymmetry in the form of an extended analysis of the rate at which an outcome loses its subjective value with change in some dimension, such as probability or magnitude. This also lends itself to the possibility of testing for shared variance between these extended level measures across the two types of tasks.

While both procedures allow for an extended level of analysis, our titrating amount procedure does not permit local level of analysis of the effect of experiencing an outcome on the next choice. In our discounting task, we were unable to isolate the effect of choosing the risky option and the experience of its outcome at a given probability in one trial on the chances of choosing the risky option again on the next trial for the same probability. The titrating amount procedure systematically varied both amount of the smaller certain and the probability of the larger uncertain from trial to trial. This meant that changes to the amounts of the smaller, certain option, the probability of the larger, uncertain option, the number of choices an individual has made, and how often the probabilistic outcome, if selected, produced a gain or a loss could not be effectively separated. The generally shallower discounting with task progression, especially with gains as compared to losses, suggested that some of the above changes from trial to trial had varying local effects on the subsequent likelihood of choosing the larger, risky option. While we were unable to separate these effects out, concurrent schedule tasks could be used to vary each factor systematically to isolate the effect of accruing experience with probabilistic gains and losses, thus allowing for such local effects to be studied separately from the extended effect of reinforcers on behaviour.

Finally, in applying the literature from the description-experience gap to our discounting procedure, we noted that the options in our experiential task necessitated descriptions, unlike most of the experiential non-discounting tasks in choice literature. Existing literature suggested that including descriptions along with experience would show patterns similar to experience alone (e.g. Jessup et al., 2008), but we may have underestimated how many trials were required to reach this point. A concurrent schedule task avoids this issue, as participants can only rely on experience from start to finish.

In Study 2, we examined the gain-loss asymmetry in experiential choice by using a concurrent schedules task. Gain-loss asymmetry based on the predictions of Prospect Theory has received limited attention in concurrent schedules tasks, and to our knowledge the description-experience gap literature has not been applied directly to concurrent-schedule tasks. We conducted both an extended and local analysis, a distinction that the data from the discounting procedure suggested might be an important part of the observed gain-loss asymmetry. Participants completed the Auckland Card Task, a rapid-acquisition concurrent schedules task. In addition, we also continued our examination of the interaction of condition order and type of outcome being discounted in experiential probability discounting by including an experiential discounting task in which balance was held constant across conditions. Finally, since each participant completed both the concurrent schedules task and the probability discounting task, we examined for any cross-task similarities at the extended level of analysis.

Gain-loss asymmetry at extended level of analysis

The behaviour on concurrent schedule tasks described above has been formalized as a quantitative principle in human choice literature. The *strict matching law* maintains that the distribution of responses in a given environment depends on all of the accessible contingencies (Herrnstein, 1961; 1970). If, at one source, reinforcers are available at a higher rate or magnitude, we would expect the individual to allocate more of their responses to this alternative relative to other. The number of responses on each alternative is recorded and is expected to follow (approximately match) the share of reinforcers the participants experienced, so that a VI7s alternative would result in roughly twice as many responses allocated to it than a VI14s one. Therefore, the strict matching law describes a pattern of responding wholly reflecting the objective probabilities of the outcomes.

The *generalized matching law* (GML) was built on this assumption, but also incorporated deviations from strict matching observed in subjects: bias to either alternative for reasons other than reinforcement rate, and a tendency towards behaviour allocation other than strict matching (Baum, 1974). It is described by the generalized matching equation:

$$\text{Log}\left(\frac{B_1}{B_2}\right) = a\,\text{Log}\left(\frac{R_1}{R_2}\right) + \text{Log}\,c \qquad \text{(Equation 3.1)}$$

B_1 and B_2 are the number of responses on each alternative. In procedures where the frequency of reinforcers that can be obtained in a given time period differs between the two alternatives, R_1 and R_2 are the number of reinforcers obtained by responding on each alternative. In procedures where the size or magnitude of the reinforcers is manipulated instead, R_1 and R_2 correspond to amount obtained by responding on each alternative. The response and reinforcement rate are each expressed as a logged ratio and plotted on the y and x axis respectively (see the black data points in Figure 3.1 sample data). If the participants behave in accordance with the strict matching law, where distribution of responses closely follows the distribution of reinforcers, the data would be along the dashed line. Equation 3.1 can then be fitted to those data using least-squares linear regression, producing a line of best fit to the reinforcer and response ratios (the solid line).

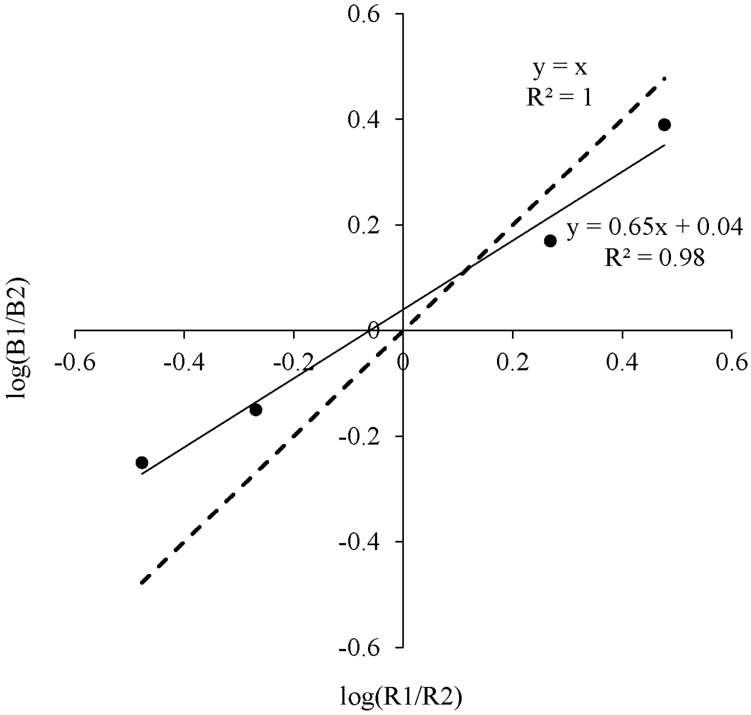

Figure 3.1. Logged ratio of responses on the y axis plotted against the logged ratio of reinforcers on the x axis. Dashed line corresponds to the predictions of the strict matching law. The solid line is the Generalized Matching Law fitted to hypothetical data points. The equations of the regression line of best fit and the R_2 are shown in the top right of the graph.

Our sample solid line data shows the common deviations from strict matching captured by the sensitivity (a) and bias (log c) parameters in the GML. a is a measure of sensitivity of the allocation of responses to the given reinforcer ratio and is the slope of the fitted GML to the data (Lobb & Davison, 1975; Davison & McCarthy, 1988), where a value of 1 indicates that the relative rate of responding matches the relative rate of reinforcement. a less than 1 indicates undermatching, where the distribution of responses across the two alternatives is less extreme than matching would predict and the organism is less sensitive to the difference between the two alternatives. Our sample data shows undermatching, with an a of 0.65. a greater than 1 is overmatching, where the distribution of responses across the two alternatives is more extreme than matching would predict. Here, the organism is more sensitive to the difference between the two alternatives. Log c (y-intercept) is a measure of bias that the participant may show for one alternative independent of the reinforcer ratios (McDowell, 1989). A bias in responding is indicative of more responding on one alternative than predicted across all of the reinforcer ratios. If log c is less than 0, then the participant is showing bias towards the R_2 side, and if it's more than 0, then the bias is towards the R_1 side. Our sample data are showing a slight bias towards the R_1 side, with a log c of 0.04. Note than when $a = 1$ and log $c = 0$, the relative rate of responses equals the relative rate of reinforcers and the individual is showing no systematic deviations; behaviour instead conforms to the strict matching law as shown by the dashed line.

In concurrent schedules task using variable intervals, Prospect Theory and the description-experience gap literature would predict that changes in some aspect of a loss outcome (e.g. its frequency or magnitude, or even its presence or absence) between the two alternatives would produce an allocation of behaviour that is not symmetrical to an equivalent change in some aspect of a gain outcome. The description-experience gap further predicts a reversal of the reflection effect, with greater risk seeking for gains than losses at most probabilities. The GML-based method of quantifying behaviour relative to reinforcement focuses on the rate at which behaviour changes after some measure of stability in responding has been achieved. In other words, the GML derives a measure of sensitivity to changes in the environment and bias based on a steady-state pattern of responding achieved over several sessions.

Although sensitivity is not a direct measure of risk taking, we can draw a link between the probability weighting function and GML, in that they both can capture a deviation in responses from the expected pattern based on the objective probability of an outcome. To our knowledge, these ways of measuring responses have not been synthesized together, however we suggest that both approaches describe similar patterns of behaviour. Figure 3.1 can be understood in the same manner as Figures 1.2 and 1.3. The x axis corresponds to expected or predicted behaviour, based on objective probability in Figures 1.2-3 and on the ratio of the chances of attaining a reinforcer in Figure 3.1, and the y axis plots the observed responses which may or may not correspond to the predicted values. Our sample data on Figure 3.1 shows the typically observed deviation (undermatching) from a linear response to the probability of an outcome in the format typically used to present data in studies using the generalized matching law.

Plotting the same data in an alternate format makes the similarities to the Prospect Theory probability weighting function more apparent. The change in ratios that results in undermatching is demonstrated in the top left graph of Figure 3.2, which plots the proportion of reinforcers on the x axis and the proportion of responses on the y axis according to Herrnstein's original strict matching law formulation without the log transformation of the GML[4]. The probability weighting function from Prospect Theory is reproduced in the top right graph to facilitate comparison. The left panel is an example of undermatching, that is, the participant is underestimating the likelihood of a reinforcer on R_1 when R_1 has a higher probability of reinforcement (right side of the graph) and on R_2 when R_2 has the higher probability of reinforcement (left side of the graph). Our participant also overestimates the likelihood of a reinforcer on R_2 when R_2 has a lower probability of reinforcement (right side of the graph) and on R_1 when R_1 has the lower probability of reinforcement (left side of the graph). The reflection effect based on descriptive choice predicts a very similar pattern for gains, with risk aversion at moderate to high probabilities and risk seeking at low probabilities (i.e. the inverse S-shaped probability weighting function on the right graph). We can interpret this as Prospect theory predicting undermatching for gains, where a high probability gain elicits fewer and a low probability gain elicits more responses than perfect matching would predict.

We can apply the same logic to losses, with the reflection effect predicting risk seeking at moderate to high probabilities and risk aversion at low probabilities. If we were to plot the left graph of Figure 3.2 such that the reinforcer ratio would instead contrast likelihood of no-loss on each alternative, then we would expect more responses on the side with the higher likelihood of no-loss, thus producing a positive slope. For example, the rightmost data point on the x axis of the left graph would correspond to a greater proportion of no-losses (75 at R_1) at the R_1 side. Since the reflection effect predicts risk seeking at moderate to high probabilities of a loss (lower likelihood of no-loss) and risk aversion at low probabilities of a loss (higher likelihood of no-loss), this corresponds to more responses on the side with lower likelihood of no-loss and fewer responses on the side with a higher likelihood of no-loss than perfect matching would predict. Thus, if we followed Prospect Theory's predictions based on choice from description, we would expect undermatching for both gains and losses. Our Figure 1.2 (reproduced in Figure 3.2) also includes greater

[4] Herrnstein's matching law (1961): $\left(\frac{B_1}{B_1 + B_2}\right) = \left(\frac{R_1}{R_1 + R_2}\right)$

deviation from linear weighting for gains rather than losses, which has sometimes been observed in literature (e.g. Abdellaoui et al., 2011; Hau et al., 2008), corresponding to a greater degree of undermatching for gains rather than losses. Thus the final predictions based solely on Prospect Theory are shown in the left graph of Figure 3.3 using hypothetical data based on the GML.

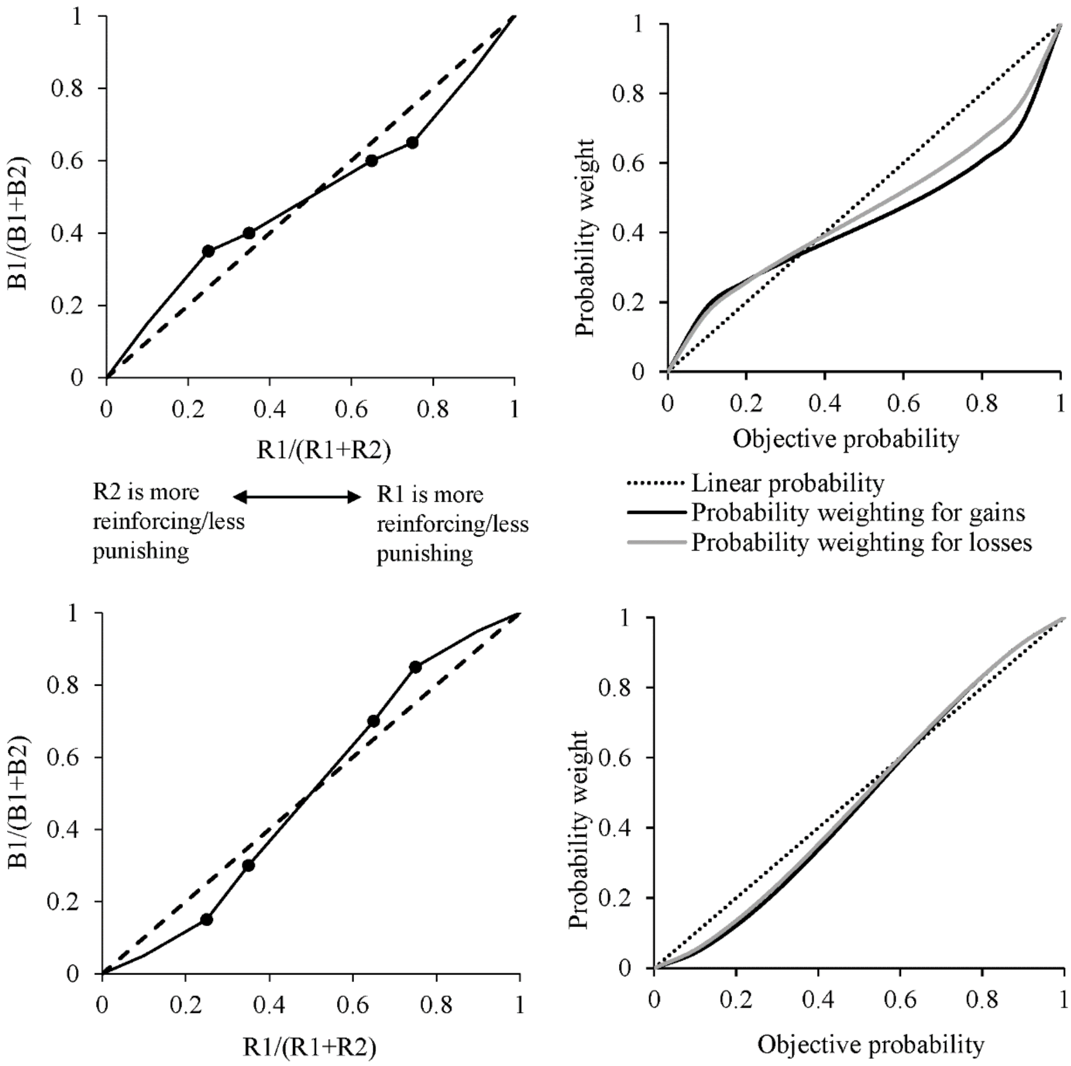

Figure 3.2. Left panel: Proportion of responses at B1 relative to total responses plotted against proportion of reinforcers at R1 relative to total reinforcers. In each graph, the black data points and the curves are hypothetical data showing undermatching (top) and overmatching (bottom) relative to the diagonal dashed line indicating perfect matching. Right panel: Figure 1.2 (top) and 1.3 (bottom) reproduced from the general introduction.

Prospect Theory describes choice from description, therefore, these predictions of undermatching derived from it also apply only to choice from description. The description-experience gap suggests that choice from experience might show a reversal of the pattern of responses seen in choice from description. Since our concurrent schedules task is to be fully experiential to avoid any complexities with including description, our predictions need to be based on choice from experience as modelled by the bottom right graph of Figure 3.2 (reproduced from Figure 1.3). Here, research into the description-experience gap has observed risk seeking for gains at moderate to high probabilities and risk aversion at small probabilities. This pattern would correspond to overmatching (see the bottom left graph of Figure 3.2), where participants allocate more responses to the side with a high chance of gain and fewer to the side with the lower chance of gain than perfect matching.

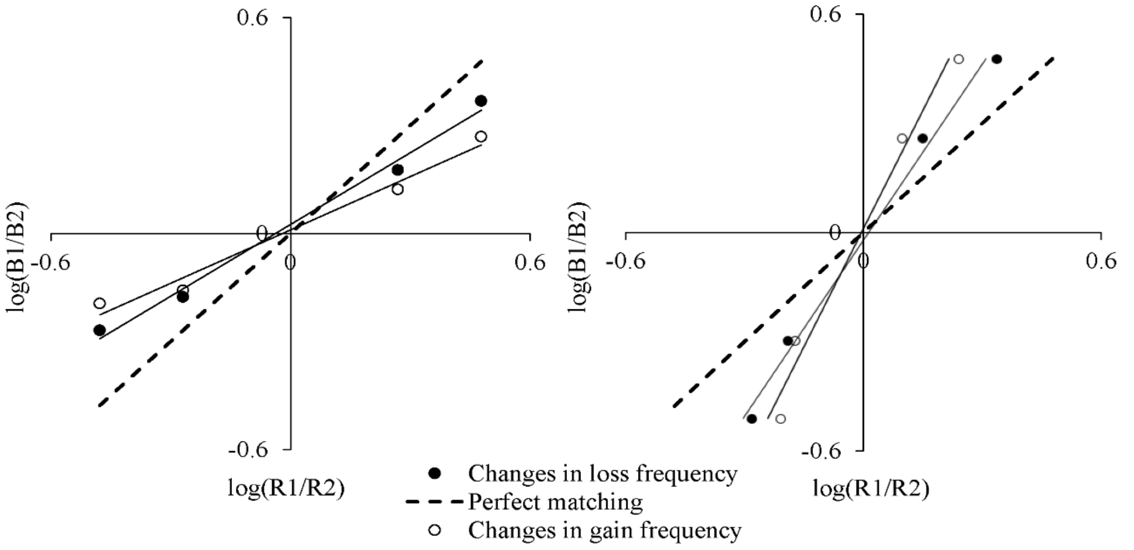

Figure 3.3. Logged ratio of responses on the y axis plotted against the logged ratio of reinforcers on the x axis. Dashed line corresponds to the predictions of the strict matching law. The solid line is the Generalized Matching Law fitted to hypothetical data points for the gain (white) and loss (black) conditions. The equations of the regression line of best fit and the R_2 are shown in the top right of the graph. The left graph corresponds to Prospect Theory's predictions for choice from description. The right graph corresponds to description-experience gap literature's predictions for choice from experience.

For losses, the reversed reflection effect would similarly predict overmatching, with more responses to the side with a high chance of no-loss and fewer responses to the side with a low chance of no-loss than perfect matching would predict. Furthermore, if we were to apply the greater deviation from linear probability in losses modelled in Figure 1.3

(reproduced in the bottom right graph of Figure 3.2), then we would expect a greater degree of overmatching with gains than losses. Thus the final prediction based on the reversed reflection that has previously been observed in choice from experience is shown in the right panel of Figure 3.3. Ideally, the extent of deviation from perfect matching in our hypothetical data would be plotted such that it corresponds closely to the parameters of the probability weighting function, but since there is little agreement in literature on the exact value of these parameters, we arbitrarily chose the extent of undermatching or overmatching.

Our application of the description-experience gap literature was so far limited to manipulation of frequency. If we were to consider behaviour in response to change in magnitude, where the frequency of a reinforcer is kept constant between the two alternatives, we would rely solely on the value function of Prospect Theory in making our predictions (see in Figure 1.1). Any probability weighting would now be held constant between the two alternatives. Given the steeper change in the value of a loss as compared to a gain with equivalent changes in magnitude, the value function of Prospect Theory would predict greater sensitivity to changes in loss magnitude rather than gain magnitude. Whether response allocation would be expected to undermatch or overmatch overall is unclear.

Thus, based on Prospect Theory and description-experience gap literature, we would expect overmatching to rate of gains and losses. However, the majority of literature does not correspond to the theoretical predictions we have outlined above and undermatching is virtually universal. In non-human animal subjects, the typical behaviour in response to varied gain frequency is slight undermatching and to gain magnitude is greater undermatching (e.g. Landon, Davison & Elliffe, 2003; see Davison & McCarthy, 1988 for a review and Elliffe et al., 2008 for a discussion). Human performance on such tasks has generally shown less sensitivity to differences in reinforcement than animal subjects, with undermatching in response to changes in gain frequency and magnitude (Kollins, Newland, & Critchfield, 1997; Schmitt, 1974; Wurster & Griffiths, 1979), although the extent of deviation from strict matching varies with procedural differences, such as whether rates of reinforcement are signalled by a discriminative stimuli (e.g. Bradshaw, Szabadi & Bevan, 1976; see Horne & Lowe, 1993 for a discussion).

The effect of introducing losses or punishers to concurrent VI schedules compared to gains has been less studied in human participants. In non-human animal subjects, reinforcers and punishers have been traditionally operationalized as delivery of food versus shock

(Deluty, 1976; de Villiers, 1980; Farley, 1980). Note that when studying the effects of punishers it is necessary to also reinforce some responses so that responding does not stop altogether. Researchers who have investigated the effects of introducing losses or punishers have examined whether effects are best understood as *subtracting* from the alternative on which punishes occur or *adding* to the other, non-punished alternative. That is, subtractive model states that punishers subtract from the reinforcing properties of the reinforcers (de Villiers, 1980), in contrast to the predictions of the additive model which assumed that punishers would increase responding on the opposite alternative (Deluty, 1976).

Results support the subtractive model of punishment. When equal frequency shock was introduced to two response keys with varying reinforcement rates, de Villiers (1980) observed that punishment had a subtractive effect on reinforcer-maintained responding. Pigeons overmatched in their distribution of responses to the distribution of reinforcers relative to responding for reinforcement only, showing an extreme preference for the alternative with higher relative reinforcement rate. This was based on a matching law formulation that did not include the systematic deviations captured by the GML that are known to be common. Critchfield, Paletz, MacAleese and Newland (2003) addressed this concern and measured performance of human subjects in response to monetary gains and losses using the GML. Across their experiments, participants responded on a two-alternative variable interval reinforcement schedule. Some of the conditions also delivered punishment on a less frequent schedule than the reinforcement. They too found support for the subtractive effect of punishment (termed the *direct-suppression model*), with participants showing increased preference for the alternative that delivered more frequent reinforcers when both alternatives were punished at an equal rate (see Experiment 2A and 2C).

Critchfield et al. (2003) further expressed the subtractive effect of punishers by creating a direct-suppression version of the GML (Equation 3.2):

$$\text{Log}\left(\frac{B_1}{B_2}\right) = a\ \text{Log}\left(\frac{R_1 - P_1}{R_2 - P_2}\right) + \text{Log}\ c \qquad \text{(Equation 3.2)}$$

P_1 and P_2 are the number of punishers (or amount, in case of changes in magnitude) obtained by responding on each alternative and are directly subtracted from the number of reinforcers obtained on the corresponding alternatives. The predictions of this model are demonstrated in Table 3.1 using a set of schedules in one of the conditions used for our experiment (details in the Method section), where the *R* and *P* values indicate the amounts of reinforcers and punishers delivered within the condition. In both of the examples, equal punishment is introduced to two alternatives with varying rate of reinforcement. If punishers are not incorporated into the model and we expect behaviour change in response to change in reinforcement rate only, then we would expect the response ratio to correspond to the reinforcer ratio in column 3. If punishers are incorporated into the reinforcer ratio (termed *net* ratio) and if they have a subtractive effect according to the direct-suppression model, then we would expect the response ratio to correspond to the net ratio in column 4. This would predict overmatching relative to the reinforcer only ratio. Notably, this model cannot accommodate instances where more punishers than reinforcers were delivered.

Table 3.1

Predictions of the direct suppression GML using a set of schedules

	$\left(\frac{R_1 - P_1}{R_2 - P_2}\right)$	Reinforcer ratio $R_1 : R_2$	Net ratio $(R_1 - P_1) : (R_2 - P_2)$
Example 1	$\left(\frac{250 - 150}{750 - 150}\right) = \frac{100}{600}$	1:3	1:6
Example 2	$\left(\frac{350 - 150}{650 - 150}\right) = \frac{200}{500}$	1:1.86	1:2.5

Using the parameters generated by the GML as well as the direct-suppression GML, a few studies have directly investigated a gain-loss asymmetry in concurrent schedules tasks. A gain-loss asymmetry was quantified as either a bias towards one alternative with the introduction of losses to the other (Magoon & Critchfield, 2008; Rasmussen & Newland, 2008) or a difference in sensitivity with the introduction of losses (Rasmussen & Newland, 2008). Notably, if Prospect Theory was mentioned by these studies, it was exclusively as the application of loss aversion observed in choice from description, and not the reversed reflection effect of the description-experience literature we have described above.

Magoon and Critchfield (2008) conceptualized a gain-loss asymmetry as a differential effect of positive and negative reinforcement on the bias parameter in the GML. Participants' choices were compared between a homogeneous condition where responses on both alternatives resulted in gains, and a heterogeneous condition where responses on one alternative produced a gain and responses on the other alternative cancelled a loss. The homogenous and heterogeneous conditions had identical reinforcement ratios and in the heterogeneous condition each participant experienced negative reinforcement at least once as the richer or leaner alternative. In the heterogeneous condition, response ratios were compared to both reinforcer ratios (Equation 3.1) and net ratios (Equation 3.2), where *R* were money gained and money losses cancelled and P were money losses that were not avoided via cancellation. A gain-loss asymmetry was defined as displaying a bias in the heterogeneous condition as compared to no bias in the homogeneous condition, while controlling for any inherent side bias (i.e. tendency to prefer responding on the left or right). Bias towards the positive reinforcement side in the heterogeneous condition would suggest that the functional value of responding for a gain is not equivalent to the functional value of responding to avoid a loss. Magoon and Critchfield (2008) observed some bias towards the negative reinforcement side, but ultimately no consistent change that would support the notion that responding for gains and to cancel losses had a different functional impact on behaviour. This was the case when using either reinforcer ratio (Equation 3.1) or net ratio (Equation 3.2).

Rasmussen and Newland (2008) similarly operationalized a gain-loss asymmetry as primarily a difference in bias. The effects of gains and losses were compared as the effect of positive reinforcement and negative punishment (losses). Individual bias parameters were compared between a reinforcement-only condition, where both sides scheduled gains at variable intervals, and a punishment condition, where one side scheduled gains and the other had a loss schedule superimposed on the existing gain schedule. In the reinforcement-only condition the response ratios were compared to the reinforcer ratios, and the punishment condition the response ratios were compared to both reinforcer and net ratios. Punishment was always scheduled to be leaner than the reinforcement schedule it was superimposed on, but it could be superimposed on either the richer or the leaner alternative in the existing reinforcement ratio. Thus in the punishment condition, responses on one side could result in either a gain or a loss, and responses on the other side that could only result in a gain.

Rasmussen and Newland (2008) hypothesized that if money gained and money lost have symmetrical functional impact on behaviour, then the participants should allocate their responses in a way that displays no bias towards the reinforcement-only alternative. Their behaviour should display a decrease in responding on the punishment side and an increase in responding on the reinforcement-only side, a distribution of responses that is only dependent on the net reinforcer ratio. If they display a bias towards the reinforcement-only alternative when punishment is introduced to the other alternative, then the functional value of experiencing a loss is not equivalent to the functional value of experiencing a gain. They also predicted that presence of punishment should reduce the participants' sensitivity to the reinforcement schedule as compared to the reinforcement-only condition.

Rasmussen and Newland (2008) observed unbiased matching in the reinforcement-only condition and a three-fold shift in bias towards the reinforcement side in the punishment condition. Additionally, for both the reinforcer and net ratio analyses, they observed generally lower sensitivity parameters for the punishment condition than the reinforcement-only condition, although this was not consistent across participants. Rasmussen and Newland concluded that experiencing punishment reduced sensitivity to existing reinforcement contingencies and biased the individual away from the punished alternative, demonstrating an asymmetry in the impact that gains and losses have on behaviour.

The lack of consensus in literature on the expected change in GML parameters in response to scheduled gains and losses makes prediction difficult. Firstly, Magoon and Critchfield (2008) and Rasmussen and Newland (2008) observed opposite findings for both bias and sensitivity in procedures were losses were associated with one of the alternatives. Although both negative reinforcement and negative punishment schedules served to deliver losses to the participant, the former encouraged and the latter discouraged responding in order to avoid a loss. Conceptually, the negative punishment is closer to our discounting task and Prospect Theory's description of negative events, as responding resulted in losses rather than avoidance of a loss. Therefore, we might expect results similar to Rasmussen and Newland. However, they scheduled punishment to only one alternative and we have been discussing a gain-loss asymmetry as choice between gains or between losses, so a concurrent schedules task like used by Critchfield et al. (2003) that schedules losses on both sides would be more relevant. Both Rasmussen and Newland and Critchfield et al. observed an increase in responses on the side that was either less punished (in Critchfield et al., when punishment rates were equal, demonstrated by the higher sensitivity parameter) or unpunished (in Rasmussen and Newland, when one side was punished, demonstrated by the higher bias parameter). Absence of consistent bias in the data collected by Critchfield et al. was presumably because losses were present on both sides, thus any bias would be cancelled out. Thus if we were to arrange losses on both alternatives, the work of Critchfield et al. and the direct-suppression model both predict greater sensitivity (i.e. a higher a parameter in the direct suppression model), and possibly, overmatching, to reinforcement.

Gain-loss asymmetry at local level of analysis

Concurrent schedule procedures allow for a distinction to be made between extended and local level of analysis. In local level of analysis, the immediate distribution of responses after outcomes is aggregated across other such instances and examined as a function of increasing time or number of responses since the last outcome. This approach does not rely on model fitting, and while it can be used to test for the individual's ability to distinguish the overall schedule in operation, it also offers additional insights into local processes that may not be observable at the extended level of analysis.

A common observation in local level analysis is that of *preference pulses*, where choice tends to favour the alternative that has just produced a reinforcer (the *just-productive* side) and gradually changes to preference for the side with the overall higher rate of reinforcement (e.g. Davison & Baum, 2000; Davison & Baum, 2002; Landon et al., 2003). Preference pulses suggest unique, transient effects of reinforcers on behaviour that are not necessarily dependent on the schedule in operation. Recent work by McLean, Grace, Pitts and Hughes (2014) has demonstrated that before drawing conclusions about the nature of unique effects of reinforcers, such preference pulses need to be corrected for the general pattern of responding, or visit structure, inherent to responding on a concurrent schedule task. The visit structure refers to the typical amount of time that an individual spends responding on one alternative before a switch. Concurrent schedule tasks arrange delivery of reinforcers that favours responding on the same alternative for a period of time as opposed to rapid switching. This is primarily due to a commonly used changeover delay (COD; Herrnstein, 1961; Shull & Pliskoff, 1967). A COD is a period of time immediately after switching when no reinforcers can be delivered and is used to ensure that frequently changing sides is not reinforced. In tasks that arranged CODs, preference pulses tended to be stronger than in no-COD tasks (e.g. Krageloh & Davison, 2003; see Gomes-Ng, Landon, Elliffe, Bensemann, & Cowie, 2018 for a recent discussion). However this observed preference for the same side as the last delivered reinforcer is conflated with a general visit structure that favours the same side as the last response.

Gomes-Ng, Elliffe and Cowie (2017) applied McLean and colleagues' proposed correction and further demonstrated that local effects of reinforcers across conditions can be effectively compared from the difference between responses after *real* reinforcers and responses at any other equivalent time in the task (i.e. the general visit structure). Responses at other equivalent times in the task were quantified as responses after *hypothetical* reinforcers, which were inserted into the acquired data sets following responses that could have produced a reinforcer, given a different sampling of the variable or random schedule, but didn't. These hypothetical reinforcers were also arranged such that they did not temporally overlap with real reinforcers. Gomes-Ng et al.'s re-analysis of several data sets from animal subjects demonstrated that preference pulses still occurred following reinforcers after correcting for the general visit structure. While this correction has not been applied to human preference pulse data, the simulations done by McLean et al. (2014) suggest that the effect of general visit structure should be controlled for in any concurrent schedule task.

Thus examining local effects, while controlling for the general tendency to respond on a concurrent schedule task, offers another level of analysis for comparing the effects of gains and losses. In Study 1, we noted that discounting rates depended on task progression, where some aspect of the repeated trials resulted in shallower discounting of both gains and losses. We could not examine the local effect of experiencing outcomes on subsequent behaviour in our titrating amount procedure, particularly, we could not determine whether there was a difference in responses after a successful probabilistic gain or loss. In concurrent schedules, we are able to examine the effect of experiencing a probabilistic gain or loss on subsequent responses for any asymmetry, such as a difference in the strength or duration of the preference pulses.

It is unclear how the description-experience gap predictions of risk seeking for gains and risk aversion for losses at most probabilities can be applied here, but we can expect differences in the strength or duration of the tendency to stay versus switch after gains as compared to losses. Given that the inverse S-shaped function proposed for experiential choice in the description-experience gap literature indicates greater deviation from linear weighting of probabilities for gains rather than losses, we might similarly expect greater deviation from the general visit-structure after the occurrence of a gain rather than a loss. If the two alternatives differ in magnitude, the application of the value function would predict that changes in the magnitude of a loss would have greater impact on subsequent responses than equivalent changes in the magnitude of a gain.

To our knowledge, local effects of gains vs. losses have not been compared in concurrent schedule tasks with humans, but we can draw some insight from local effects analysis in a related procedure that has also been used with human participants: a signal detection task (e.g. Johnstone & Alsop, 2000). Typically, the participant is presented with one of two stimuli at a time and responds by choosing one of two available response alternatives. One of the alternatives is the correct response to the given stimulus and the participant can be reinforced for giving the correct response. The rate of reinforcement for the two stimuli can be varied, such that correct responses for one stimulus can be reinforced more frequently than for the other. The difficulty of distinguishing the two stimuli can be adjusted to create the desired frequency of correct vs. incorrect responses. Lie and Alsop (2009) introduced punishment to a signal detection procedure in the form of point-loss for incorrect responses for given stimuli. This rate of punishment was varied between the two stimuli while the rate of reinforcement for correct responses was held constant. A re-analysis of these data for local effects showed that responses after receiving a punisher on one alternative tended to favour the other alternative, independent of whether the other alternative had overall a more or less frequent rate of punishment (Lie, 2010). This was in contrast to a different group of participants, who experienced varying reinforcement rate and were not punished for errors; their responses were to favour the alternative associated with higher rate of reinforcement. Notably, this analysis was done without correcting for the general pattern of responding on such tasks, so the true extent of these local effects is not known.

Kubanek, Snyder and Abrams (2015) examined asymmetry in local effects of reinforcers and punishers in a task similar to a signal detection procedure. Participants responded to auditory or visual stimuli and received hypothetical monetary gains and losses based on the accuracy of their response. Task difficulty was set to result in a 60% accuracy. The amount gained or lost was varied so that responses could also be analysed for the effect of magnitude, ranging from 5 to 25 cents. Local effects of experiencing a gain or loss outcome were compared by examining subsequent tendency to repeat the same choice relative to the expected baseline of a 50% chance, since the task was set-up such that the outcome on the previous trials did not predict the outcome on the next trial.

Kubanek et al. (2015) observed an asymmetry in participants' response repetition tendencies after equal magnitude rewards and punishers. Losses occasioned avoidance and gains occasioned repetition of the just-productive alternative, but the extent of avoidance was approximately 2-3 times greater than the extent of repetition. This effect was transient, akin to preference pulses, as outcomes two trials prior to the test trial did not significantly affect current choice. The authors noted that this echoes the findings of Rasmussen and Newland (2008), who observed greater bias shown to the reinforcement-only side. Kubanek et al. also observed an effect of magnitude with gains, but not with losses, which was contrary to our predictions based on the value function of Prospect Theory. Losses elicited a similar avoidance tendency of the last choice regardless of its magnitude, while repetition tendency increased as magnitude of the gain increased. Notably, this effect of magnitude bears similarities to the effect of changing magnitude in probability discounting, where discounting of gains was affected by their magnitude, but discounting of losses was not affected by their magnitude (e.g. Estle et al., 2006). Larger gains were discounted more steeply than smaller gains and in Kubanek et al. larger gains elicited a greater chance of repetition. The link between choice repetition and discounting is not clear, but both suggest that we are more likely to see a gain-loss asymmetry with larger amounts.

Auckland Card Task: Rapid-acquisition choice procedure

In most of the concurrent schedule tasks we have discussed so far (e.g. Rasmussen et al., 2008), the participants' responses to change in reinforcement schedules were measured over several sessions until a defined stability criterion in their responding was reached. This approach limits the number of conditions as well as the total sample size that is practical and cost-efficient to collect with human participants. Behaviour under concurrent-schedules has also been examined in *rapid-acquisition* choice procedures using animal subjects, where changes in reinforcer contingencies were not signalled and occurred rapidly during the sessions (Davison & Baum, 2000). This procedure demonstrated a rapid increase in sensitivity with each experienced reinforcer while keeping the task duration much shorter than typical. Subsequent work by Lie, Harper and Hunt (2009) and Krageloh, Zapanta, Shepherd, and Landon (2010) has demonstrated that this method generates reasonable GML fits for human participants responding for reinforcers. Their average sensitivity and R^2 values are summarized in Table 3.2 and were in the reasonable range for human participants (Kollins et al., 1997). Lie et al. and Krageloh et al. also included an analysis of local effects of reinforcers, showing preference pulse patterns largely consistent with literature. Upon receiving a reinforcer, participants responded more on the just-productive side and this preference gradually decreased with time or number of responses since reinforcer. Notably, these were not corrected for the general pattern of responding proposed by McLean et al. (2014).

Table 3.2

Summary of average sensitivity and R^2 in Lie et al. (2009) and Krageloh et al. (2010)

	a median (Q_1, Q_3)	R^2 median (Q_1, Q_3)
Lie et al. (last block of 5 reinforcers)	0.35 (0.17, 0.66)	0.62 (0.23, 0.85)
Krageloh et al. (2nd half of the sessions)	0.39 (0.16, 0.83)	0.61 (0.42, 0.70)

Note. Q_1 = first quartile, Q_3 = third quartile.

Based on the rapid-acquisition tasks above, Bull, Tippett and Addis (2015) created a novel task which measured participants' behaviour in response to changes in both frequency and magnitude of gains, as well as losses. The Auckland Card Task consisted of four concurrent variable-interval schedule conditions, all of which were designed to be completed within a few hours by each participant. Within each condition, one dimension of a gain or loss was varied between the two alternatives while the rest were kept constant. The four conditions varied gain frequency, loss frequency, gain magnitude and loss magnitude. For example, in the gain frequency condition, gains were scheduled on a more frequent concurrent VI schedule, were of fixed magnitude, but varied in probability on each alternative (e.g. 75% vs. 25% chance of a gain). In the same condition, losses were scheduled on a less frequent concurrent VI schedule, and were of fixed magnitude and probability on each alternative. For clarity, we will subsequently refer to the scheduled outcome that varied between the two alternatives in each condition as the primary outcome (e.g. gains in the gain frequency condition) and the other scheduled outcome that did not vary between the two alternatives as the secondary outcome (e.g. losses in the gain frequency condition).

Re-analysis of Bull et al. (2015) for gain-loss asymmetry. Bull et al. piloted their task on thirty participants and while they did not investigate a gain-loss asymmetry, we can use the available sensitivity parameters to determine whether there were any differences between the conditions. We accessed supplementary data from Bull (2013) thesis dissertation listing each participants' sensitivity values in each condition of the same study reported in Bull's 2015 paper. Bull reported estimates of sensitivity parameters using two methods. First, the GML was fitted to all of the responses and primary outcomes received in each condition, producing an *overall* measure of sensitivity (mean sensitivity and median R^2 are summarized in the first two columns of Table 3.3). Second, the data were split into five blocks of primary outcomes and the GML was fitted to each block (we conducted the same analysis, and detail on this is provided in the methods section below). The sensitivities derived from the last few blocks were then averaged to produce an *asymptotic* measure of sensitivity that excluded changes in the response ratios in the early/learning trials (mean sensitivities are summarized in the third column of Table 3.3; R^2 were not reported). All of the sensitivity estimates were derived from fitting the GML to the primary outcomes only and no net analysis (e.g. using the direct-suppression version of the GML) was conducted.

Comparison of overall and asymptotic sensitivity estimates in Table 3.3 suggested that sensitivity to gains changed to a greater degree over the course of the session than sensitivity to losses. Furthermore, all of the conditions showed overmatching in their asymptotic sensitivity, which corresponds to our predictions of overmatching based on the description-experience gap, but not to the typical pattern of undermatching in human literature.

Table 3.3

Summary of average sensitivity and R^2 reported in Bull (2013)

	Overall a mean (SEM)	Overall R^2 median (Q_1, Q_3)	Asymptotic a mean (SEM)
Gain frequency	0.54 (0.07)	0.74 (0.53, 0.81)	1.31 (0.19)
Loss frequency	1.15 (0.12)	0.90 (0.84, 0.93)	1.30 (0.15)
Gain magnitude	0.90 (0.13)	0.80 (0.46, 0.96)	1.12 (0.21)
Loss magnitude	1.16 (0.13)	0.90 (0.72, 0.95)	1.32 (0.20)

Note. Q_1 = first quartile, Q_3 = third quartile

We conducted additional analyses on these data for any gain-loss asymmetry in the overall and asymptotic estimates of sensitivity. Analysis of overall sensitivity parameters showed a gain-loss asymmetry for frequency, but not magnitude conditions. The mean sensitivity parameter for loss frequency was significantly higher than for gain frequency ($t(29) = -5.29$, $p < .001$) but the mean sensitivity parameter for loss magnitude was not significantly different from gain magnitude ($t(28) = -1.90$, $p = .067$). To be certain of the asymmetry we also assessed gain-loss asymmetry in a sample limited to participants with R^2 of at least 0.50; data were not normally distributed and the Wilcoxon Signed Rank Test showed a significant difference between the loss and gain frequency conditions ($Z = -3.73$, $p < .001$; $N = 22$) and no significant difference between the magnitude conditions ($Z = -0.60$, $p = .546$; $N = 21$). In contrast to overall sensitivity, asymptotic sensitivity showed no gain-loss asymmetry for frequency ($t(27) = -0.08$, $p = .935$) and magnitude ($t(21) = -0.43$, $p = .673$) conditions. Thus, a gain-loss asymmetry was present only if preferences in the earlier blocks were included and did not correspond to our predictions of greater sensitivity to gains than losses based on the description-experience gap.

The current study

In Study 2, we examined gain-loss asymmetry in a rapid-acquisition choice procedure that arranged gains and losses via concurrent schedules. We decided to use a rapid-acquisition concurrent schedule task to examine gain-loss asymmetry. The rapidly changing schedule in such tasks is suitable for collecting large quantities of data over short periods of time, especially with human participants where costs of prolonged data collection are prohibitive of acquiring sample sizes sufficiently powered for parametric analysis. We have selected the ACT as one such procedure as it has successfully trialled delivery of scheduled gains and losses, as well as distinguished between frequency and magnitude, thus separating the four decision-making dimensions also present in our probability discounting procedure. While Bull (2013) data did not show our predicted pattern of gain-loss asymmetry, the available data does not allow for both an extended and a local level of analysis, as well as an evaluation of goodness of fit of the asymptotic estimates of sensitivity. Therefore, our specific goals were to:

1. Examine the reliability of ACT by replicating Bull et al. (2015) procedure and determine sensitivity parameters for gain-loss asymmetry analysis.

We examined goodness of fit of GML to ACT data in our replication by comparing R^2, percentage of missing data and whether sensitivity increased linearly as a function of Block compared to Bull et al.'s data. Bull et al. observed average R^2 values that were higher than other rapid-acquisition choice procedures for the overall sensitivity measure while goodness of fit of the GML to each block of data, and by extension the asymptotic sensitivity, is not known. We expected average R^2 values and proportion of missing data comparable to Bull et al., as well as a linear increase in sensitivity as a function of block.

Prior to extended level analysis for gain-loss asymmetry, we considered the method of deriving final GML fits Bull et al. used, as well as alternative estimates of both the response and reinforcer ratios in order to derive reliable GML parameters. We considered adjustments to the response ratio based on responses in the last four blocks and based on responses to the last five primary outcomes experienced, akin to the overall and asymptotic estimates used by Bull et al. We also considered whether model fit can be improved with inclusion of secondary outcomes according to the direct-suppression GML (Equation 3.2). In each condition of the ACT, the primary outcome ratio differed from the net outcome ratio (see Table 3.4), which produced a different range of predictor values while the measured behaviour remained unchanged. Since the number of parameters does not change between Equations 3.1 and 3.2, and thus does not necessitate adjustments for such changes (e.g. AIC; Akaike, 1998), we directly compared R^2 values in order to determine whether participants' behaviour was better described by the net rather than primary outcome ratios. Based on the acquired data, we determined a final estimate of sensitivity to arranged outcomes for each individual and each condition in order to conduct a gain-loss asymmetry analysis.

2. Examine gain-loss asymmetry at extended level of analysis.

At the extended level of analysis, we quantified gain-loss asymmetry as a difference in sensitivity parameters between the gain and loss condition within the frequency and magnitude tasks. ACT arranges gains and losses on both alternatives, but only one property (frequency or magnitude) of one of these outcomes varied between the two alternatives. The four reinforcer and net ratios were also equated across conditions, thus the total number of points that can be acquired by the end of the task did not differ. Due to the presence of gains and losses on both alternatives, we would not expect a gain-loss asymmetry to manifest itself as a difference in bias (Critchfield et al., 2003). Based on our application of Prospect Theory and the description-experience gap literature, as well as the more consistently observed difference in sensitivity (Critchfield et al., 2003; Magoon & Critchfield, 2008; Rasmussen & Newland, 2008; Bull, 2013), we expected a gain-loss asymmetry as a difference in sensitivity to changes in gain as opposed to loss delivery. We expected greater sensitivity to changes in gain as opposed to loss frequency to emerge with increasing experience of the scheduled outcomes. In the magnitude task, we expected greater sensitivity to changes in loss as opposed to gain magnitude. Lastly, we also examined whether participants would show the same pattern of overmatching observed by Bull (2013), or whether data would resemble the more commonly seen undermatching in human participants.

3. Examine gain-loss asymmetry at local level of analysis.

At the local level of analysis, we quantified a gain-loss asymmetry as a difference in the pattern of responses after primary outcomes that has been corrected for visit structure. To our knowledge, the McLean et al. (2014) correction has not been applied to the study of preference pulses generated by humans. We expected an immediate response to losses to be opposite to that of gains and the extent of deviation from the general visit-structure to be greater after gains than after losses. In the magnitude task, we examined whether magnitude had a greater impact on responses after a loss in accordance with the value function, or whether data were consistent with Kubanek et al. (2015), where the opposite was observed.

4. Examine whether an order effect would remain after resetting the balance and separating the conditions.

In addition to completing the ACT, participants also discounted gains and losses in a safe-risky experiential probability discounting task. Our observation of order effects in the experiments in Study 1 was also confounded with change in balance. In Study 2, we removed this confound by starting each condition with $3000. Due to time constraints of including four ACT conditions and two conditions of the discounting task, the gain and loss conditions were split across the two sessions. Hence, the differences to Study 1 included balance reset and a day separation between the conditions. We examined gain-loss asymmetry under these modified conditions for any changes in the effect of order observed in Study 1.

5. Examine if sensitivity to gains and losses can predict discounting after controlling for expected order effects.

Given the theoretical relevance of the probability discounting and concurrent schedule procedures, we also considered whether performance on one task was predictive of performance on the other. We conducted a linear regression with condition order and sensitivity to frequency and magnitude as predictors of discounting in order to determine whether sensitivity predicted significant change in discounting after controlling for the effect of order.

Method

Participants

We recruited a total of 123 participants attending Victoria University of Wellington through the School of Psychology Research Programme online tool. Participants who dropped out of the study after completing the first session were removed from analysis, resulting in the final sample of 103 participants. Participants received course credit and completed an informed consent procedure.

Materials

Auckland Card Task. We programmed the ACT using Microsoft Visual Basic based on Bull et al. (2015). It began with a series of instruction screens with content tailored to each condition (see Appendix E for exact wording). After the instructions, participants were presented with images of two decks of cards, with only the back of the cards visible (see Figure 3.4, A). Below the decks was a bar titled “Winnings” and its length was indicative of the current total money won. The actual number was not displayed to the participant.

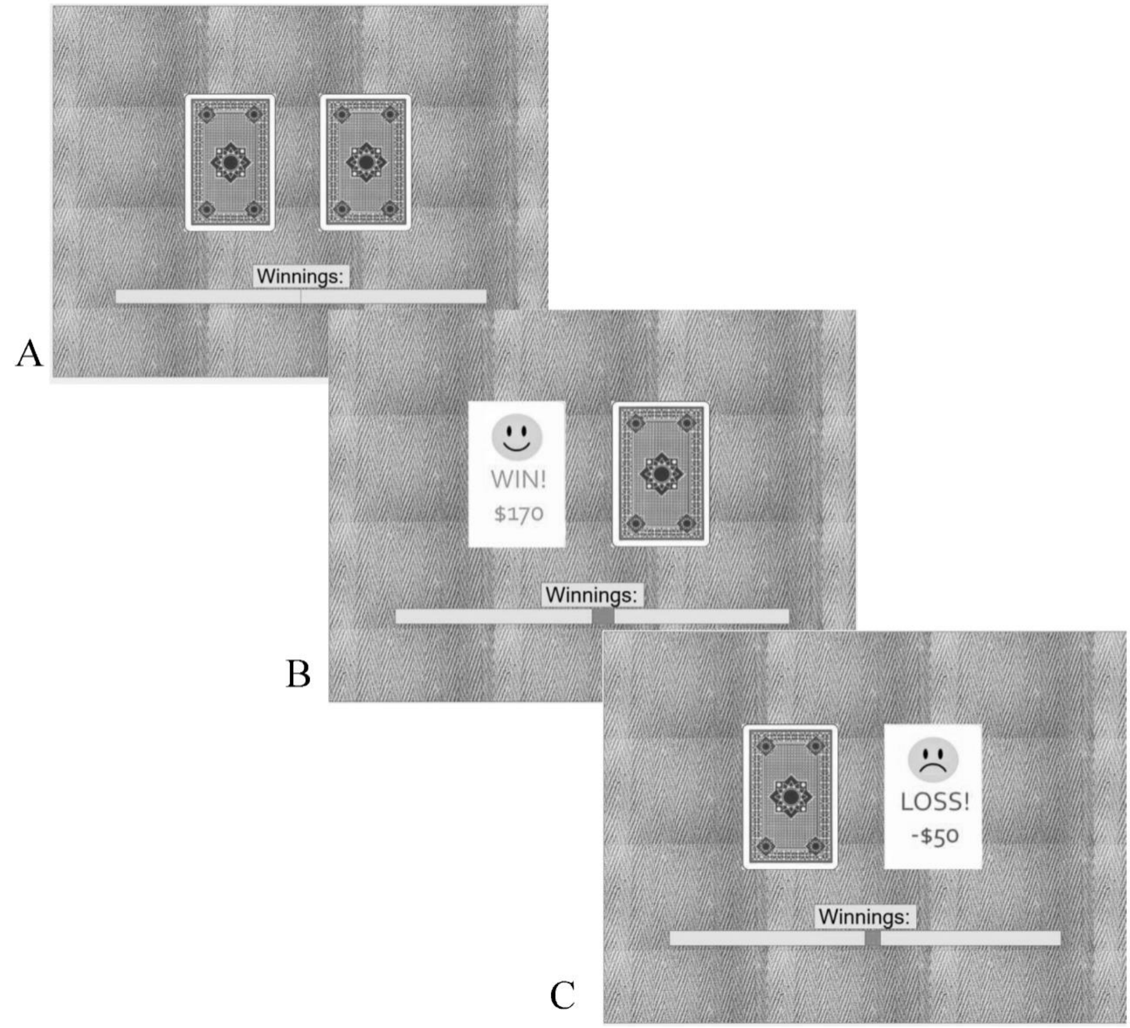

Figure 3.4. Screen capture of the ACT. The participants were presented with images of two decks of cards, with only the back of the cards visible (A). Choosing a deck via key press resulted in either the scheduled reward/penalty animation sequence or the no reward/penalty animation sequence. Panel B is demonstrating a scheduled reward probabilistically allocated to the left deck. Panel C is demonstrating a scheduled penalty probabilistically allocated to the right deck. The winnings bar is updated after each outcome, and is coloured green if the balance is above zero and red if below zero.

Once the trials began, deck selection was made with their dominant hand by pressing caps lock for the left deck and enter key for the right deck[5]. Choosing a deck via key press resulted in a card flipping sound being played and either the scheduled gain/loss animation sequence, or the no gain/loss animation sequence. If a gain/loss was due, a card with the amount won or lost was shown for 1000ms accompanied by a "ding" sound for gains and a

[5] Bull et al. (2015) used the left and right control keys to record responses, which were the same distance apart on the keyboard as caps lock and enter keys.

“buzz” sound for loss (see Figure 3.4, B and C for examples of cards). The “Winnings” bar was updated accordingly, with gains shifting the net amount to the right and loss to the left. If the current net amount was above 0, it was coloured green, if below, it was coloured red. If no gain/loss was due, a random playing card was shown for 200ms. During these sequences, no responses were recorded. After the sequences ended, the participant was able to choose a deck of cards again. A two second changeover delay was used to ensure that frequently changing decks was not reinforced.

Each of the four conditions had concurrent gains and losses scheduled independently of one another. Table 3.4 (adapted from Bull et al., 2015) lists the schedules arranged within each condition. In the frequency conditions, the independent variable of interest was either gain or loss frequency and was the primary outcome. The primary outcome was set to one magnitude on both decks, but varied in frequency. Each condition also had a secondary outcome, either a loss (for gain frequency condition) or a gain (for loss frequency condition), which was of the same magnitude and frequency on both decks. The primary outcomes were scheduled according to a richer concurrent variable-interval schedule and were more numerous, while the secondary outcomes were scheduled according to a leaner concurrent variable-interval schedule and were less numerous. Practically, both the primary and secondary outcomes were arranged using dependent scheduling (see Stubbs & Pliskoff, 1969 for a related procedure). The primary outcome was coded according to a single base VI 4s schedule and the secondary outcome according to a single base VI 10s schedule. Once a primary or secondary outcome was due, it was probabilistically allocated to either the left or right deck using the arranged probabilities for each component. The outcome then needed to be collected by the participant before a new outcome was scheduled. If a primary or secondary outcome were both scheduled to the same deck, their order of appearance was randomly determined. For example, in the 1st component of the gain frequency task, the gains ($50) would be allocated to the left deck with a 25% (with a maximum of 5 gains) chance or to the right deck with a 75% chance (with a maximum of 15 gains). Concurrently, losses ($30) would be allocated to the left or right deck with a 50% probability (maximum of 5 on each side).

Table 3.4

Summary table of the four conditions in the Auckland Card Task adapted from Bull et al. (2015)

Condition	Component	Gains (Probability/Amount/Number of outcomes available)		Losses (Probability/Amount/Number of outcomes available)		Net reward		Primary outcome ratio	Net ratio
		Deck 1	Deck 2	Deck 1	Deck 2	Deck 1	Deck 2		
Gain frequency	1	.25/$50/5	.75/$50/15	.50/$30/5	.50/$30/5	$100	$600	1:3	1:6
	2	.35/$50/7	.65/$50/13	.50/$30/5	.50/$30/5	$200	$500	1:1.86	1:2.5
	3	.65/$50/13	.35/$50/7	.50/$30/5	.50/$30/5	$500	$200	1.86:1	2.5:1
	4	.75/$50/15	.25/$50/5	.50/$30/5	.50/$30/5	$600	$100	3:1	6:1
Loss frequency	1	.50/$170/5	.50/$170/5	.75/$50/15	.25/$50/5	$100	$600	1:3	1:6
	2	.50/$170/5	.50/$170/5	.65/$50/13	.35/$50/7	$200	$500	1:1.86	1:2.5
	3	.50/$170/5	.50/$170/5	.35/$50/7	.65/$50/13	$500	$200	1.86:1	2.5:1
	4	.50/$170/5	.50/$170/5	.25/$50/5	.75/$50/15	$600	$100	3:1	6:1
Gain magnitude	1	.50/$25v/10	.50/$75v/10	.50/$30/5	.50/$30/5	$100	$600	1:3	1:6
	2	.50/$35v/10	.50/$65v/10	.50/$30/5	.50/$30/5	$200	$500	1:1.86	1:2.5
	3	.50/$65v/10	.50/$35v/10	.50/$30/5	.50/$30/5	$500	$200	1.86:1	2.5:1
	4	.50/$75v/10	.50/$25v/10	.50/$30/5	.50/$30/5	$600	$100	3:1	6:1
Loss magnitude	1	.50/$170/5	.50/$170/5	.50/$75v/10	.50/$25v/10	$100	$600	1:3	1:6
	2	.50/$170/5	.50/$170/5	.50/$65v/10	.50/$35v/10	$200	$500	1:1.86	1:2.5
	3	.50/$170/5	.50/$170/5	.50/$35v/10	.50/$65v/10	$500	$200	1.86:1	2.5:1
	4	.50/$170/5	.50/$170/5	.50/$25v/10	.50/$75v/10	$600	$100	3:1	6:1

Note. Primary outcome ratios for the loss conditions are reversed, such that the side with more losses predicts fewer responses.

The logic for arranging outcomes in the magnitude conditions was similar but the primary outcome and independent variable of interest was either gain or loss magnitude. Here, the primary outcome was set to one frequency for both decks, but varied in magnitude. The secondary outcome was either a gain (for loss magnitude condition) or a loss (for gain magnitude condition), set to the same magnitude and frequency on both decks. Hence, while the primary outcome was scheduled based on the same richer VI schedule, the probability of allocating an outcome to either side was identical (50%) but varied in magnitude. For example, in the 1st component of the loss magnitude task, the mean magnitude of losses allocated to the left was $25 (ranging from $3-48) and to the right was $75 (ranging from $8-143), with a maximum of ten loss outcomes on each side. Concurrently, gains were the secondary outcome and were set to a single magnitude ($170) and allocated with equal probability to each side (maximum of five on each side).

After experiencing the scheduled outcome, the participant continued choosing decks and experiencing gains/losses until all 30 outcomes were experienced or six minutes had passed since the start of the component. Notably, this was a shorter duration than used by Bull et al. (2015), who terminated the components after eight minutes and was done to constrain the total experiment time to two hours. Once the component terminated, the participants were given a rest break with a screen that repeated the condition hint (see Appendix E) and their total winnings were displayed. The total winnings did not reflect the actual winnings as these were identical from component to component (assuming all 30 outcomes were experienced). Instead, the actual winnings had a random amount added to or subtracted from it before being displayed in order to mask from the participant that each component resulted in the same net total of $700. The four components per task were presented in random order.

Experiential Money Task. This task was programmed as described in Study 1, Experiment 2.1, with a few exceptions. The gain and loss conditions were split across the two data collection sessions which were a day apart. Each condition began with the same initial balance: $3000.

Procedure

Data collection was spread over two one-hour sessions that were a day apart. Participants had to sign up for both sessions, but were instructed that they can terminate their participation at any time. During each session, the participants were assigned to a cubicle with a computer terminal and headphones and tested in groups of four to fifteen. Each session consisted of two ACT conditions and one discounting condition, so that by the end of the two sessions each participant completed all of the tasks and conditions. The order of the conditions was counterbalanced across the two sessions, but always occurred such that the

ACT conditions preceded the discounting condition. Upon completion of the second session, they were debriefed on the purpose of the study.

Results

Does the ACT produce reliable parameters following Bull et al. analysis?

Our first aim was to test the goodness of fit of the GML to responding in the ACT and its reliability for examining subsequent gain-loss asymmetry analysis. We compared average R^2, percentage of missing data and whether sensitivity increased linearly as a function of block relative to Bull et al. (2015).

Fitting GML to data. Following Lie et al. (2009) and Bull et al. (2015), participants' responses in each component were split into five blocks of data according to the number of successive primary outcomes they had experienced: from the start of the component until the receipt of the 4th outcome (block 1), from the 4th outcome to receipt of the 8th outcome (block 2), from the 8th outcome until the receipt of the 12th outcome (block 3), from the 12th outcome until the receipt of the 16th outcome (block 4), and lastly from the 16th outcome until the receipt of the 20th outcome (block 5). For each of these blocks, a response ratio and a primary outcome ratio were calculated. The response ratio consisted of logged ratio of left to right responses made within that block. The primary ratio consisted of logged ratio of left to right primary outcomes from the start of the component to the end of that block (e.g. reinforcer ratio for block 2 consisted of primary outcomes experienced during blocks 1 and 2). If, in some blocks, no responses or primary outcomes occurred at either the left or right side, resulting in a zero value and an error for the logged ratio calculations, we added 0.25 to both the right and left alternatives (Hautus, 1995; Brown & White, 2005). This adjustment was exclusive to either the response or the primary outcome ratios that had a zero value; if the participant had a zero value in their primary outcome but not their response ratio, their primary outcome ratio's numerator and denominator were adjusted by 0.25 and their response ratio was not. Additionally, each primary outcome ratio in the loss conditions was multiplied by -1 in order to produce a ratio that predicted more responses on the side with fewer/lower losses. This adjustment resulted in positive slopes, facilitating comparison to the gain conditions.

The GML was fitted to response and primary outcome ratios such that for each participant, five GML fits were derived for five blocks of data across the four components using least squares linear regression. For example, GML fitted to block 1 would consist of four response and their respective primary outcome ratios from the first block across each of the four components. This method of fitting the GML to each block produced parameters that describe behaviour at a comparable number of primary outcomes experienced across each component within a given condition. At each block, for each condition and participant we derived five measures of sensitivity (a) and bias (log c). Sensitivity and bias were normally distributed (see the means and standard error of the mean in Table 3.5).

Table 3.5

Summary of Study 2 GML-derived parameters for each block of primary outcomes.

			Participants who have reached that block of primary outcomes in all 4 components			
Condition	Block	% of participants that did not reach block in at least one component (count)	N	Mean a (SEM)	Mean log c (SEM)	Median R^2 (Q_1, Q_3)
Gain Frequency	1	0 (0)	103	0.12 (0.05)	-0.08 (0.03)*	0.33 (0.11, 0.70)
	2	2.91 (3)	100	0.84 (0.09)	-0.05 (0.04)	0.63 (0.23, 0.84)
	3	12.62 (13)	90	1.07 (0.10)	-0.13 (0.05)*	0.61 (0.23, 0.80)
	4	24.27 (25)	78	1.12 (0.13)	-0.05 (0.04)	0.64 (0.20, 0.85)
	5	39.81 (41)	62	1.34 (0.19)	-0.11 (0.05)*	0.61 (0.35, 0.84)
Loss Frequency	1	0 (0)	103	0.47 (0.08)	-0.11 (0.03)*	0.66 (0.30, 0.87)
	2	11.65 (12)	91	0.74 (0.09)	0.02 (0.04)	0.69 (0.38, 0.85)
	3	37.86 (39)	64	0.50 (0.11)	-0.10 (0.04)*	0.54 (0.26, 0.79)
	4	57.28 (59)	44	0.65 (0.15)	-0.06 (0.05)	0.63 (0.35, 0.87)
	5	73.79 (76)	27	0.33 (0.14)	0.02 (0.07)	0.37 (0.19, 0.76)
Gain Magnitude	1	0 (0)	103	0.12 (0.05)	-0.08 (0.03)*	0.33 (0.09, 0.68)
	2	6.80 (7)	96	0.56 (0.08)	-0.02 (0.04)	0.43 (0.11, 0.78)
	3	19.42 (20)	83	0.61 (0.10)	-0.07 (0.04)	0.51 (0.21, 0.78)
	4	36.89 (38)	65	0.76 (0.13)	0.02 (0.04)	0.55 (0.31, 0.82)
	5	51.46 (53)	50	0.88 (0.15)	-0.10 (0.04)*	0.51 (0.30, 0.68)
Loss Magnitude	1	0 (0)	103	0.19 (0.04)	-0.11 (0.03)*	0.47 (0.13, 0.77)
	2	14.56 (15)	88	0.63 (0.11)	-0.05 (0.04)	0.42 (0.15, 0.69)
	3	37.86 (39)	64	0.53 (0.13)	-0.03 (0.05)	0.36 (0.13, 0.63)
	4	51.46 (53)	50	0.46 (0.15)	-0.12 (0.04)*	0.37 (0.11, 0.78)
	5	66.99 (69)	34	0.56 (0.19)	-0.01 (0.07)	0.53 (0.19, 0.63)

Note. * one sample t-test $p < .05$. Q_1 = first quartile. Q_3 = third quartile.

The fit of the model to data was determined by calculating the R^2, which were not normally distributed, with negative kurtosis across the blocks (see the medians and quartiles in Table 3.5). The values were overall lower than the median R^2 reported by Bull (2013). However, we are not able to compare directly to Bull's block by block R^2 as the reported R^2 were for the overall sensitivity measures, which fits the GML to responses across blocks. In our data, by the 5th block only the gain frequency condition achieved median R^2 values close to medians reported by Lie et al. (2009) and Krageloh et al. (2010), with the other median R^2 values generally lower.

Missing data. Fitting the GML to successive blocks of primary outcomes replicated Bull et al.'s (2015) observation of missing data: participants did not always complete five blocks of primary outcomes for some of the components. The dependent schedule relies on the delivery of the arranged outcomes before a new one can be arranged. If the participant develops exclusive preference for one side or repeatedly switches while remaining on each side for a shorter duration than the COD, then the arranged outcome will not be delivered and the session will terminate at the time limit. The third column of Table 3.5 lists the percentage of participants that did not reach the specified block of reinforcers for at least one of their components in each condition, resulting in at least one of the four log response ratios that could not be calculated (zero responses on both sides).

Our rate of missing data was higher than Bull (2013). Missing data at each block was not reported, but this can be approximated from the number of participants without final estimates of sensitivity measures. Bull's method of averaging the last 2-3 blocks of sensitivities resulted in one uncalculated sensitivity in each of the gain frequency, gain magnitude and loss frequency conditions, and seven in the loss magnitude condition. Uncalculated sensitivities imply that at least one of the final 2-3 blocks could not be calculated, as in our data set, but their rate of 3-23% missing data was generally lower than our rate of 12.62-39.81% of participants who did not reach block 3 in at least one of their components (see Table 3.5 for exact percentages).

Trends in bias. Examination of Table 3.5 showed a slight, but, if present, consistently negative bias among the means, indicative of bias towards the R_2 side and more responses on the right side across the components. Means were significantly different from zero consistently for all of the first blocks, and only some of the subsequent blocks (indicated by an asterisk in Table 3.5). This was likely due to being instructed to use their dominant hand only to respond on either the Caps Lock key for the left or Enter key for the right, making a right-hand side bias more likely given that most of the participants would have been righthanded. This was more likely to happen when the participant had little information on the distribution of the primary outcomes at the start of each component, and less likely at subsequent blocks.

Trends in sensitivity. We examined whether sensitivity increased linearly as a function of block for participants that have reached the 5th block of primary outcomes in all four of their components for each condition (see the fourth column of Table 3.5). Figure 3.5 plots the average sensitivities and their standard errors of the means listed in Table 3.5 as a function of blocks for each condition. Similar to Bull et al. (2015), we observed larger error bars at later blocks in all conditions due to decreasing sample size of available data. Bull et al.'s analysis identified linear trends in sensitivities across blocks for all four conditions, with participants showing increasing sensitivity to reinforcement as they progressed through the blocks of reinforcers in each component. The patterns and range of sensitivities for the two gain conditions in our sample were comparable to Bull et al., showing linear increase in sensitivity, unlike the two loss conditions. Our participants showed lower average sensitivities to losses relative to Bull et al., and the loss frequency condition in particular did not show a pattern of increasing sensitivity across blocks.

Repeated-measures ANOVAs showed a significant difference between blocks for the two gain conditions and loss frequency, but not for loss magnitude (see Table 3.6). Contrasts showed a significant linear trend for the two gain conditions only, with increasing sensitivity from 1st to 5th block. Examination of contrasts showed significant linear trends for gain frequency (F(1, 61) = 45.82, p < .001, η_p^2 = .429) and gain magnitude (F(1, 49) = 21.67, p < .001, η_p^2 = .307), indicating that sensitivity increased for participants that successfully progressed until the 5th block of reinforcers.

Table 3.6

Study 2 sphericity and ANOVA statistics for a by-block analysis of sensitivities

	Sphericity				ANOVA			
	χ^2	df	p	ε	F	df	p	η_p^2
Frequency								
Gains	32.34	9	<.001	.81[b]	17.26	3.22, 196.48	<.001	.221
Losses	12.49		.188		2.62	4, 104	.039	.091
Magnitude								
Gains	36.49	9	<.001	.87[b]	6.53	3.49, 171.05	<.001	.118
Losses	28.09		.001	.70[a]	1.61	2.81, 92.76	.194	.047

Note. a = Greenhouse-Geisser correction; b = Huynh-Feldt correction.

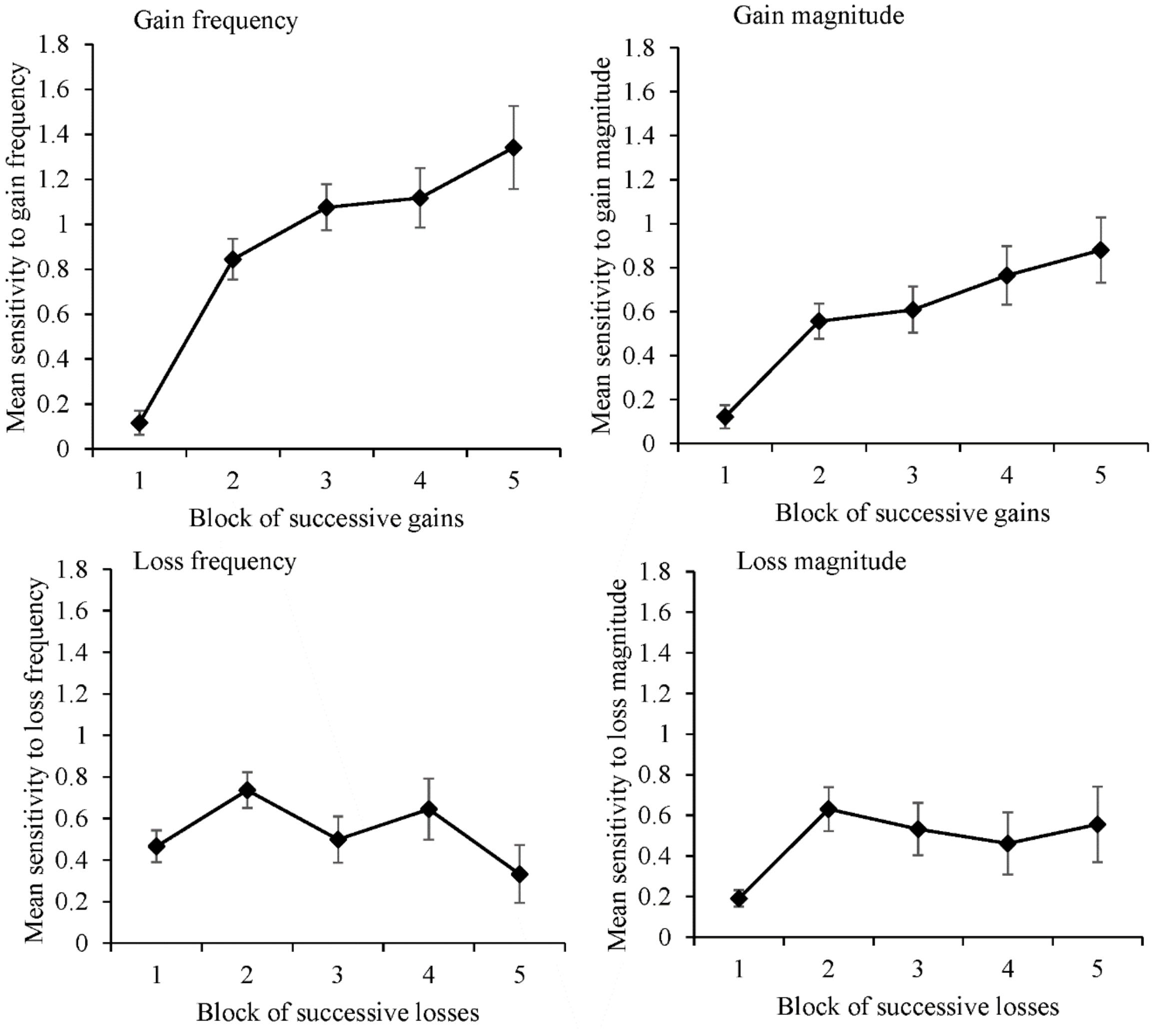

Figure 3.5. Mean sensitivity derived from Generalized Matching Law plotted as a function of successive blocks of primary outcomes. Error bars are the standard error of the mean.

In contrast to Bull et al. (2015), the two loss conditions in our sample did not show significant linear trends. For loss frequency, we observed a significant effect of block, but examination of contrasts showed no significant linear trend across the blocks. For loss magnitude, we observed no significant effect of block.

Following Bull (2013), we derived a group measure of sensitivity by averaging individual data and re-fitting the GML to these data. Bull totalled the number of responses and primary outcomes received on each alternative across the components and participants, logged these group ratios and fitted the GML to generate four estimates of sensitivity, bias and R^2. In our version of this analysis, we averaged individual response and reinforcer ratios and fitted the GML to this group data in order to preserve the individual ratios rather than the absolute number of responses/reinforcers. Figure 3.6 plots the gain and loss frequency data on the left, and gain and loss magnitude data on the right. Aggregate data were well described by the GML, similar to Bull, but did not correspond to our gain-loss asymmetry predictions. However, it did correspond to the patterns in group data observed by Bull, with higher sensitivity to losses than gains and about equivalent sensitivity between gain and loss magnitude.

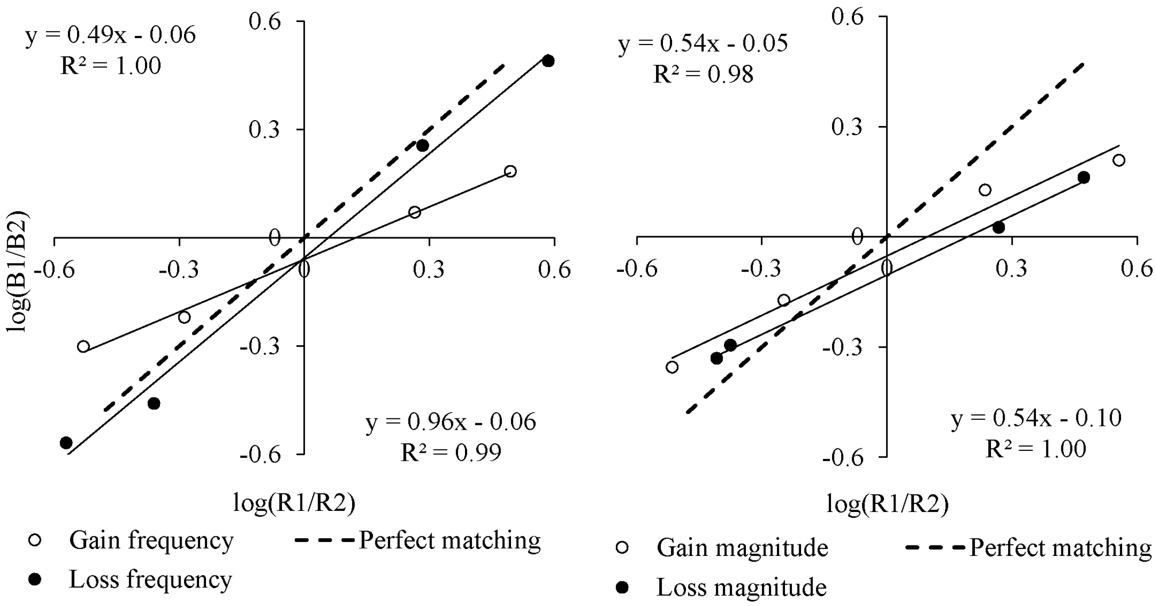

Figure 3.6. Logged ratio of responses on the y axis plotted against the logged ratio of reinforcers on the x axis, derived from average response and reinforcer ratios across components and participants. Dashed line corresponds to perfect matching. The solid lines are the Generalized Matching Law fitted to gain (white) and loss (black) data points. The equations of the regression line of best fit and the R_2 are shown for gains (top left of each graph) and for losses (bottom right of each graph). Left graph plots group data sensitivity to gain and loss frequency. Right graph plots group data sensitivity to gain and loss magnitude.

Following Bull et al. (2015), we attempted to derive a single estimate of each individuals' sensitivity to arranged outcomes that we could subsequently analyse for a gain-loss asymmetry. We considered whether averaging sensitivities across the last few blocks was suitable for our sample. All four conditions in Bull et al. showed a significant linear trend, with increasing sensitivities as the participants progressed through the blocks of reinforcers. In comparison, we observed significant linear trends for only two of the four conditions, suggesting that, in the loss conditions, the participants who progressed until the 5th block of outcomes did not necessarily increase in their sensitivity to changes in primary outcomes. Our rate of missing data was also higher than Bull et al., so averaging sensitivity estimates across the last few blocks would not only exclude a large proportion of participants, but it also wouldn't correspond to a point in the task where participants reached their highest sensitivity to obtained outcomes. Lastly, our R^2 were on average lower than Bull et al., and generally lower than average R^2 reported by Lie et al. (2009) and Krageloh et al. (2010), suggesting poor fit of GML to the reinforcer and response ratios split by block. Overall, our replication of Bull et al.'s analysis did not generate reliable sensitivity parameters for most blocks of responses.

Is there a gain-loss asymmetry in global sensitivity to reinforcement vs. punishment?

Due to the generally poor fit and a large proportion of missing data, we did not adapt Bull et al.'s (2015) approach of visually inspecting the graphs for final few blocks at which sensitivity reached stability. Below are several methods of deriving sensitivity that were aimed to include most of the participants' data and to produce better model fit based on higher R^2 values. We considered changes to both response and reinforcer ratios and ran a gain-loss asymmetry analysis where appropriate.

Sensitivity based on responses across last four blocks. We attempted an approach similar to Bull (2013) overall sensitivity estimate by adding up the responses at the final few blocks and re-fitting the GML to this range of data. Block 1 was excluded as a learning block, and the range of 2 to 5 was taken to account for a large number of participants not completing later blocks. This increased the number of participants with some measure of sensitivity after an initial exposure to primary outcomes, as we could now include participants who have at least reached their second block.

Table 3.7

Summary of GML-derived parameters using the last four blocks method.

	N	Mean a (SEM)	Median R^2 (Q_1, Q_3)	Mean log c (SEM)	% of incomplete 2nd blocks across components
Gain frequency	100	0.69 (0.07)	0.71 (0.30, 0.86)	-0.05 (0.03)	2.00 - 5.00
Loss frequency	91	0.85 (0.07)	0.83 (0.61, 0.95)	-0.08 (0.03)	6.59 - 13.19
Gain magnitude	96	0.67 (0.08)	0.69 (0.35, 0.88)	-0.05 (0.02)	3.13 - 7.29
Loss magnitude	88	0.64 (0.10)	0.57 (0.21, 0.84)	-0.08 (0.03)	6.82 - 10.23

Note. Q_1 = first quartile, Q_3 = third quartile.

Sensitivity, bias and R^2 are presented in Table 3.7. For the two loss conditions, mean sensitivity was higher than that derived from the 5th block of reinforcers in Table 3.5, but mean sensitivity for gains was lower. This most likely reflected the significant linear trend observed for the two gain but not the two loss conditions. For the two gain conditions, deriving a sensitivity measure across the increasing trend shown in Figure 3.5 resulted in decreased sensitivity parameters. Furthermore, examination of the means matched the pattern seen in Figure 3.6, with greater sensitivity to losses than gains and similar sensitivity to gain and loss magnitude.

Using responses from blocks 2-5 did not show a gain-loss asymmetry in either the frequency or the magnitude conditions (see top graph of Figure 3.7). A mixed measures ANOVA, with two within-subjects factors (task: frequency vs. magnitude and condition: gain vs. loss) showed no significant main effects of condition ($F(1, 77) = 2.45$, $p = .122$, , $\eta_p^2 = .031$), task ($F(1, 77) = 1.87$, $p = .176$, $\eta_p^2 = .027$) or interaction ($F(1, 77) = 1.55$, $p = .217$, $\eta_p^2 = .020$). This is demonstrated by the inset in the top left corner of the graph. This inset shows overlapping means for both tasks that are situated over the diagonal line indicating symmetrical sensitivities. Limiting the sample to GML fits with at least 0.50 R^2 similarly showed no significant difference between frequency gains ($M = 0.93$, $SE = 0.11$) and losses ($M = 1.00$, $SE = 0.08$; $t(49) = -0.63$, $p = .531$), and magnitude gains ($M = 0.86$, $SE = 0.11$) and losses ($M = 1.19$, $SE = 0.16$; $t(34) = -1.99$, $p = .055$).

Sensitivity based on responses at the last experienced block. We have established that most participants did not complete five blocks of primary outcomes. Furthermore, most of the components would have been terminated with incomplete blocks, where the participant experienced less than four of the scheduled primary outcomes for that block. Failing to complete the block due to time-out could result in relatively fewer responses than on other blocks. Alternatively, it could be an indication of developing exclusive preferences for one alternative based on the experience of outcomes in the previous blocks and relatively more responses after the last experienced outcome than on other blocks. Basing the response ratio on responses from the last four blocks avoids the former (with the exception of participants who did not complete their second block in at least one of the components; see the last column of Table 3.7), but not the latter issue. Lastly, while the last four blocks method does provide a measure of sensitivity for participants who did not complete five blocks, it results in less accurate final estimates for participants that did complete them. This is evident in the gain conditions in particular, with decreased sensitivity due to averaging across an increasing trend.

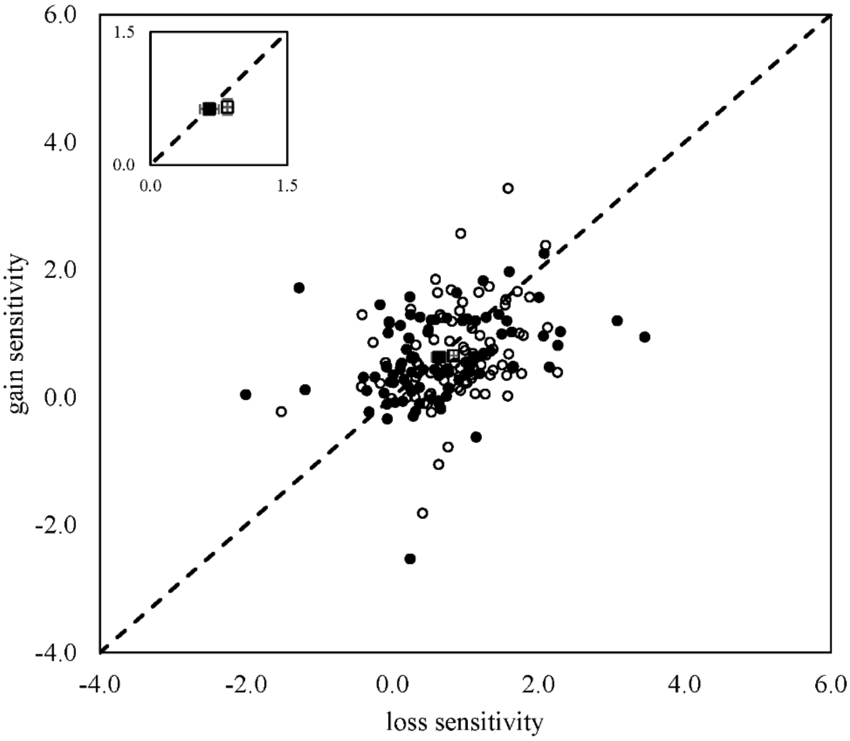

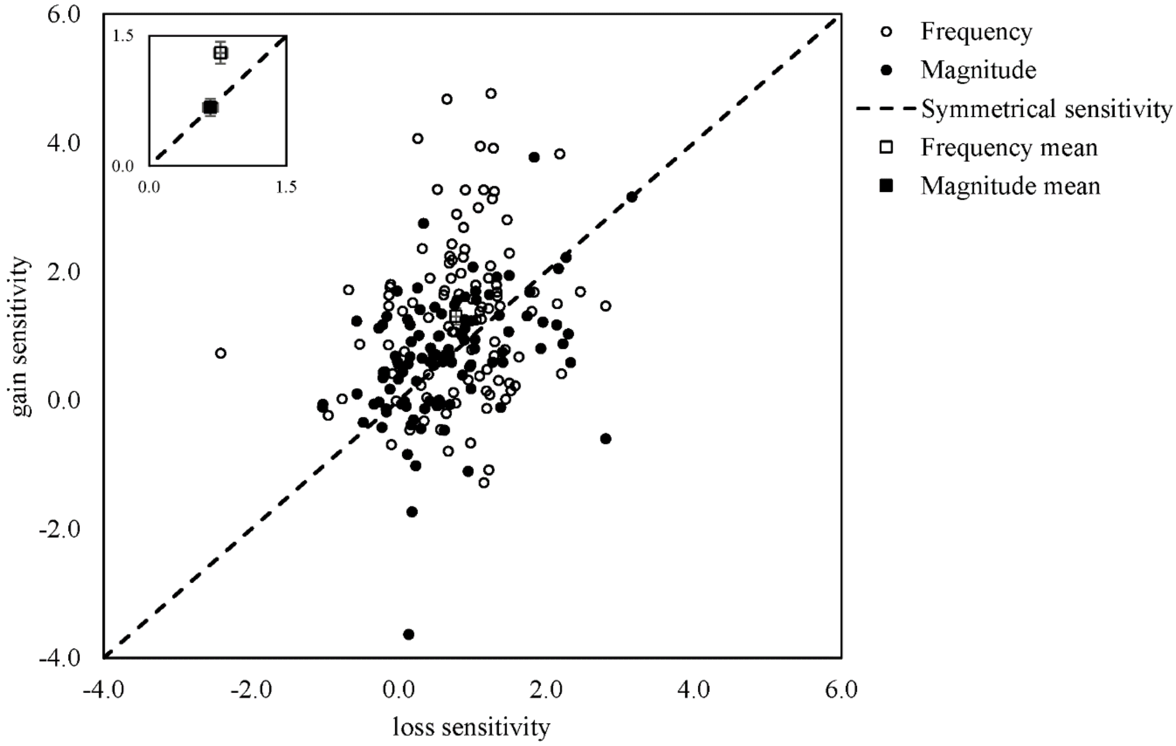

Figure 3.7. Sensitivity to gains plotted against sensitivity to losses for each individual. Black circles correspond to sensitivity to magnitude and white circles to sensitivity to frequency. The dashed diagonal line corresponds to symmetrical gain and loss sensitivity. Top graph plots sensitivity parameters derived from the last four block method and the bottom graph from the last experienced block method.

To get a measure that better accounts for final sensitivity to primary outcomes, we extracted the participants' allocation of responses that lead to experiencing their last five primary outcomes for each component, regardless of which block they were in, and combined them into one range of data (*last experienced* block). For example, the last experienced block for one component could comprise of responses after 15th-19th primary outcomes, while for another component it could comprise of responses after 4th-8th primary outcomes. In this example, the participant experienced 20 primary outcomes in their first component and it terminated upon collecting the 20th primary outcome, and in the second component they experienced 9 primary outcomes, which terminated via time-out. Note that this method calculates responses after a maximum of five primary outcomes, but excludes responses after the last primary outcome experienced. In the case of experiencing all the 20 primary outcomes, assuming all secondary outcomes have also been collected, the procedure would terminate and no responses were possible after the last outcome. In the case of experiencing fewer than 20 primary outcomes, the session would have terminated via timeout and the participant would have had a varying number of responses after the last outcome (relatively fewer or more than after other primary outcomes, as described above). Extracting their responses that lead to the last five primary outcomes standardizes the range of data by excluding the variation that would have been introduced by including responses after the last primary outcome experienced.

The GML was fitted to responses from the last experienced block of primary outcomes for each condition and each participant and the primary outcome ratio which consisted of all the primary outcomes experienced. Sensitivity, bias and R^2 for the last experienced block are presented in Table 3.8. We were able to calculative sensitivity for almost all the participants, with the exception of two in the loss magnitude condition, which was an improvement over the last four blocks method. The last column of Table 3.8 specifies the percentage of sensitivity estimates that relied on less than five primary outcomes experienced. As shown in Table 3.5 above, participants were generally less likely to complete the loss rather than the gain conditions. Notably, although most individuals now had a sensitivity estimate, the median R^2 were significantly lower than in the last four blocks method for the loss frequency (Bonferroni adjusted alpha of .0125; $z = -3.59$, $p < .001$) and the gain magnitude ($z = -4.13$, $p < .001$) conditions, but not for the gain frequency ($z = -0.89$, $p = .371$) and loss magnitude ($z = -1.46$, $p = .145$) conditions.

Table 3.8

Summary of GML-derived parameters using the last experienced block method.

Last experienced block					
	N	Mean a (SEM)	Median R^2 (Q_1, Q_3)	Mean log c (SEM)	% of less than 5 primary outcomes experienced across components
Gain frequency	103	1.31 (0.12)	0.55 (0.32, 0.83)	-0.05 (0.04)	0 - 1.94
Loss frequency	103	0.80 (0.07)	0.66 (0.34, 0.88)	-0.05 (0.03)	3.88 - 9.71
Gain magnitude	103	0.66 (0.10)	0.51 (0.14, 0.79)	-0.05 (0.04)	1.94 - 3.88
Loss magnitude	101	0.67 (0.08)	0.46 (0.14, 0.77)	-0.08 (0.03)	6.80 - 10.68

Note. Q_1 = first quartile, Q_3 = third quartile.

Using the last experienced block method showed a gain-loss asymmetry in sensitivity when responding to frequency changes, but not magnitude. The bottom graph of Figure 3.6 plots loss against gain sensitivity parameters for frequency and magnitude tasks. A mixed measures ANOVA showed a significant main effect of condition (F(1, 100) = 11.28, p = .001, η_p^2 = .101), task (F(1, 100) = 16.66, p < .001, η_p^2 = .143) and a significant interaction between condition and task (F(1, 100) = 8.81, p = .004, η_p^2 = .081). This significant interaction is demonstrated by the inset in the top left corner. Simple main effect analysis showed that in the frequency conditions, participants had higher sensitivity to gains (*M* = 1.30, *SE* = 0.13) than losses (*M* = 0.78, *SE* = 0.07; p < .001), and in the magnitude conditions there was no significant difference between the sensitivity to gains (*M* = 0.67, *SE* = 0.10) and losses (*M* = 0.67, *SE* = 0.08; p = .961). Limiting the sample to GML fits with at least 0.50 R^2 confirmed these relationships: sensitivity to gain frequency (*M* = 1.75, *SE* = 0.22) was higher than loss frequency (*M* = 1.15, *SE* = 0.08; t(42) = 2.58, p = .013), and no significant difference between gain (*M* = 1.16, *SE* = 0.16) and loss (*M* = 1.31, *SE* = 0.16) magnitude sensitivities was observed (t(28) = -0.79, p = .438).

Sensitivity to net outcomes. We considered whether model fit can be improved by adjusting our reinforcer ratios with inclusion of both primary and secondary outcomes according to the direct-suppression GML (Equation 3.2).

We fitted the direct-suppression version of the GML to the responses using both the last four blocks and the last experienced block methods. As noted by Critchfield et al. (2003) and Klapes, Riley and McDowell (2018), combining gains and losses presents additional challenges as it introduces the possibility of negative net outcomes and a log reinforcer ratio that cannot be calculated. For most participants, net outcomes were generally positive for both left and right alternatives (that is, they gained more points than they lost on each alternative). In cases where either left or right alternatives resulted in a zero (either no outcomes experienced or participants gained as many point as they lost) or negative (gained fewer points than they lost) value and the log reinforcer ratio could not be calculated, it was adjusted such that the resultant ratio would correspond to expecting more responses on the side with higher net. If one of the net outcomes was zero, both alternatives were adjusted by 0.25 in the direction of the valence of the other non-zero alternative: if it was negative, then both sides had 0.25 subtracted, and if it was positive, then both sides had 0.25 added. If the net for one alternative was positive and for the other negative, the negative alternative was set to 0.25. If the net for both alternatives was negative, the final ratio was multiplied by -1 because this produced a ratio that predicted more responses on the side with fewer/lower negative outcomes. Notably, this approach is unlike that of Critchfield et al. and Klapes et al., who chose to exclude such cases from analysis.

Comparison of primary outcome to net ratios showed a wider range of ratios obtained in the loss as compared to gain conditions. Figure 3.8 plots the primary outcome ratios on the x axis against the net ratios on the y axis for each of the conditions, with the black solid line corresponding to the ratios where all of the scheduled outcomes were obtained. Since frequency and magnitude graphs showed similar patterns, we will focus on the comparison of gains and losses in the frequency graphs. In the gain condition, most of the data were along the black line, corresponding to most individuals collecting both the primary and secondary outcomes. In the loss condition, there was greater variability in outcome ratios, corresponding to most individuals not experiencing all of the outcomes and producing acquired reinforcer ratios different from the scheduled. In both gain and loss conditions, there was a greater variability in ratios in the more extreme component (25:75), which most likely corresponded to a higher chance of developing exclusive preference, therefore suspending further outcome delivery, in the condition where it was easier to determine which side had more gains or fewer losses.

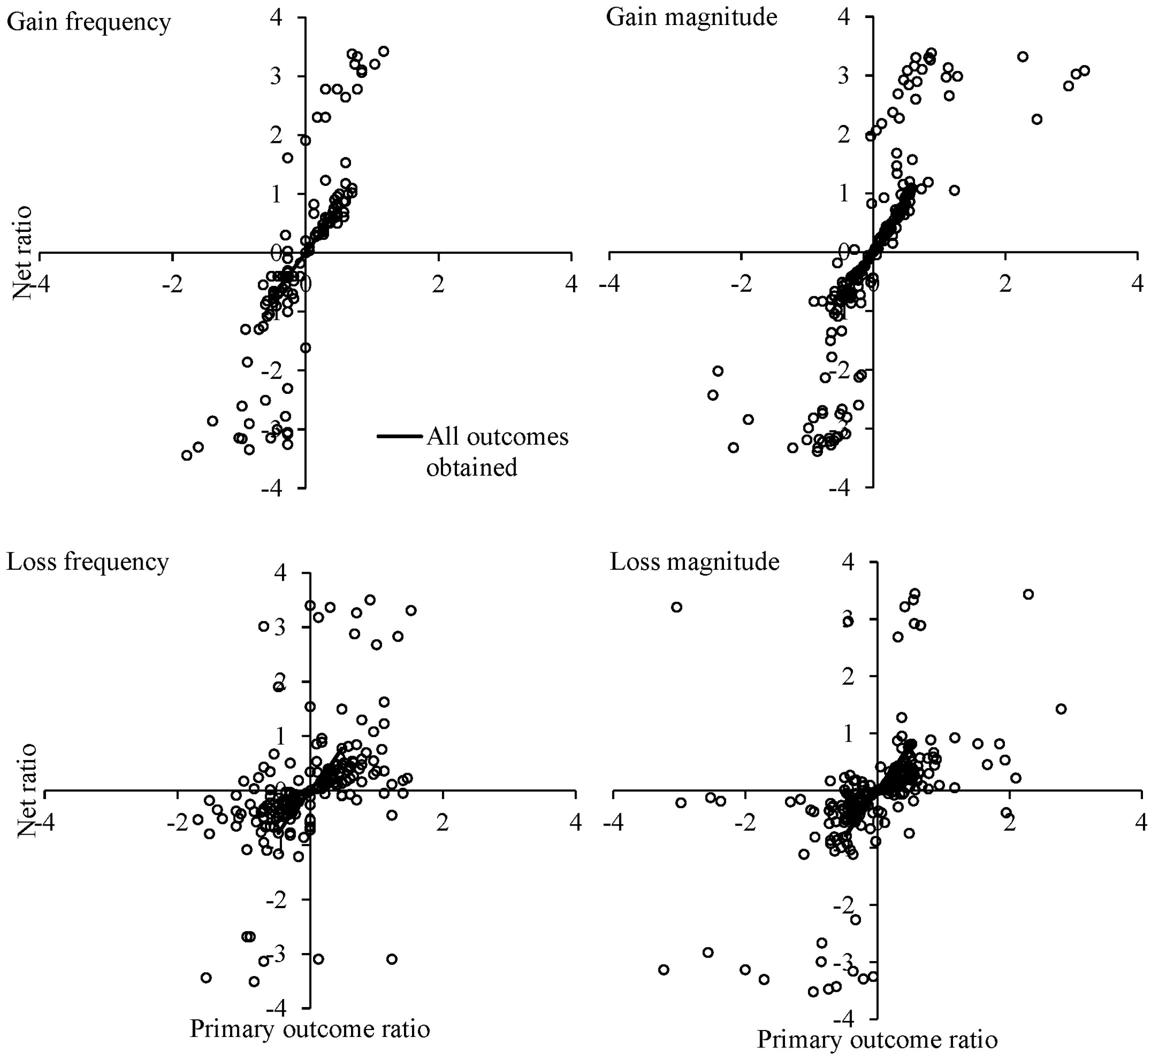

Figure 3.8. Net outcome ratios plotted against primary outcome ratios. Each data point corresponds to one individual's net and primary outcome ratios at one of the four components experienced in a given condition. Each individual has four data points representing their ratios in four components in each graph. The solid line corresponds to ratios where all of scheduled outcomes were obtained.

Using net as compared to primary outcome ratio did not improve model fit when using the last four blocks or the last experienced block methods. Median R^2 in Table 3.9 show that for both estimates of response allocation, using net ratio did increase R^2 relative to Tables 3.7 and 3.8. For the last four blocks method, using net instead of primary outcome ratio resulted in significantly lower R^2 for the loss frequency (Bonferroni adjusted alpha of .0125; z = -2.71, p = .007), but not the gain frequency (z = -2.17, p = .030), gain magnitude (z = -.29, p = .773) and loss magnitude (z = -2.47, p = .014) conditions. For the last experienced block, using net reinforcers produced significantly lower R^2 than using the primary outcome ratio for the loss frequency (Bonferroni adjusted alpha of .0125; $z = -2.83$, $p = .005$), but not the gain frequency ($z = -1.21$, $p = .227$), gain magnitude ($z = -0.94$, $p = .348$) and loss magnitude ($z = -2.08$, $p = .038$) conditions. Given the generally lower R^2 in the net as compared to the primary outcome analysis for both methods of measuring response allocation, we did not calculate gain-loss asymmetry using this method.

Table 3.9

Summary of GML-derived parameters for the last four blocks and the last experience block methods using net outcome ratio.

	N	Mean *a* (SEM)	Median R^2 (Q_1, Q_3)	Mean log c (SEM)
Last four blocks relative to net outcome ratio				
Gain frequency	100	0.35 (0.04)	0.65 (0.24, 0.86)	-0.06 (0.02)
Loss frequency	91	0.88 (0.10)	0.74 (0.46, 0.89)	-0.10 (0.03)
Gain magnitude	96	0.30 (0.04)	0.70 (0.34, 0.88)	-0.05 (0.03)
Loss magnitude	88	0.56 (0.10)	0.47 (0.16, 0.74)	-0.09 (0.04)
Last experienced block relative to net outcome ratio				
Gain frequency	103	0.74 (0.08)	0.52 (0.32, 0.80)	-0.07 (0.03)
Loss frequency	103	0.69 (0.11)	0.52 (0.23, 0.80)	-0.05 (0.03)
Gain magnitude	103	0.34 (0.05)	0.44 (0.16, 0.77)	-0.04 (0.04)
Loss magnitude	103	0.58 (0.09)	0.43 (0.12, 0.70)	-0.05 (0.03)

Note. Q_1 = first quartile, Q_3 = third quartile.

Is there a gain-loss asymmetry in the local effects of gains vs. losses?

We conducted an analysis focusing on behaviour after any response that produced a primary outcome (real outcomes) and after responses that could have produced a primary outcome, but didn't (hypothetical outcomes), for patterns that suggested a gain-loss asymmetry in the immediate response to gains and losses.

Responses after real outcomes. Responses after each primary outcome were numbered in ascending order starting from one for the immediate response after outcome delivery and resetting upon delivery of the next outcome. For our analysis, we chose to focus on up to thirty responses after each primary outcome. Within each condition, the number of left and right responses at each response count (1-30) was counted separately for the outcomes occurring on the left and on the right across the four components. The number of responses at higher counts were expected to be fewer than at lower counts since the count restarted each time the participant received another primary outcome. Therefore, the response counts were binned in order to ensure that sufficient and roughly equivalent number of responses occurred at each bin relative to the average total responses made. Examination of approximately 20 data sets showed that grouping responses into the following response bins would satisfy these criteria: 1, 2-4, 5-7, 8-10, 11-13, 14-17, 18-21, 22-25 and 26-30. As above, we added 0.25 to responses on both alternatives if either had zero responses. Two logged response ratios of left to right responses were calculated for each condition at each response bin, one after a left outcome and the other after a right outcome.

Responses after hypothetical outcomes. We wrote an Excel VBA Macro to program hypothetical primary outcomes (HPOs) and arranged them after responses that did not produce a primary outcome in each participants' acquired data, but could have occurred if the participant experienced a different sampling of the VI schedule. For each participant, HPOs were simulated a 100 times, producing a 100 instances of possible sequences of arranged HPOs. The procedure for each simulation was identical.

The HPOs were arranged largely as the primary outcomes[6] in the ACT. A single base VI4 schedule (randomly generating a number between 1 and 7) determined the interval after which the HPO was available and it was probabilistically allocated to either the left or right deck using the arranged contingency in each component (e.g. 25% chance on the left and 75% on the right; see Table 3.4 for exact contingencies). The macro then read each line of the participant's responses until the specific schedule conditions were met. If the scheduled time has passed and a response was made on the deck the HPO was scheduled to, then a HPO could be assigned within certain additional constraints. As in the ACT, a two-second COD was implemented and a HPO could not be allocated within two seconds of switching decks. A HPO also had to be more than eight seconds after a primary outcome. This constraint was based on the work of McLean et al. (2014) and Gomes-Ng et al. (2017) and was aimed to minimize the effect of primary outcomes on the HPOs. Gomes-Ng et al. employed a 30-second constraint on data derived from non-human animal subjects with multiple sessions. Our rapid-acquisition procedure had fewer sessions and primary outcomes were arranged on a generally richer schedule. Therefore, a shorter duration of constraint was more appropriate in order to balance minimizing the impact of primary reinforcers with ensuring sufficient opportunities for a HPO to be arranged. We also carried out the same procedure, but without

[6] Only the primary outcome that was manipulated in each condition was arranged and no secondary outcomes were scheduled. We chose not to schedule secondary outcomes as our local effects analysis was concerned with the effects of the manipulated gains/losses on responding in each condition. In this data analysis, we also did not schedule different magnitudes of gains and losses for the magnitude conditions (see below for the magnitude specific analysis).

an 8-second constraint. The patterns discussed below were largely maintained, but the difference between responses after real and hypothetical outcomes was reduced due to the introduced possibility of overlap. Since the patterns were maintained with and without the constraint, we chose to present data with the constraint for consistency with literature.

If the conditions for assigning a HPO at a given response were not met, the macro checked the next response for the same criteria until the specific conditions were met. If the necessary conditions were met and a HPO was recorded, the time and side for the next HPO was recomputed. This continued until the end of the participant's responses or if the maximum number of HPOs was obtained. Response distributions for HPOs were calculated as for primary outcomes.

Frequency conditions. The top panel of Figure 3.9 plots the mean log response ratios of left to right responses after a primary outcome as a function of successive response bins. Considering responses after a primary outcome, for both gains and losses there was an apparent preference for the just productive alternative. This preference was greater following gains than losses. As noted in the introduction, whether this is an effect unique to primary outcomes needs to be established by comparing to responses after HPOs. The middle panel of Figure 3.9 plots the mean log response ratios after HPOs. The patterns of responding after HPOs were similar in shape across gain and loss conditions, with a preference for the just-productive side that gradually decreases. Responding on the loss condition also showed a slightly lower overall tendency to respond on the just-productive side as compared to the gain condition. This pattern in the responses after HPOs shows that participants tended to switch more in the losses rather than gains condition, with overall shorter visits to each side. Importantly, both of these patterns differed from responding after primary outcomes, indicating that the primary outcome did have unique effects on behaviour. The difference between responses after primary outcomes and after HPOs are plotted in the bottom panel of Figure 3.9. The further the data points were from a mean difference of zero (indicated by the dotted line) the greater the difference between responses after primary and hypothetical outcomes.

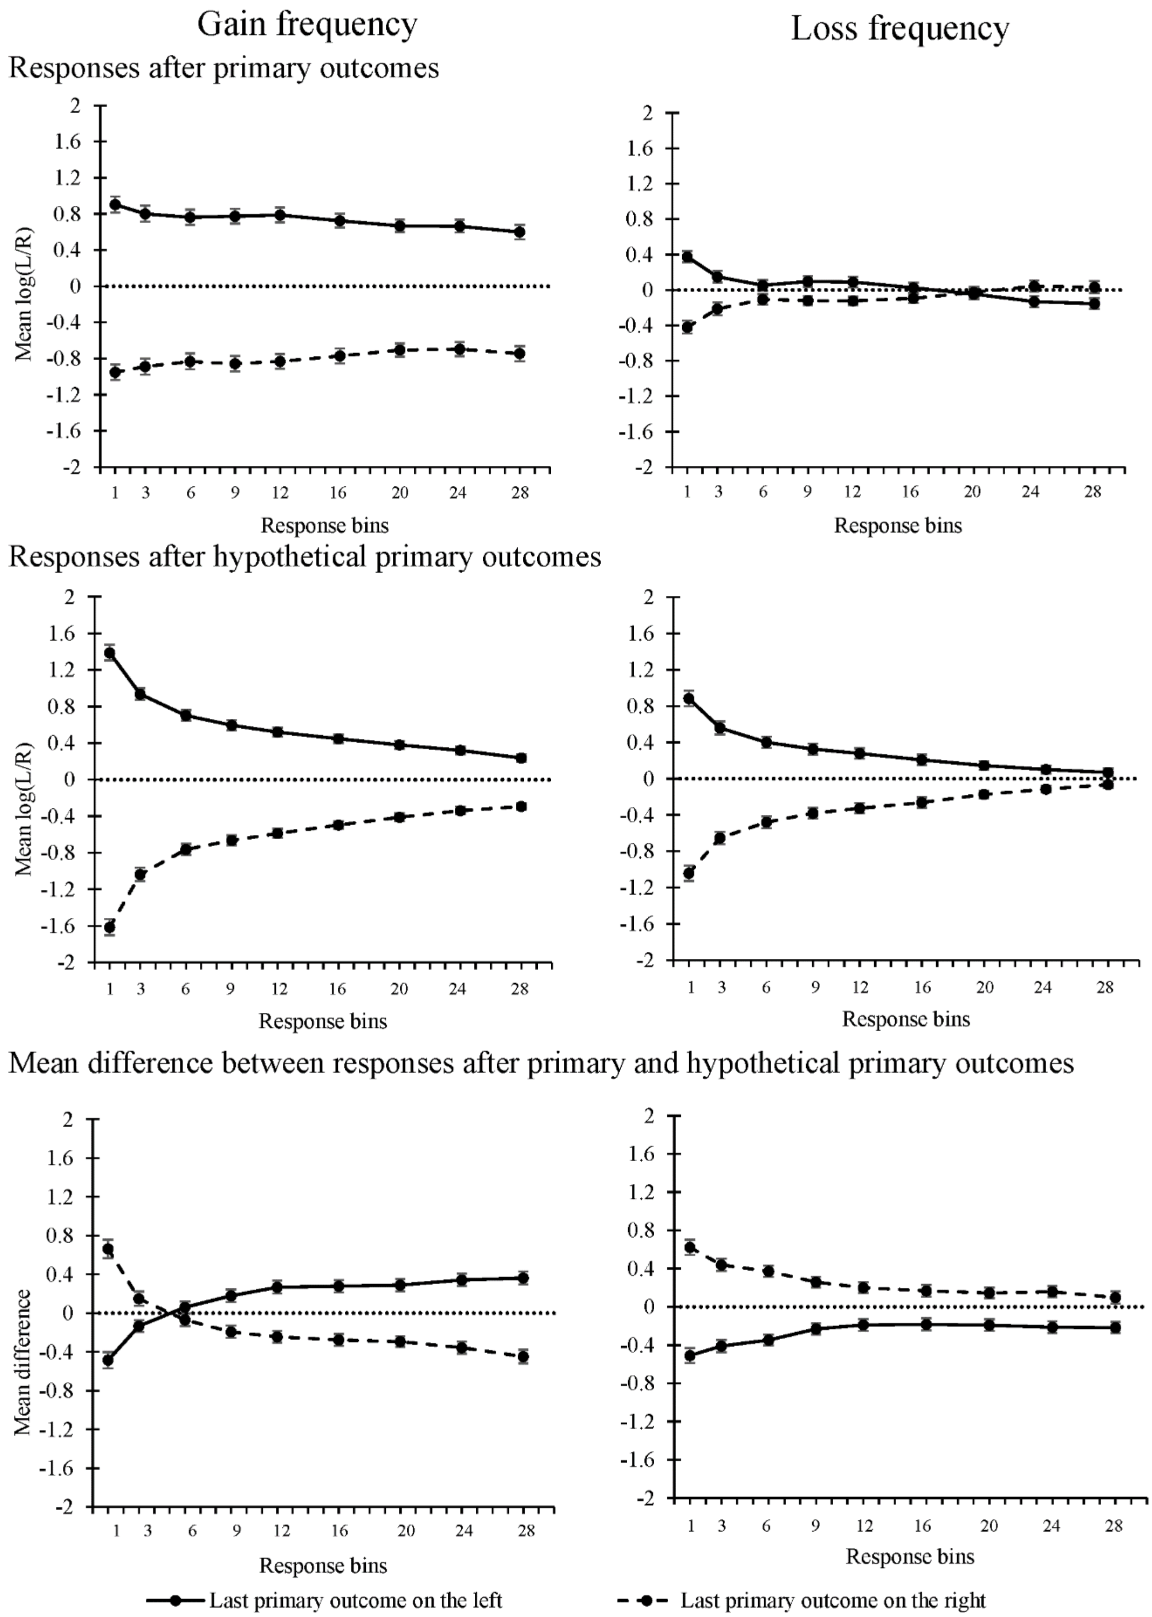

Figure 3.9. Top two rows: Mean logged ratio of left to right responses after an outcome on the left (solid line) and on the right (dashed line) as a function of successive response bins. The dotted line at y axis value of zero demarcates a ratio where number of left responses equals to number of right responses. Error bars are standard error of the mean. Bottom row:

Mean difference between response after hypothetical primary outcomes and primary outcomes after an outcome on the left (solid line) and right (dashed line) as a function of successive response bins. The dotted line at y axis value of 0 demarcates a mean difference of zero. Error bars are standard error of the mean. For all rows, data on the left are from the gain frequency task and data on the right are from the loss frequency task.

Comparison of the mean difference graphs showed similar initial response to primary outcomes followed by a gain-loss asymmetry in the pattern of subsequent responses. Immediate responses after both gains and losses showed a decreased tendency to respond on the just-productive side. That is, encountering any primary outcome resulted in a higher likelihood of switching to the other side, while the overall ratio of responses still favoured the just-productive side. Subsequent responding for gains showed an increasing tendency to respond on the just-productive side, while for losses the decreased tendency to respond on the just-productive side maintained across the 30 responses. Notably, the effect of gains on behaviour seemed to increase as a function of responses (in that the mean difference was increasing further away from the dotted line), but the effect of losses on behaviour seemed to decrease as a function of responses (in that the mean difference was decreasing closer to the dotted line).

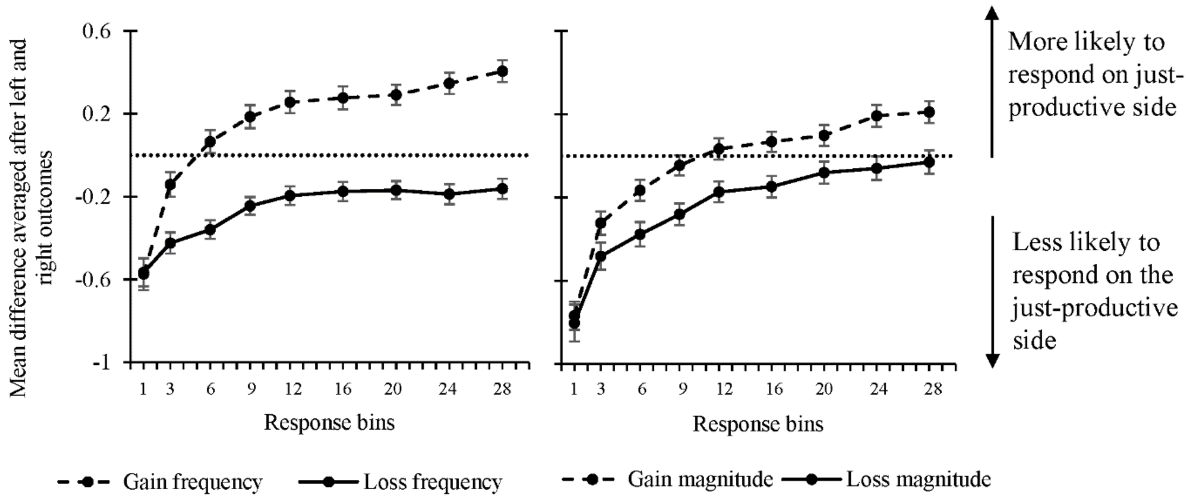

Figure 3.10. Mean differences between the ratio of responses on the just-productive side over responses on the alternative side after hypothetical primary outcomes and primary outcomes as a function of successive response bins. Note that unlike previous graphs, left and right outcomes are collapsed together. Dashed lines correspond to gain frequency (left) and magnitude (right) data, and sold lines correspond to loss frequency (left) and magnitude

(right) data. The dotted line at y axis value of zero demarcates a mean difference of zero. Error bars are standard error of the mean.

The observed asymmetry was examined using a repeated measures ANOVA. For each condition, the mean difference after an outcome on the right was reversed in sign and did not significantly differ from mean difference after outcomes on the left. We therefore averaged the mean difference after right (sign reversed) and left outcomes (see left panel of Figure 3.10) and ran a repeated measures ANOVA, with block (9 levels) and condition (gain vs. loss; see Table 3.10). Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of block and the interaction, therefore degrees of freedom were corrected using the Greenhouse-Geisser estimate of sphericity. There were significant main effects of block, condition and a significant interaction. Pairwise comparisons using the Sequential Bonferroni correction showed no difference between gains and losses at the first response (p = .967) and a significantly higher mean difference for gains than losses at all of the other response bins (p < .001 at subsequent bins). Figure 3.10 demonstrates this with a steeper decline in mean difference for gains as opposed to losses after the first response bin.

Table 3.10

Study 2 sphericity and ANOVA statistics for the gain-loss asymmetry analysis of the mean differences in the frequency conditions.

	Sphericity				ANOVA			
	χ^2	df	p	ε	F	df	p	η_p^2
Block	690.21	35	<.001	0.32[a]	93.00	2.60, 254.28	<.001	.487
Condition	-	-	-	-	50.23	1, 98	<.001	.339
Block by Condition	654.70	35	<.001	0.34[a]	13.12	2.74, 267.99	<.001	.118

Note. a = Greenhouse-Geisser correction.

Magnitude conditions. Figure 3.11 shows the mean log response ratios after primary outcomes (top panel), HPOs (second panel) and the mean difference (third panel) for the two magnitude conditions. The differences between the gain and loss conditions were less distinct in both the primary outcomes and HPO graphs as compared to the frequency conditions. However, both of the primary outcomes response patterns differed from HPOs response patterns.

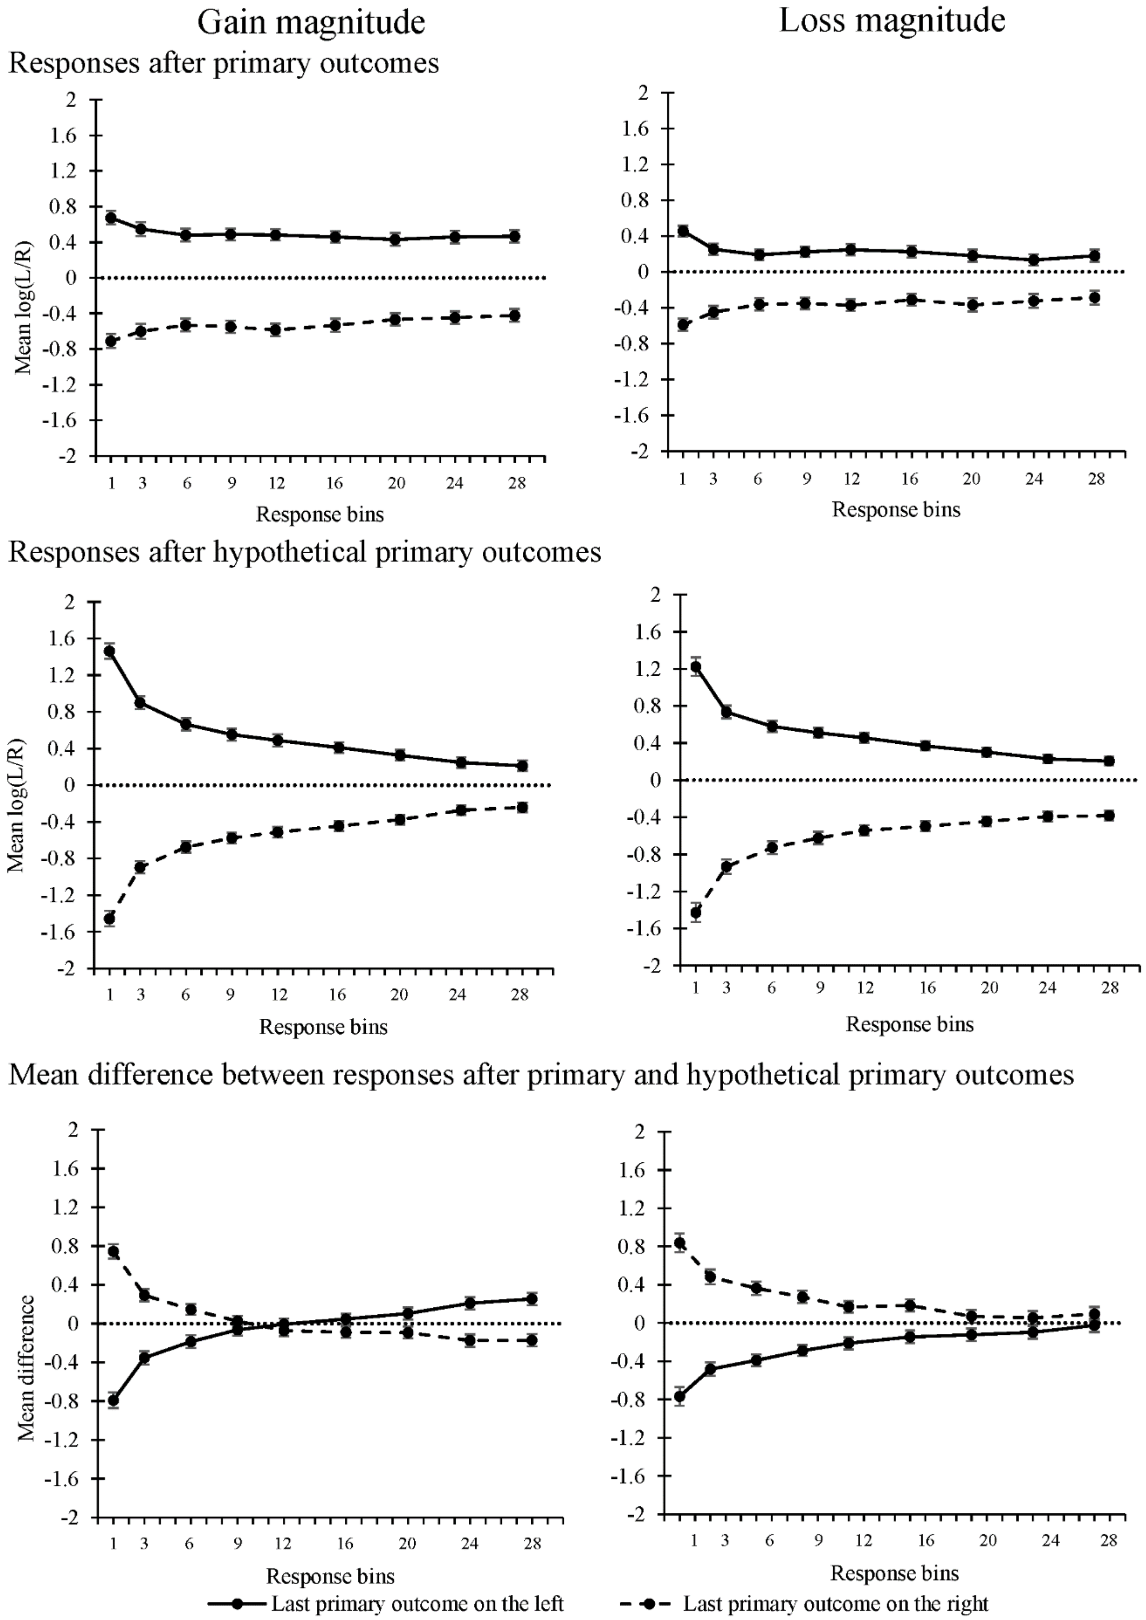

Figure 3.11. Top two rows: Mean logged ratio of left to right responses after an outcome on the left (solid line) and on the right (dashed line) as a function of successive response bins. The dotted line at y axis value of zero demarcates a ratio where number of left responses equals to number of right responses. Error bars are standard error of the mean. Bottom row:

Mean difference between response after hypothetical primary outcomes and primary outcomes after an outcome on the left (solid line) and right (dashed line) as a function of successive response bins. The dotted line at y axis value of 0 demarcates a mean difference of zero. Error bars are standard error of the mean. For all rows, data on the left are from the gain magnitude task and data on the right are from the loss magnitude task.

Comparison of the mean difference graphs indicated an initial lower preference for the just-productive side after both gains and losses and asymmetry in subsequent responses. As with the frequency conditions, immediate responses after both gains and losses showed a lower tendency to respond on the just-productive side, but this seemed to last across more responses than in the frequency conditions. Subsequent responding for gains showed an increasing tendency to respond on the just-productive side, which was more gradual than observed in the gain frequency condition. Subsequent responding for losses showed a decreased tendency to respond on the just-productive side, as with frequency losses. The effect of gains and losses not differentiated by magnitude followed the patterns observed in the frequency conditions, but were overall less extreme relative to the general visit structure.

Table 3.11

Study 2 sphericity and ANOVA statistics for the gain-loss asymmetry analysis of the mean differences in the magnitude conditions.

	Sphericity				ANOVA			
	χ^2	df	p	ε	F	df	p	η_p^2
Block	868.34	35	<.001	0.27[a]	119.74	2.16, 207.71	<.001	.555
Condition	-	-	-	-	7.45	1, 96	.008	.072
Block by Condition	761.24	35	<.001	0.31[a]	2.45	2.48, 238.39	.076	.025

Note. a = Greenhouse-Geisser correction.

The delayed start to the asymmetry and its general lesser extent was confirmed with a repeated measures ANOVA with block (9 levels) and condition (gain vs. loss) comparing mean difference averaged across right (sign reversed) and left outcomes as described above (see the right panel of Figure 3.10). Mauchly’s test indicated that the assumption of sphericity had been violated for the main effect of block and the interaction, therefore degrees of freedom were corrected using the Greenhouse-Geisser estimate of sphericity. There were significant main effects of block, condition and a non-significant interaction. Examination of mean differences for the nine blocks collapsed across the conditions showed a decrease from

1st to 9th, with a steeper decline in the first few blocks and shallower in the last few blocks, the pattern of which did not differ between gains and losses (right panel of Figure 3.10). Overall, mean differences for the gain blocks (M = 0.07, SE = 0.04) were lower than for the loss blocks (M = 0.26, SE = 0.05).

Is there a gain-loss asymmetry in the local effects of small vs. large magnitudes?

Our local analysis of responses after gains and losses on the magnitude conditions above did not differentiate by magnitude of the outcomes, hence we conducted a separate examination of a gain-loss asymmetry split by magnitude.

Data analysis. In Kubanek et al. (2015), participants experienced all of the possible magnitudes in each condition. This was not true of the distribution of magnitudes in the ACT, where each component had a different range of magnitudes, although the participants would have experienced all of the possible magnitudes by the end of the condition. Table 3.12 lists the scheduled magnitudes in each component; for example, component 1 (25v:75v) of the gain magnitude condition arranged the ten magnitudes in the first column on one of the alternatives and in the fourth column on the other alternative. Upon examination of approximately twenty data sets, we chose to create two magnitude bins (indicated by the bolded black line in the table) and five response bins in order to ensure that we have a number of responses in the response bins that is consistent with the local analysis done above. The two magnitude bins were created by splitting the magnitude range at its median (41) and the following response bins were used: 1, 2-6, 7-12, 13-18, and 19-30. Logged left over right responses at each response bin for each magnitude were generated for each participant. The HPOs were generated for the two magnitude conditions as described above, which this time also included varying magnitudes of primary outcomes as in the ACT.

Figure 3.12 plots the mean logged response ratios of left to right responses after primary outcomes (top row), HPOs (middle row) and the mean differences (bottom row) for the two magnitude bins (black circles for smaller magnitude and white squares for larger magnitude). Responses after HPOs showed a parallel decrease in responding on the just-productive side between the two magnitudes that differed in the extent of preference for the just-productive choice. This greater tendency to respond on the just-productive side after a larger than a smaller gain, and a smaller than a larger loss, was indicative of global patterns in behaviour that corresponded to how the magnitude bins were created. Majority of the larger magnitudes occurred on the side scheduling primary outcomes in the higher magnitude range (i.e. 65 and 75), resulting in a general tendency to switch less in the gain (white squares further from zero than black circles) and more in the loss condition (white squares closer to zero than black circles). Conversely, majority of the smaller magnitudes occurred on the side scheduling primary outcomes in the lower magnitude range (i.e. 25 and 35), thus resulting in more switching behaviour after a gain and less after a loss. Overall, responses after HPOs for the two magnitudes differed in the overall chances of switching, but not in the rate at which preference for the just-productive side decreased.

Table 3.12

Range of primary outcomes in each VI schedule of the magnitude conditions

VI Schedule	25	35	65	75
Magnitude 1	3	4	7	8
Magnitude 2	8	11	20	23
Magnitude 3	13	18	33	38
Magnitude 4	18	25	46	53
Magnitude 5	23	32	59	68
Magnitude 6	28	39	72	83
Magnitude 7	33	46	85	98
Magnitude 8	38	53	98	113
Magnitude 9	43	60	111	128
Magnitude 10	48	67	124	143

Examination of the mean difference graphs suggested an asymmetry in the effect of magnitude on responses after gains and losses. Responses after both gains and losses showed an immediate decreased tendency to respond on the just-productive side, with subsequent responses showing an increasing tendency to respond on the just-productive side. For gains, the difference between two magnitude bins was small across responses: no difference in immediate response, and subsequently some indication of a greater tendency to respond on the just-productive side after a larger magnitude outcome. For losses, immediate and subsequent responses after a larger magnitude loss showed lesser tendency to respond on the just-productive side as compared to smaller magnitudes. Furthermore, the difference between magnitudes across the response bins appeared to be consistent for losses, a parallel pattern that was also observed in the responses after HPOs.

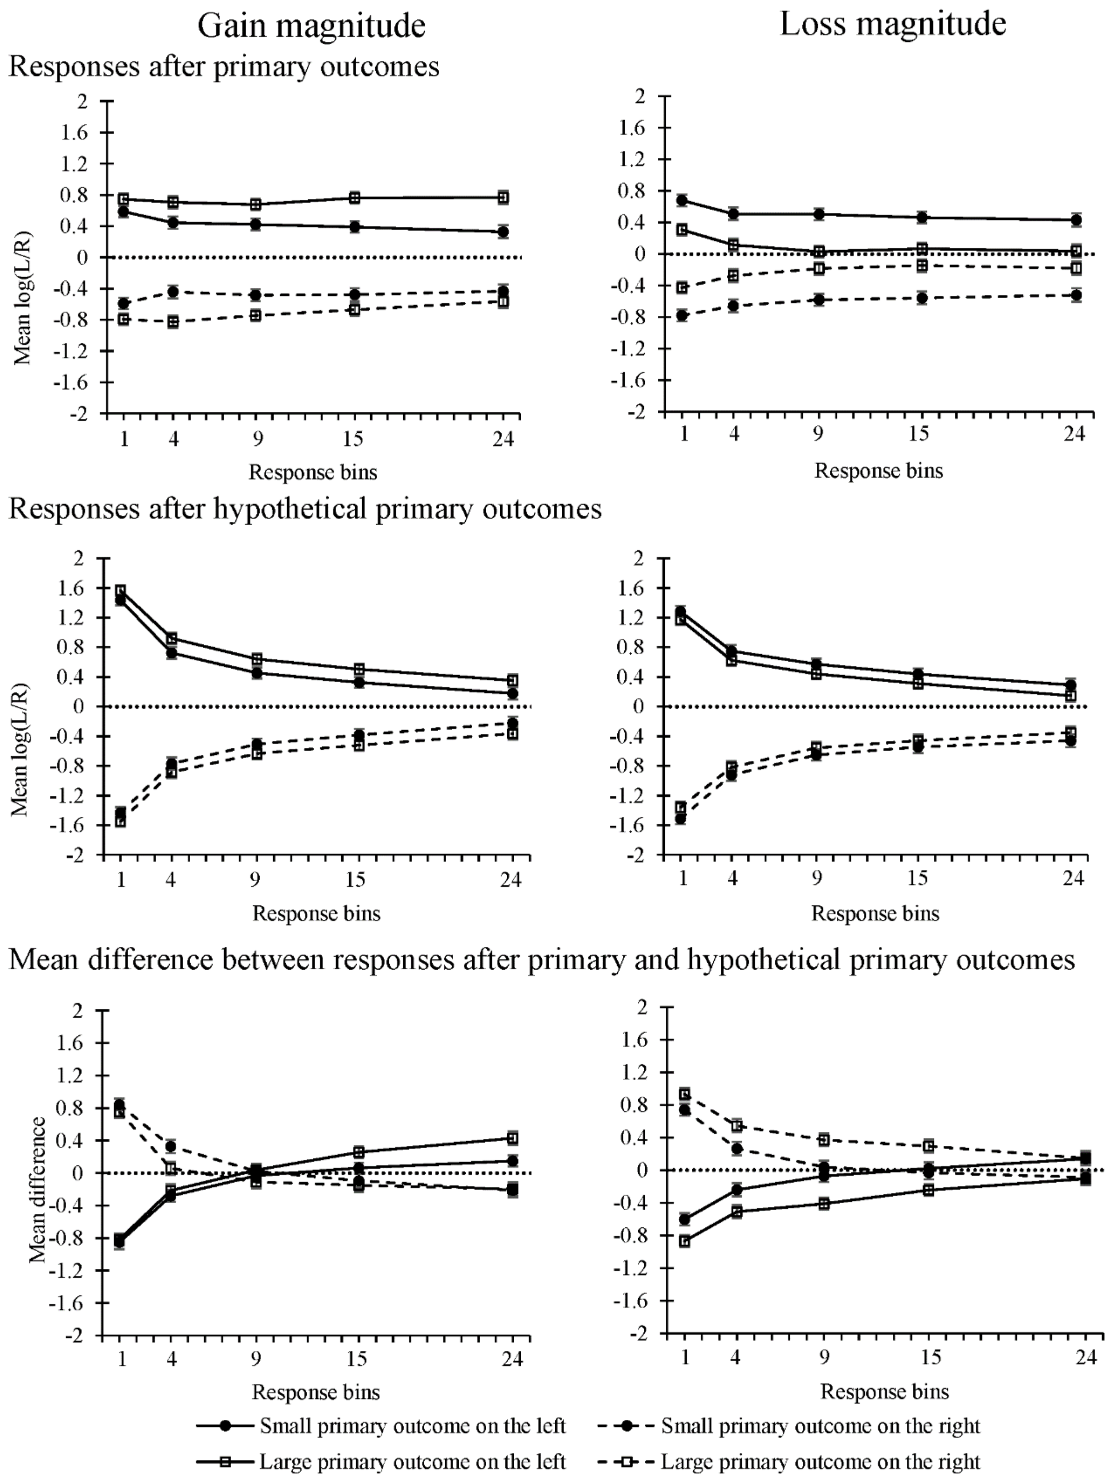

Figure 3.12. Top two rows: Mean logged ratio of left to right responses after an outcome on the left (solid line) and on the right (dashed line) as a function of successive response bins. Black circles correspond to primary outcomes in the small magnitude range and white squares correspond to primary outcome in the large magnitude range. The dotted line at y axis value of zero demarcates a ratio where number of left responses equals to number of right responses. Error bars are standard error of the mean. Bottom row: Mean difference

between response after hypothetical primary outcomes and primary outcomes after an outcome on the left (solid line) and right (dashed line) as a function of successive response bins. The dotted line at y axis value of zero demarcates a mean difference of zero. Error bars are standard error of the mean. For all rows, data on the left are from the gain magnitude task and data on the right are from the loss magnitude task.

We averaged the mean difference after right (sign reversed) and left outcomes and tested the relationship between magnitude, condition and block in a repeated measures ANOVA (see Table 3.13). Mauchly's test indicated that the assumption of sphericity had been violated for the main effect of block and all of the interactions that included block, therefore degrees of freedom were corrected using the Greenhouse-Geisser estimate of sphericity. There were significant main effects of condition, magnitude and block, as well as significant interactions between condition and magnitude, and condition and block[7]. We focused on the interaction between condition and magnitude.

Table 3.13

Study 2 sphericity and ANOVA statistics for the gain-loss asymmetry analysis of the mean differences split by magnitude of the primary outcome.

	Sphericity				ANOVA			
	χ^2	df	p	ε	F	df	p	η_p^2
Condition	-	-	-	-	3.98	1, 90	.049	.042
Magnitude	-	-	-	-	7.53	1, 90	.007	.077
Block	288.55	9	<.001	0.41[a]	158.44	1.64, 147.88	<.001	.638
Condition by Magnitude	-	-	-	-	44.41	1, 90	<.001	.330
Condition by Block	272.76	9	<.001	0.44[a]	5.84	1.75, 157.55	.005	.061
Magnitude by Block	107.03	9	<.001	0.64[a]	1.28	2.58, 231.73	.281	.014
Condition by Magnitude by Block	93.30	9	<.001	0.66[a]	1.49	2.62, 235.95	.223	.016

Note. a = Greenhouse-Geisser correction.

A pairwise comparison showed that a gain-loss asymmetry in the mean difference data, averaged across five blocks, emerged only at higher magnitudes. The top panel of

[7] Simple main effect analysis showed the same patterns already discussed in the local effects not differentiated by magnitude, where the difference between gain and loss conditions emerged at later blocks.

Figure 3.13 demonstrates a comparison between gains and losses at smaller magnitude (left) and larger magnitude (right). When magnitude of both gains and losses was small, there was no significant difference between gains (*M* = 0.18, *SE* = 0.05) and losses (*M* = 0.13, *SE* = 0.05; *p* = .512). This is demonstrated by the overlapping gain and loss data points in the top left graph. When the magnitude of both gains and losses was large, there was a significant difference between gains (*M* = 0.04, *SE* = 0.06) and losses (*M* = 0.42, *SE* = 0.06; *p* < .001). The top right graph illustrates this difference with the loss magnitude points generally higher than the gain magnitude points.

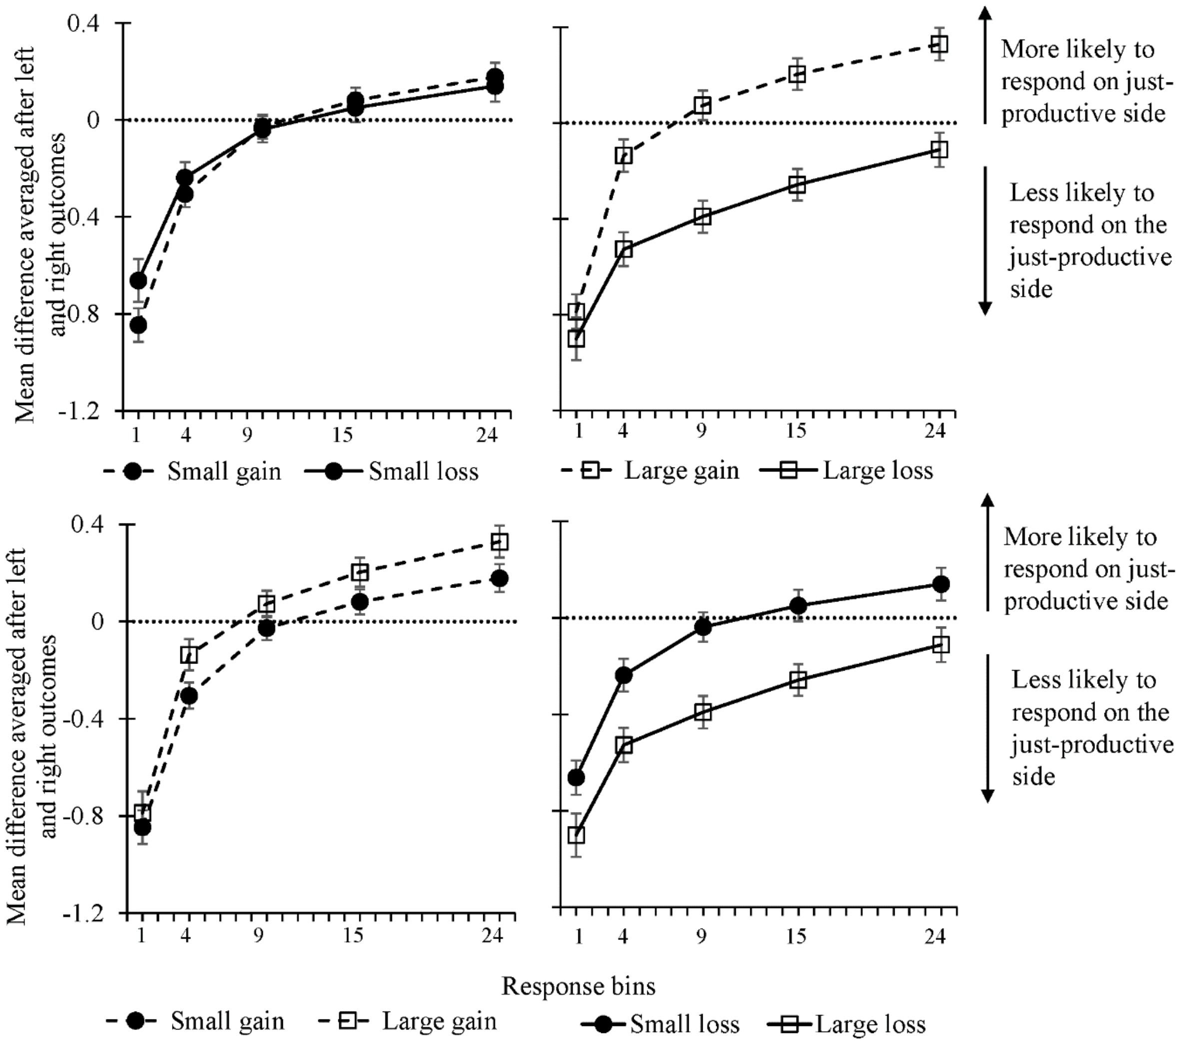

Figure 3.13. Mean differences between the ratio of responses on the just-productive side over responses on the alternative side after hypothetical primary outcomes and primary outcomes as a function of successive response bins. Note that left and right outcomes are collapsed together. Black circles correspond to primary outcomes in the small magnitude range and white squares correspond to primary outcome in the large magnitude range. The dotted line at

y axis value of zero demarcates a mean difference of zero. Error bars are standard error of the mean. Note that the bottom graphs re-plot the same data given on the top graphs to facilitate comparisons of magnitudes within gain and loss conditions.

Within each condition, responses after smaller magnitudes were significantly different from responses after larger magnitudes. The bottom panel of Figure 3.13 demonstrates a comparison between smaller and larger magnitudes for gains (left) and losses (right). In the gain condition, participants had significantly higher data points in the smaller than larger magnitude (p = .001). In the loss condition, participants had significantly lower data points in the smaller than larger magnitude (p < .001). Comparing across the four graphs, the difference between gains and losses emerged at higher magnitudes, which was largely due to a greater effect of magnitude on losses as opposed to gains.

Does order effect in discounting persist after balance reset and condition separation?

Data were analysed as described in Study 1, Experiment 1. Table 3.14 specifies the number of participants excluded based on the Johnson and Bickel (2008) criteria. Analysis below uses participants who had fully systematic data across conditions. Individual h parameters were not normally distributed (Shapiro-Wilk's test for normality all p < .001), with high positive skew and kurtosis for both gains and losses. Data were log transformed, with Shapiro-Wilk's test for normality remaining significant for gains (W = 0.96, p = .005) and losses (W = 0.95, p = .001). Examination of histograms showed no skew and slight positive kurtosis, but the shape of the distribution was otherwise normal.

Table 3.14

Number of participants (n = 103) in Study 2 who had unsystematic data by criterion, and total participants with systematic data.

	Gain	Loss
Criterion 1	3	3
Criterion 2	5	4
Total systematic	94 (91.26)	

Note. Percentages in parenthesis.

The hyperbolic model provided good fits to individual and group median indifference points (see Table 3.15), with medians for both group indifference points and individual h parameters showing steeper discounting of losses than gains as with the experiential money tasks in Study 1. Examination of mean residuals showed no systematic trends (see Appendix H).

Table 3.15

Study 2 money task h, R^2 and AUC values

	Group median indifference points			Individual participants		
Condition	*h*	R^2	*AUC*	*h* median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Gain	1.38	0.96	0.27	1.54 (0.99, 2.54)	0.85 (0.75, 0.92)	0.28 (0.01)
Loss	2.10	0.96	0.21	1.96 (1.34, 3.64)	0.87 (0.79, 0.93)	0.22 (0.01)

Note. Q_1 = first quartile, Q_3 = third quartile.

As in Study 1, we observed significant interaction between order and condition, and unlike Study 1, we also observed a significant main effect of condition (see top right section of Figure 3.14). Mixed measures ANOVA with order as the between-subjects factor and condition as the within-subjects factor showed a significant main effect of condition ($F(1, 92) = 9.23$, $p = .003$, $\eta_p^2 = .091$), non-significant main effect of order ($F(1, 92) = 1.65$, $p = .202$, $\eta_p^2 = .018$), and a significant interaction ($F(1, 92) = 4.72$, $p = .032$, $\eta_p^2 = .049$). The inset in the left panel of Figure 3.14 shows this interaction of order and condition; when losses were presented first, gains ($M = 0.24$, $SE = 0.05$) were discounted significantly less steeply than losses ($M = 0.28$, $SE = 0.05$; $p = .001$), but when gains were presented first, there was no significant difference between discounting of gains and discounting of losses ($p = .524$; Bonferroni correction applied). The log(*h*) parameters were significantly higher than zero for gains in the gains first order ($t(50) = 5.49$, $p < .001$, $d_z = 0.77$), losses in the gains first order ($t(50) = 5.24$, $p < .001$, $d_z = 0.73$), gains in the losses first order ($t(42) = 3.35$, $p = .002$, $d_z = 0.51$), and losses in the losses first order ($t(42) = 8.63$, $p < .001$, $d_z = 1.32$).

Contrary to Study 1 and the analysis using log(*h*) above, the same analysis conducted with AUC showed that losses were discounted more steeply than gains regardless of order. Mixed measures ANOVA showed a significant main effect of condition ($F(1, 92) = 13.95$, $p < .001$, $\eta_p^2 = .132$), non-significant main effect of order ($F(1, 92) = 1.48$, $p = .227$, $\eta_p^2 = .016$) and a non-significant interaction ($F(1, 92) = 3.23$, $p = .076$, $\eta_p^2 = .034$).

Gains and losses were not significantly correlated using log(*h*) ($r(92) = .02$, $p = .863$) or AUC ($r(92) = -.02$, $p = .817$), unlike the negative correlation in Study 1.

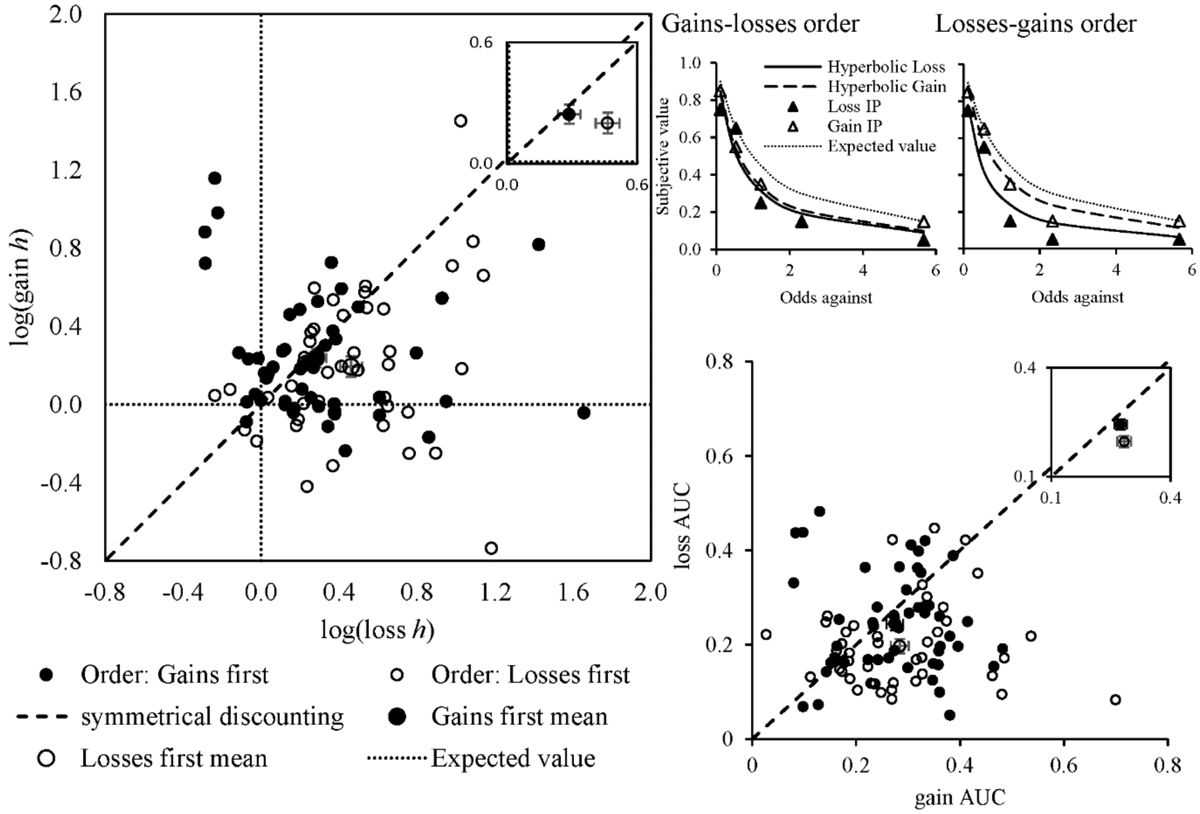

Figure 3.14. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and white circles represent the means for each order, which are also shown in the inset. The error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged *h* value (0) when decisions are made based on expected value. Top right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles). For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Bottom right panel: AUC for losses plotted against AUC for gains for each individual. Order split and group means are plotted as described for the left panel graph.

Does sensitivity predict discounting after controlling for condition order?

For gains and losses separately, we tested whether a model with order and sensitivity to last experienced block for both frequency and magnitude was a better predictor of discounting than order alone. We also conducted this analysis for both log(*h*) and AUC discounting measures, but since they supported the same conclusions, AUC results are reported in the Appendix F.

Losses. We conducted a two-stage hierarchical regression with log(*h*) for losses as the dependent variable and condition order entered into the first block (Model 1) and the sensitivities to loss magnitude and loss frequency derived from the last experienced block entered into the second block (Model 2).

We observed no high correlations between independent variables (see Table 3.16), with the only significant correlation between order and log(*h*), suggesting no multicollinearity between the predictors. Two cases were identified as outliers based on their standardized residuals being outside the -2 to 2 range (Field, 2009), but they were retained in the sample as the Cook's distance was less than 1 (Cook & Weisberg, 1982), Mahalanobis distance were less than 15 (Barnett & Lewis, 1978; Field, 2009) and the centred leverage values were less than twice the size of the average leverage (Hoaglin & Welsch, 1978), indicating that the identified cases were not undue influence on the model parameters. One additional case was identified by manual examination that exceeded both the Mahalanobis distance and leverage cut-offs, but with Cook's distance less than 1; we excluded this participant, re-ran the regression and it produced no change in the results. We chose to retain this case. Analysis of residual and scatter plots showed that the data were normally distributed and met the assumptions of linearity and homoscedasticity.

Table 3.16

Correlations between log(h) for losses, condition order, and sensitivities to loss frequency and magnitude

	Log(losses *h*)	Order	Loss frequency	Loss magnitude
Log(losses *h*)	-	-0.27*	-0.12	0.06
Order	-	-	-0.07	0.02
Loss frequency	-	-	-	-0.02
Loss magnitude	-	-	-	-

Note. *$p < .01$.

The hierarchical regression showed that adding loss frequency and magnitude sensitivity did not account for significantly more variance in discounting than condition order alone (Table 3.17). Both Models 1 ($F(1, 90) = 7.02$, $p = .010$) and 2 ($F(3, 88) = 3.14$, $p = .029$) were significantly better at predicting discounting than the mean. Model 1 showed that condition order was a significant predictor, accounting for 6.20% of the variation in log(*h*) parameter for losses. In Model 2, adding loss frequency and loss magnitude sensitivities did not explain significantly more variance in log(h) than Model 1.

Table 3.17

Results of a two-stage hierarchical regression analysis using log(h) for losses

					Change statistics			
	B (SE)	β	t	p	F	df	p	adj. R^2
Model 1: Order alone					7.02	1, 90	.010	.06
Order	-0.20 (0.08)	-0.27	-2.65	.010				
Model 2: Order and sensitivity					1.18	2, 88	.311	.07
Order	-0.21 (0.08)	-0.28	-2.75	.007				
LF	-0.07 (0.05)	-0.14	-1.38	.170				
LM	0.03 (0.05)	0.07	0.65	.520				

Gains. As above, we conducted a two-stage hierarchical regression with log(h) for gains as the dependent variable and condition order entered into in the first block (Model 1) and the sensitivities to gain magnitude and gain frequency derived from the last experienced block entered into the second block (Model 2).

Table 3.18

Correlations between log(h) for gains, condition order, and sensitivities to gain frequency and magnitude

	Log(gains h)	Order	Gain frequency	Gain magnitude
Log(gains h)	-	0.06	-0.02	-0.04
Order	-	-	0.09	-0.19
Gain frequency	-	-	-	0.22*
Gain magnitude	-	-	-	-

Note. *$p < .05$

We observed no high correlations between independent variables (see Table 3.18), with the only significant correlation between gain frequency and magnitude, but the VIF and tolerance statistics showed no multicollinearity. Four cases were identified as outliers based on their standardized residuals being outside the -2 to 2 range (Field, 2009), but they were retained in the sample as the Cook's distance was less than 1 (Cook & Weisberg, 1982), Mahalanobis distance were less than 15 (Barnett & Lewis, 1978; Field, 2009) and the centred leverage values were less than twice the size of the average leverage (Hoaglin & Welsch, 1978). One additional case was identified by manual examination that exceeded both the

Mahalanobis distance and leverage cut-offs, but with Cook's distance less than 1; we excluded this participant, re-ran the regression and it produced no change in the results. We chose to retain this case. Analysis of residual and scatter plots showed that the data were normally distributed and met the assumptions of linearity and homoscedasticity.

The hierarchical regression showed that neither order alone ($F(1, 92) = 0.37$, $p = .544$) nor the inclusion of sensitivities ($F(3, 90) = 0.16$, $p = .923$) were significantly better at predicting discounting than the mean. Table 3.19 shows that none of the predictors significantly explained variance in the discounting of gains.

Table 3.19

Results of a two-stage hierarchical regression analysis using log(h) for gains

					Change statistics			
	B (SE)	β	*t*	*p*	*F*	*df*	*p*	adj. R^2
Model 1: Order alone					0.37	1, 92	.544	-0.01
Order	0.04 (0.07)	0.06	0.61	.544				
Model 2: Order and sensitivity					0.06	2, 90	.942	-0.03
Order	0.04 (0.07)	0.06	0.56	.577				
LF	-0.01 (0.03)	-0.02	-0.19	.849				
LM	-0.01 (0.04)	-0.03	-0.23	.817				

Discussion

We examined gain-loss asymmetry in a rapid-acquisition concurrent schedules task at both extended and local level of analysis. Our replication of Bull et al.'s (2015) Auckland Card Task procedure produced reasonable data quality relative to other such tasks with humans, but overall of poorer quality than the original study. According to Prospect Theory and the description-experience gap literature (e.g. Hertwig & Erev, 2009), we expected an asymmetry at both extended and local level of analysis in both frequency and magnitude tasks. We identified a consistent asymmetry when the participants responded to changes in frequency at the extended level, as well as in the immediate, local response to occurrence of gains and losses. We found less consistent results in the magnitude task, where participants were not sensitive to changes in magnitude at the extended level, but did show differences in local level of analysis. The specific nature of this asymmetry and its relation to our predictions will be discussed below.

ACT as a rapid-acquisition task

The behaviour of most participants showed control by the changes in primary outcome delivery, with data that was moderately well described by the GML. Our participants responded at a high rate and no participant was observed to stop responding prior to component completion. Furthermore, the high rate of exclusive preference for the more reinforcing or less punishing sides suggested control by the hypothetical money in that task (Horne & Lowe, 1993). However, while the average R^2 from the last four blocks and the last experienced block methods were comparable to other rapid-acquisition tasks with human participants which have reported moderately good fits (Lie et al., 2009; Krageloh et al., 2010), they were generally lower than reported by Bull (2013) for their overall sensitivity measure (no R^2 were provided for the asymptotic sensitivities). Furthermore, a sub-sample (approximately 40% with R^2 below 0.50 using the last experienced block method) of participants showed little control by the differential reinforcement rates, with an inconsistent distribution of response relative to the distribution of reinforcers. Since our adaptation of the ACT had several procedural changes from the original, we considered whether these could explain the reduced data quality.

Firstly, we considered differences that could have affected motivation to engage with the task. In our data collection, our participants were given course credit. In contrast, Bull et al. (2015) paid their participants a total of $30 for two 1-2 hour sessions, which the participants were told was dependent on their performance in the task. We also tested them in groups of four to fifteen, while Bull et al.'s participants completed the task alone. While our testing conditions were identical to Study 1, the more repetitive and longer ACT may have decreased engagement in the absence of perceived rewards for accurate performance and increased distraction due to the presence of other participants. Bowen and Kensinger (2017) observed that in tasks where participants were rewarded for correct recall of high vs. low reward stimuli, behaviour differed in the group responding for course credit as compared to the group responding for money. While the rate of correct recall of low and high reward stimuli did not differ for the credit group, the cash group showed higher rate for the high as compared to low reward stimuli. When data was collapsed across high and low reward stimuli, the performance of the credit group was not significantly different from the performance of the money group. This suggests that while overall engagement with the task goals did not depend on compensation type, the monetary rewards that the participants understood to be contingent on performance in Bull et al. may have encouraged increased discrimination of differential reinforcement rates.

Second, the changes we made to shorten the ACT for inclusion into two one-hour sessions could have contributed to the higher number of incomplete components and therefore limited information on the true likelihood of outcomes. Bull et al. (2015) provided their participants with a practice gain magnitude task that lasted for two components, with ten gains and no losses each, prior to data collection sessions. They also coded each component of the ACT to terminate after a maximum of eight minutes, while we shortened ours to six. On average, our participants who did not experience all of their primary outcomes, had their task terminated by timeout after 10-15 primary outcomes. The loss conditions were at the lower end of this range, which corresponds to the third and fourth block. This was somewhat earlier than Bull et al., who observed that the participants started to develop exclusive preference for one alternative by the fourth (13-16) and fifth (17-20) block, thus suspending the delivery of further primary outcomes until component time-out. The tendency of our participants to develop exclusive preference earlier in the task suggested that providing them with more time may not have resulted in a significant increase in component completion rates.

Overall, despite the features of a rapid-acquisition task that make it more suitable to human participants and therefore likely to increase engagement with the task goals, we did not obtain model fits in the high range reported by Bull et al. (2015). Our range of R^2 was comparable to Lie et al. (2009), whose participants were not paid and participated for course credit in a relatively short task, and to Krageloh et al. (2010), who paid their participants based on the accuracy of their performance and tested them over approximately 5 hours each. The similarity in our moderate R^2 values suggested that the differences discussed above, such increasing task duration and paying the participants, may not suffice in improving data quality, although we will return to the issue of ways to increase discrimination of reinforcement rates in the general discussion.

Gain-loss asymmetry in response to changes in frequency

When responding to changes in gain frequency, participants allocated their responses to the more frequently rewarding side at an increasing rate relative to the less frequently rewarding side as they received more reinforcers. Consistent with performance on other rapid acquisition tasks, participants became more sensitive to the differences between the two sides with more gains received in each component (Davison & Baum, 2000; Lie et al., 2009; Krageloh et al., 2010; Bull et al., 2015). Unlike Bull et al., experiencing more losses in each component did not result in increasing allocation of responses to the side with higher chance of no-loss as compared to the side with the lower chance of no-loss. Participants' sensitivity to changes in loss frequency remained low, on average, and did not improve with more losses experienced.

Increasing sensitivity for gain frequency and relatively unchanging sensitivity for loss frequency resulted in a gain-loss asymmetry that emerged with increasing experience of each component. Using responses at the last experienced block of primary outcomes, participants had higher sensitivity to changes in gain frequency than loss frequency. This was in contrast to using responses across the last four blocks of primary outcomes, which showed no significant differences in sensitivity. Thus an asymmetry emerged at later as opposed to earlier blocks, which was opposite to Bull (2013), where participants' sensitivity to gains and losses differed when using responses across all of the blocks (overall) and not when using averaged sensitivity to the last few blocks (asymptotic). In Bull, this was largely due to a greater change in sensitivity to gains across blocks, while sensitivity to losses was high at the first few blocks and did not greatly change across blocks. This change in gain sensitivity and relatively unchanging loss sensitivity corresponds to our data, although our sensitivity to losses began and ended low for most participants.

Participants showed greater sensitivity to changes in the frequency of gains rather than losses, consistent with the predictions of a reversed reflection effect in experiential choice (e.g. Hertwig & Erev, 2009), but overmatching was only observed in the gain frequency task. Using the last experienced block method, participants, on average, overmatched in the allocation of behaviour in response to changes in gain frequency and undermatched in response to changes in loss frequency. Overmatching for gains indicated that across the four components, the side with high probability of a gain elicited more responses and the side with low probability of a gain elicited fewer responses than perfect matching would predict. For losses, since we reversed our primary outcome ratio such that the side with fewer losses (higher chance of no-loss) ought to have more responses, undermatching indicated that the side with a low chance of no-loss elicited more responses and the side with a high chance of no-loss elicited fewer responses than perfect matching would predict. Participants were sensitive to which side was less punished, since the average sensitivity was above zero, but not to the same extent as determining which side was more reinforced on the gain task. Notably, when we limited our sample to participants with higher R^2 and better fit of the GML to data, we replicated overmatching for both gain and loss frequency tasks observed in Bull (2013), consistent with the predictions of the probability weighting function in experiential choice. Participants whose rate of response allocation showed a more consistent linear pattern were also showing a rate of allocation that was more extreme than the rate of change in primary outcomes.

Analysis of local effects showed consistency with the gain-loss asymmetry at the extended level. Our data supported McLean et al. (2014), showing that the increased preference for the just-productive side was partially a product of the general visit structure to each alternative. After any response and regardless of its outcome, participants were more likely to respond on the just-productive side than the alternative. Examining the difference between responses after real and hypothetical outcomes showed unique effects of real outcomes, supporting Gomes-Ng et al.'s (2017) re-analysis of existing data for corrected preference pulses. The occurrence of both gains and losses disrupted the general visit structure, with the initial response showing a decreased tendency to respond on the just-productive side (more on that below) and subsequently the likelihood of responding on the just-productive side increased linearly. However, the steeper rate of increasing preference for the just-productive side after gains meant that gains had a more lasting effect on responding, while for losses behaviour shifted sooner to the general visit structure.

Analysis of local effects showed additional properties; participants were more likely to switch sides and had shorter stays on each side in the context of frequently received losses as compared to gains. The generally less extreme preferences for the just-productive side in their general visit structure on the loss frequency task corresponded to lower sensitivity at the extended level. This reduced sensitivity and the general propensity to switch sides was consistent with the local analysis in Lie (2010), which identified that participants tended to favour the not just-productive alternative after receiving a loss, regardless whether this alternative was overall more or less punished, while for gains they were less likely to switch if the just-productive alternative had overall a higher reinforcement rate. Overall, our data suggested that a context with frequently delivered losses reduces the effect of a COD on the rate of switching.

Gain-loss asymmetry in response to changes in magnitude

Analysis of responses to changes in magnitude at the extended and local level replicated several patterns from the frequency analysis above. Participants' sensitivity increased across blocks with more gains experienced within each component and did not significantly change with more losses experienced. On average, participants undermatched in their allocation of responses, while limiting the sample to participants with better GML fits showed overmatching in both the gain and loss conditions. Local effects not differentiated by magnitude also replicated the overall pattern of disrupted general visit structure, with the initial response showing a decreased tendency to respond on the just-productive and subsequent responses showing an increasing tendency to respond on the just-productive side.

However, the differences in both extended and local effects between gains and losses were of a lesser extent than seen with frequency tasks. The average sensitivity to changes in magnitude was, in most estimates, lower than the corresponding sensitivity to frequency, consistent with the observation that human participants tend to be better at detecting differences in frequency rather than magnitude (e.g. Schmitt, 1974). No significant differences were detected between the gain and loss magnitude sensitivities when using responses to the last few outcomes or responses to the last four blocks. Limiting the data to a sub-sample with better GML fit similarly did not show a gain-loss asymmetry. This was consistent with our analysis of gain-loss asymmetry in Bull (2013), which also did not show an asymmetry using either of the estimates of sensitivity. Furthermore, in local effects analysis that did not differentiate by magnitude, the rate at which preference for the just-productive side increased after an outcome did not differ for gains and losses. Responses after a gain showing an overall greater likelihood to respond on the just-productive side across the thirty responses, than after a loss. General visit structure on the two tasks also did not differ, with a similar rate of switching, unlike the greater rate of switching on the loss as compared to gain frequency tasks.

Local analysis that differentiated responses by magnitude of the outcome that preceded them showed an asymmetry inconsistent with Kubanek et al. (2015), but consistent with our predictions based on the value function of Prospect Theory. Consistent with the low sensitivity to magnitude at the extended level, responses after smaller and larger gains showed a small effect of magnitude and a similar rate at which preference for the just-productive side increased for both magnitudes. The pattern here resembled the pattern seen in the local effects analysis not differentiated by magnitude, confirming that participants did not respond strongly to differences in gain magnitude. In contrast, responses after smaller and larger losses showed a larger effect of magnitude on the tendency to respond on the just-productive side. Participants were more likely to switch after a larger loss than a smaller loss. However, this did not translate to higher sensitivity to changes in loss magnitude at the extended level, where the sensitivities between gains and losses were not significantly different. Furthermore, after small magnitude gains and losses, responses were virtually identical, while after larger gains participants showed a steeper rate of increasing preference for the just productive side than after larger losses.

Corrected preference pulses

In our analysis of local effects of gains and losses, we noted that the immediate response to any primary outcome was to respond less on the just-productive side relative to the general visit structure. In the case of gains, preference for the just-productive side was lowest right after a gain and increased gradually. This appears to be contrary to the typical pattern of preference pulses in the literature, which was described as immediate heightened responses on the just-productive side which gradually decreases in intensity (Davison & Baum, 2002). Importantly, this pattern was not absent in our data when we consider the general visit structure and uncorrected responses after a primary outcome separately. After any given response, productive or not, we observed preference pulses as described in literature, which corresponds to McLean et al.'s (2014) observation that the general visit structures on concurrent schedule tasks tend to resemble preference pulses with or without occurrence of reinforcers. Our analysis showed that the unique effect of primary outcomes, when separated from the general visit structure, was to immediately reduce preference for the just-productive side (but not reverse it to the not just-productive side), which subsequently gradually increased.

Gomes-Ng et al. (2017) re-analysed select data derived from animal subjects by applying McLean et al.'s (2014) preference pulse correction and found the typical pattern of heightened preference for the just-productive side that gradually decreased. However, a similar pattern to ours was observed in conditions 2 and 4 of Phase 1 in Gomes-Ng et al. (2018), where the immediate response of animal subjects to a scheduled outcome showed reduced preference for the just-productive side compared to general visit structure (S. Gomes-Ng, personal communication, September 17, 2018). We have printed the relevant data below in Figure 3.15. Gomez-Ng et al. (2018) plotted the distribution of responses as a function of time since reinforcer, while we plotted the distribution of responses as function of ordinal response count since last outcome. Given the high rate of responding on our task and our focus on the first thirty responses, our data would roughly correspond to the distribution of responses in the first few seconds of Figure 3.15. It appears that in both our data, the occurrence of an outcome disrupted the participants' ongoing pattern of responding, prompting a brief increase in responses on the not just-productive side. The more non-productive time or responses that passed after the last productive response, the more likely the participants were to allocate more of their responses to the just-productive side. In Gomez-Ng et al., given the leaner schedule of reinforcement and longer sessions, they were able to observe behaviour returning to the general visit structure, which is likely what occurred in our data as well.

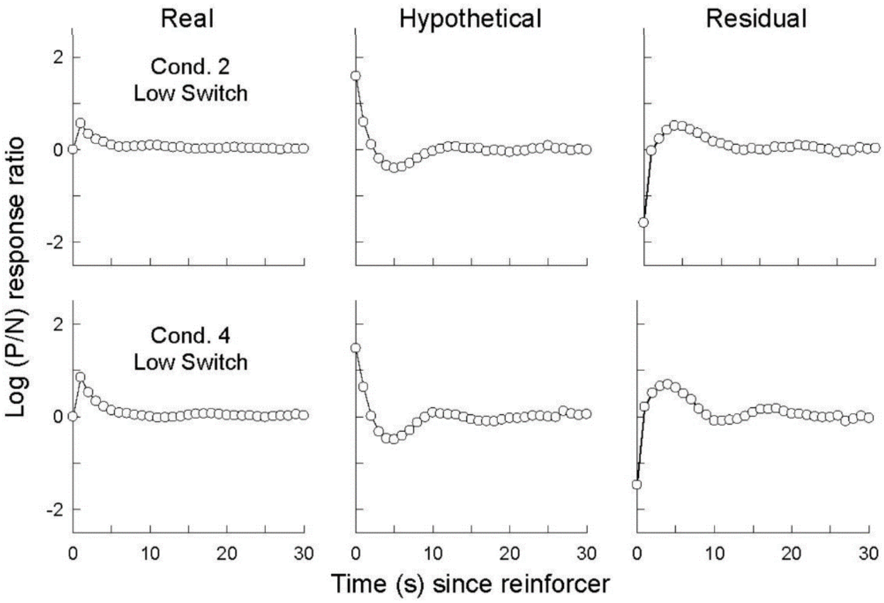

Figure 3.15. Logged ratio of responses on the just-productive side (P) over the not just-productive side (N) as a function of increasing time since last reinforcer. Responses are plotted after real outcomes (left), hypothetical outcomes (middle) and the difference between the two (right). Data from conditions 2 and 4 of Phase 1 in Gomes-Ng et al. (2018; S. Gomes-Ng, personal communication, September 17, 2018).

Our findings are preliminary given the limited literature that has applied McLean et al.'s (2014) correction to preference pulses (Hachiga, Sakagami, & Silberberg, 2015; Gomes-Ng et al., 2017; Gomes-Ng et al., 2018) and no application to human data to our knowledge. It does suggest that unique effects of outcomes might be to disrupt responding relative to the general visit-structure that is opposite to the pattern shown by uncorrected preference pulses. The occurrence of an outcome appeared to signal that an immediate occurrence of another was less likely, which was true given our dependent scheduling procedure. Upon receiving the outcome, the punishing effects on switching imposed by the COD appeared to lessen since another outcome on the just-productive side was also unlikely. Hence, at this point in the task, switching became slightly more preferable than in the general visit structure, but the participants overall allocated more of their responses to the just-productive side.

Direct-suppression GML

Using a net ratio of primary and secondary outcomes did not predict distribution of responses better than using the ratio of primary outcomes only. The better fit using primary outcome ratio suggested that the participants attended more to the difference in primary outcomes rather than tracked overall how much they won/lost on each side. Furthermore, in the case of loss frequency, the fit using net ratios was significantly worse. Comparison of primary outcome ratios vs. net ratios showed that on loss tasks, there was a wider range of obtained ratios than on the gain tasks, most likely due to the higher rate of incomplete components. On both of the loss conditions participants developed exclusive preference for one of the alternatives earlier in each component as compared to the gain conditions, thus suspending the delivery of further losses and resulting in incomplete blocks. In cases where participants did not receive all of the scheduled outcomes, the inclusion of a few secondary outcomes (loss of $30) in the gain conditions would not skew the ratio as much as the inclusion of a few secondary outcomes (gain of $170) in the loss conditions. The large difference in the magnitude of the secondary outcomes served to standardize the primary outcome and net ratios across conditions, but resulted in a greater range of outcome ratios for participants in the loss condition who did not experience all of the scheduled outcomes.

Discounting, sensitivity and gain-loss asymmetry

Contrary to Study 1, participants' discounting showed a reduced interaction between order and outcome type and a stronger effect of condition. Although analysis from fitting the hyperbolic model showed a significant interaction as in Study 1, the effect size was lower, while analysis using AUC showed a non-significant interaction. In contrast, participants discounted losses significantly more steeply than gains using both the h parameter and AUC, consistent with the predictions of the description-experience gap literature on the reversed reflection effect (e.g. Hertwig & Erev, 2009). Given that the fit of the hyperbolic model was good, it is likely that the AUC discrepancies were capturing noise in the distribution of indifference points and the results based on $\log(h)$ are more reliable. Nevertheless, the reduced effect size of the interaction and increased effect size of the main effect of condition both indicated a novel effect in the data.

The task differences in Study 2 that could have contributed to this stronger effect of condition included resetting the balance at the start of each condition, separating conditions by a day and the experience of the ACT prior to the discounting task. The possible effect of resetting the balance was discussed in Study 1 where we noted that magnitude effects in literature do not appear to be consistent with the direction of the order effects we observed. While we observed an order effect in Study 2, it was not to the same extent as in Study 1, which could have been caused removing the variation in the starting balance that altered the subjective value of the options. A similar effect could have been caused by separating the conditions by a day, which was made in order to accommodate both ACT and discounting task into two one-hour sessions. Any carry-over effects of discounting gains or losses, or variation in balance, on the subjective value would be reduced by separating the conditions by a long delay. This also corresponds to a lack of a significant correlation in Study 2 sample, where an individual's tendency to choose the larger, uncertain option did not vary to the same extent across the gain and loss conditions. Literature that has assessed discounting of real rewards has shown some carry-over effects over sessions up to one week apart (Matusiewicz et al., 2013), so while the day separation is also a possible cause of reduced carry over effects, it was not likely to be sufficient.

Lastly, sensitivity to frequency or magnitude in the ACT was not predictive of discounting after accounting for the effects of order. The greater sensitivity to gains than losses and the steeper discounting of losses than gains both supported the reversed reflection effect predictions of the description-experience gap, but variation in the former did not predict variation in the latter. Although no significant relationship was observed at the extended level, experience of chance outcomes on the ACT, regardless of eventual sensitivity to them, could have contributed to an increased experience with probabilistic events that was lacking in our Study 1, Experiment 3. The prolonged experience where each response had a chance of producing an outcome, and did not do so after most responses, could have comprised sufficient experience of chance outcomes to transition behaviour on the discounting task to that of decision from experience sooner than was observed in Study 1.

To summarize, we noted that the theoretical predictions of the description experience gap literature appeared to be at odds with some of the observed human behaviour on the concurrent schedules tasks. Given the lack of direct synthesis between the two fields of research, our data offers novel findings on the properties of a gain-loss asymmetry in human operant choice. Overmatching that was predicted by the probability weighting function was observed only in a minority of participants, and the majority showed undermatching typical in human subjects. While we did observe the predicted greater sensitivity to gain than loss frequency, sensitivity to magnitude at the extended level did not show a gain-loss asymmetry. Local level of analysis revealed additional properties of behaviour in the task, confirming the asymmetry at the extended level in the frequency tasks, but also showing an asymmetry in the magnitude tasks. Participants behaved according to the predictions of the value function, with a greater effect of changes in loss magnitude rather than gain magnitude on behaviour. Furthermore, all of our local level of analysis showed a novel pattern in corrected preference pulses, supporting the importance of applying a general visit structure correction to any future study of preference pulses. Lastly, we also observed a possible effect of ACT on subsequent probability discounting, which will be examined in Study 3.

Study 3: Investigation into the novel discounting pattern of Study 2

In Study 3, we addressed follow-up questions raised in Studies 1 and 2. Our procedure in Study 2 produced a novel pattern of results in the discounting task. Participants showed a weaker effect of order and a stronger effect of condition that was consistent with the reversed reflection effect described in the description-experience gap literature (Hertwig & Erev, 2009). In Study 3, we examined several procedural changes in Experiments 1 and 2 that could have contributed to the gain-loss asymmetry observed in Study 2. One possible contributor was that participants experienced chance outcomes in the ACT that preceded the discounting conditions. Another was that we set the balance to $3000 at the start of each condition, and a third was that we split the conditions over two test days. We considered that the most likely contributing factor for the increased gain-loss asymmetry in Study 2 was that participants had recently completed the ACT, but the impact of the other two factors on the more consistent steeper discounting of losses could not be separated out in Study 2.

In Study 3, Experiment 1 we ran an experiential safe-risky probability discounting task with reset balance and no delay between tasks. This was aimed to test whether resetting the balance was sufficient to show the significant effect of condition seen in Study 2. In Study 3, Experiment 2 we reset the balance in the same way as we did in Study 3, Experiment 1, but participants completed a short version of the ACT (consisting of gain and loss frequency tasks) before the discounting tasks, and there was no delay between the discounting conditions. This was to determine whether the experience of the ACT was necessary to show a significant effect of condition.

Experiment 1: Effect of balance reset on experiential discounting

Study 3, Experiment 1 aimed to determine whether resetting the balance before each condition was sufficient to produce the steeper discounting of losses than gains seen in Study 2. Participants made choices about gains and losses with a starting balance of $3000 for both conditions. If data resemble the patterns in the discounting tasks in Study 2, with a weaker interaction effect and a stronger main effect of condition, then resetting the balance was sufficient to produce the novel pattern. If data resemble discounting patterns in Study 1, with a significant interaction effect and a non-significant main effect of condition, then resetting the balance was not sufficient to produce the novel pattern. Furthermore, if the order effect persists with resetting of the balance, we can also rule out the effect of changing balance magnitude as a cause of order effects discussed in Study 1.

Method

Participants

Ninety-nine participants attending Victoria University of Wellington were recruited through the School of Psychology Research Programme online tool and participated in partial fulfilment of a course requirement.

Materials

The experiential safe-risky money task was as described in Study 1, Experiment 2 and included the reset balance as implemented in Study 2.

Procedure

The participants were tested as described in Study 1, Experiment 2.

Results

Data were analysed as described in Study 1, Experiment 1. Table 4.1 specifies the number of participants excluded based on the Johnson and Bickel (2008) criteria. The analysis below uses participants with fully systematic data. Individual *h* parameters were not normally distributed (Shapiro-Wilk's test for normality all $p < .001$), with high positive skew and kurtosis for both gains and losses. Data were log transformed, with Shapiro-Wilk's test for normality remaining significant for gains ($W = 0.92$, $p < .001$) only. Examination of histograms showed no skew and slight positive kurtosis and the shape of the distribution was otherwise normal. Individual AUC parameters showed significant Shapiro-Wilk's tests for gains ($W = 0.97$, $p = .034$) and losses ($W = 0.95$, $p = .001$), but the examination of histograms showed slight positive kurtosis for gains and slight positive skew for losses, but the shape was otherwise normal.

Table 4.1

Number of participants (n = 99) in Study 3, Experiment 1 who had unsystematic data by criterion, and total participants with systematic data

	Gain	Loss
Criterion 1	3	3
Criterion 2	0	2
Total systematic	93 (93.94)	

Note. Percentages in parenthesis.

The hyperbolic model provided good fits to group median indifference points and somewhat lower median R^2 derived from individual participants compared to Study 1, Experiment 3 and Study 2 (see Table 4.2). Furthermore, the direction of steeper discounting was inconsistent between *h* derived from group median indifference points and from individual participants, most likely due to somewhat poorer fit. Examination of residuals showed no systematic trends (see Appendix G). Although there was not a systematic trend in the residuals, the poorer fits reduce our confidence in the *h*-based analysis and we examined AUC. AUC from group median indifference points and individual participants were consistent in their direction of slightly steeper discounting of losses than gains, but overall were roughly equivalent between gains and losses.

Table 4.2

Study 3, Experiment 1 money task h, R^2 and AUC values

	Group median indifference points			Individual participants		
Condition	*h*	R^2	AUC	*h* median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Gain	1.82	0.97	0.25	1.46 (1.07, 2.66)	0.79 (0.82, 0.94)	0.27 (0.01)
Loss	1.46	0.96	0.24	1.73 (1.02, 2.83)	0.77 (0.77, 0.94)	0.25 (0.01)

Note. Q_1 = first quartile, Q_3 = third quartile.

As in Study 1, discounting was affected by an interaction of condition and order, and did not show a main effect of condition seen in Study 2 (see top right section of Figure 4.1). A mixed measures ANOVA with order as the between-subjects factor and condition as the within-subjects factor showed no significant main effect of condition ($F(1, 91) < 0.001$, $p = .992$, $\eta_p^2 < .001$), order ($F(1, 91) = 0.02$, $p = .889$, $\eta_p^2 < .001$), and a significant interaction ($F(1, 91) = 23.80$, $p < .001$, $\eta_p^2 = .207$). The inset in the left panel of Figure 4.1 shows this interaction of order and condition; when losses were presented first, gains ($M = 0.12$, $SE = 0.05$) were discounted significantly less steeply than losses ($M = 0.36$, $SE = 0.05$; $p = .001$), and when gains were presented first, gains ($M = 0.36$, $SE = 0.05$) were discounted significantly more steeply than losses ($M = 0.12$, $SE = 0.05$, $p = .001$; Bonferroni correction applied). The log(*h*) parameters were significantly higher than zero for gains in the gains first order ($t(46) = 6.67$, $p < .001$, $d_z = 0.97$), losses in the gains first order ($t(46) = 3.10$, $p = .003$, $d_z = 0.45$), gains in the losses first order ($t(45) = 2.86$, $p = .006$, $d_z = 0.42$), and losses in the losses first order ($t(45) = 7.12$, $p < .001$, $d_z = 1.24$). Again, given the poorer fits to the hyperbolic model, the AUCs were further analysed.

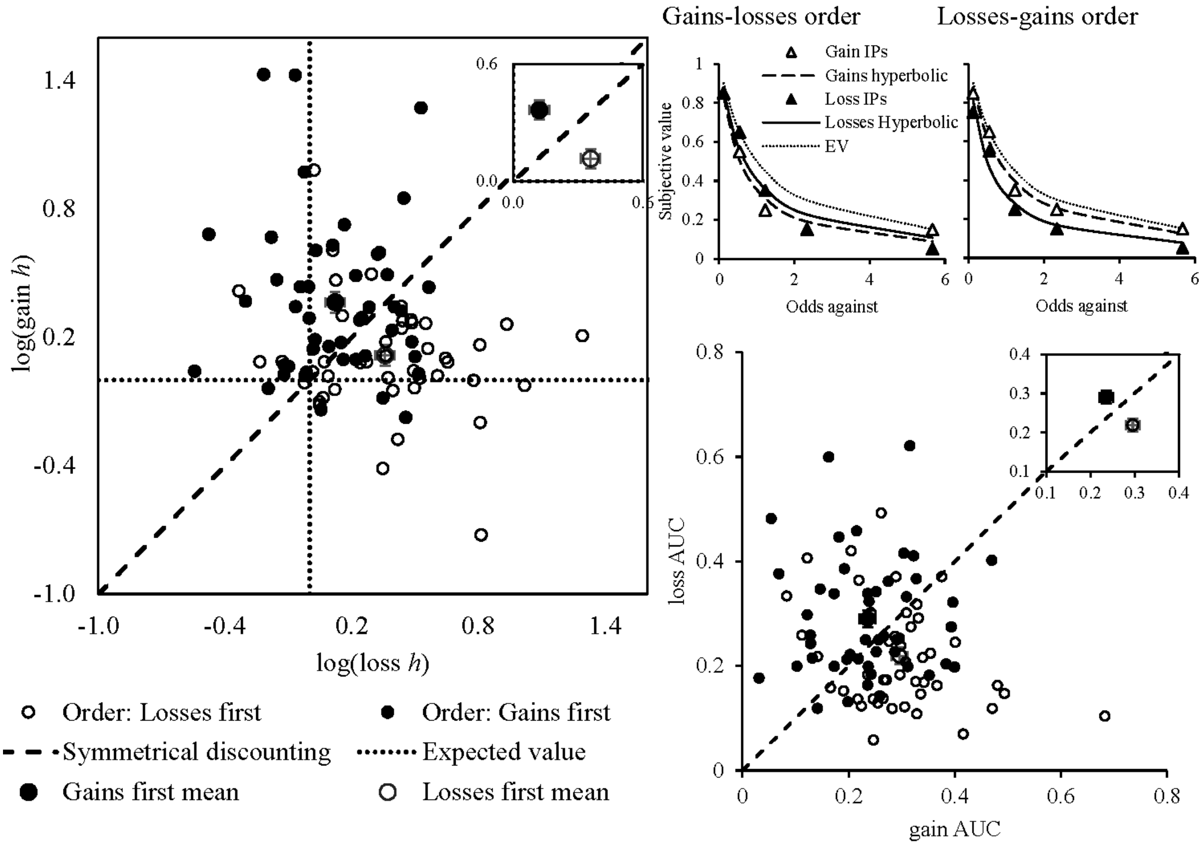

Figure 4.1. Left panel: Logged *h* parameters for gains plotted against logged *h* parameters for losses for each individual, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and white circles represent the means for each order, which are also shown in the inset. The error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged *h* value (0) when decisions are made based on expected value. Top right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss (black triangles). For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Bottom right panel: AUC for losses plotted against AUC for gains for each individual. Order split and group means are plotted as described for the left panel graph.

Analysis using AUC confirmed the log(h) findings, with discounting affected by the interaction of order and condition only. Mixed measures ANOVA showed a non-significant main effect of condition (F(1, 91) = 0.55, p = .462, η_p^2 = .006), order (F(1, 91) = 0.14, p = .710, η_p^2 = .002) and a significant interaction (F(1, 91) = 16.03, p < .001, η_p^2 = .150). When losses were presented first, gains (*M* = 0.30, *SE* = 0.02) were discounted significantly less steeply than losses (*M* = 0.22, *SE* = 0.02; *p* = .001), and when gains were presented first, gains (*M* = 0.24, *SE* = 0.02) were discounted significantly more steeply than losses (*M* = 0.29, *SE* = 0.02, *p* = .022; Bonferroni correction applied).

As in Study 1 and unlike Study 2, gains and losses were significantly negatively correlated when using log(*h*) (*r*(91) = -.25, *p* = .016) and AUC (*r*(91) = -.21, *p* = .041). Overall, data supported the conclusion that resetting balance was not sufficient to produce the reduced effect of order observed in Study 2, making the influence of the ACT and/or separating the conditions the more probable causes.

Experiment 2: Effect of recent experience of probabilistic outcomes on experiential discounting

Given that resetting the balance alone did not replicate the novel pattern from Study 2, in Study 3, Experiment 2 we aimed to determine whether experience of probabilistic outcomes in the ACT was a necessary component. Participants completed a short version of the ACT which included gain and loss frequency conditions only, after which they discounted gains and losses with a starting balance of $3000 for both conditions. If data resemble discounting patterns in Study 2, with a weaker interaction effect and a stronger main effect of condition, then experience with the ACT was a necessary component. If data resemble discounting patterns in Study 1, with a significant interaction effect and a non-significant main effect of condition, then neither the ACT nor the reset balance are necessary and the novel result in Study 2 was likely due to a day separation.

Method

Participants

Ninety participants attending Victoria University of Wellington were recruited through the School of Psychology Research Programme online tool and participated in partial fulfilment of a course requirement.

Materials

Shortened Auckland Card Task. We programmed the shortened version of ACT using Microsoft Visual Basic largely as described in Study 2. Participants completed gain and loss frequency conditions only and the arranged schedules are described in Table 4.3. Each condition was shortened to include 3 components and the component length was further shortened to last a maximum of five minutes. The 25:65 component was replaced with a 50:50 condition. The three components per task were presented in random order.

Table 4.3

Summary table of the two conditions in the shortened Auckland Card Task adapted from Bull et al. (2015)

Condition	Component	Gains (Probability/Amount/Number of outcomes available)		Losses (Probability/Amount/Number of outcomes available)		Net reward		Primary outcome ratio	Net ratio
		Deck 1	Deck 2	Deck 1	Deck 2	Deck 1	Deck 2		
Gain frequency	1	.25/$50/5	.75/$50/15	.50/$30/5	.50/$30/5	$100	$600	1:3	1:6
	2	.50/$50/10	.50/$50/10	.50/$30/5	.50/$30/5	$350	$350	1:1	1:1
	3	.75/$50/15	.25/$50/5	.50/$30/5	.50/$30/5	$600	$100	3:1	6:1
Loss frequency	1	.50/$170/5	.50/$170/5	.75/$50/15	.25/$50/5	$100	$600	1:3	1:6
	2	.50/$170/5	.50/$170/5	.50/$50/10	.50/$50/10	$350	$350	1:1	1:1
	3	.50/$170/5	.50/$170/5	.25/$50/5	.75/$50/15	$600	$100	3:1	6:1

Note. Primary outcome ratios for the loss condition are reversed, such that the side with more losses predicts fewer responses.

Experiential Money Task. Experiential safe-risky money task was as described in Study 1, Experiment 2 and included reset balance implemented in Study 2.

Procedure

The participants were tested as described in Study 1, Experiment 2. The order of the ACT conditions was counterbalanced, but always occurred such that the ACT conditions preceded the discounting condition. The order of discounting conditions was also counterbalanced.

Results

Discounting data were analysed as described in Study 1, Experiment 1. Table 4.4 specifies the number of participants excluded based on the Johnson and Bickel criteria (2008). Analysis below uses participants who had fully systematic data across conditions. Individual h parameters were not normally distributed (Shapiro-Wilk's test for normality all $p < .001$), with high positive skew and kurtosis for both gains and losses. Data were log transformed, with Shapiro-Wilk's test for normality remaining significant for gains (W = 0.97, p = .048) and losses (W = 0.97, p = .029). Examination of histograms showed slight positive kurtosis for losses, but the shape of the distribution was otherwise normal. Individual AUC parameters showed significant Shapiro-Wilk's test for losses (W = 0.94, p = .001) with slight positive kurtosis but otherwise normal distribution, and a non-significant result for gains (W = 0.99, p = .492). As the focus of Study 3 was on isolating the source of the novel effect on discounting rates, we did not include an analysis of the shortened ACT here.

Table 4.4

Number of participants (n = 90) in Study 3, Experiment 2 who had unsystematic data by criterion, and total participants with systematic data.

	Gain	Loss
Criterion 1	0	1
Criterion 2	1	2
Total systematic	87 (96.67)	

Note. Percentages in parenthesis.

The hyperbolic model provided good fits to group median indifference points and individual data (see Table 4.5), with steeper discounting of losses than gains as with the experiential money tasks in Study 2. Examination of residuals showed no systematic trends (see Appendix H).

Table 4.5

Study 3, Experiment 2 money task h, R^2 and AUC values

	Group median indifference points			Individual participants		
Condition	h	R^2	AUC	h median (Q_1, Q_3)	R^2 median (Q_1, Q_3)	AUC mean (SE)
Gain	1.24	0.98	0.31	1.35 (0.88, 2.11)	0.88 (0.72, 0.93)	0.31 (0.01)
Loss	1.83	0.96	0.23	1.70 (1.08, 2.62)	0.87 (0.75, 0.93)	0.26 (0.01)

Note. Q_1 = first quartile, Q_3 = third quartile.

Discounting was affected by condition order and did not show a significant main effect of condition as seen in Study 2 (see top right section of Figure 4.2). Mixed measures ANOVA with order as the between-subjects factor and condition as the within-subjects factor showed a non-significant main effect of condition (F(1, 85) = 3.25, p = .075, η_p^2 = .037), non-significant main effect of order (F(1, 85) = 0.56, p = .456, η_p^2 = .007), and a significant interaction ($F(1, 85) = 7.58$, $p = .007$, $\eta_p^2 = .082$). The inset in the left panel of Figure 4.2 shows this interaction of order and condition; when losses were presented first, gains (M = 0.10, SE = 0.05) were discounted significantly less steeply than losses (M = 0.34, SE = 0.05; $p = .002$), but when gains were presented first, there was no significant difference between discounting of gains and discounting of losses ($p = .501$; Bonferroni correction applied). The log(h) parameters were significantly higher than zero for gains in the gains first order ($t(43) = 5.36$, $p < .001$, $d_z = 0.81$), losses in the gains first order ($t(43) = 3.84$, $p < .001$, $d_z = 0.58$), losses in the losses first order ($t(42) = 6.54$, $p < .001$, $d_z = 1.00$), but not for gains in the losses first order ($t(42) = 1.66$, $p = .104$, $d_z = 0.25$).

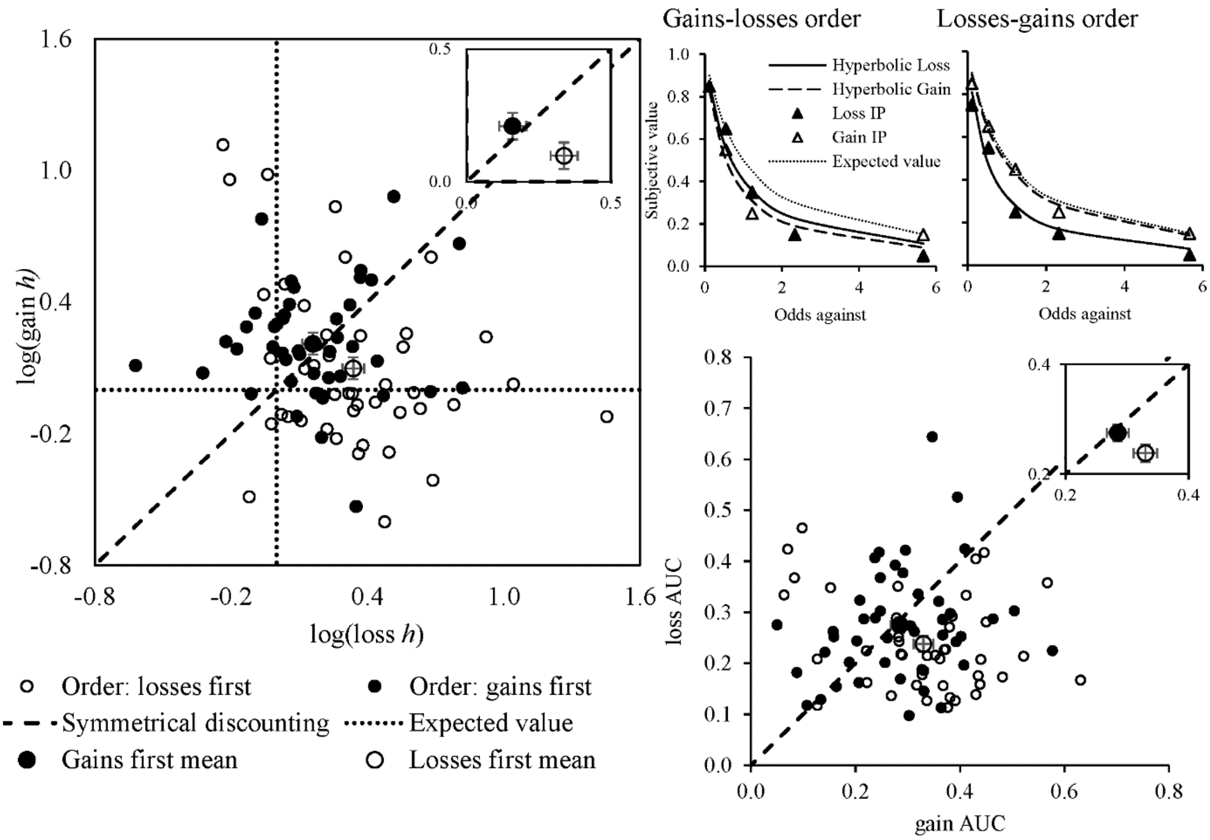

Figure 4.2. Left panel: Logged h parameters for gains plotted against logged h parameters for losses for each individual, split by order; black circles are gain-loss and white circles are loss-gain order. The larger black and white circles represent the means for each order, which are also shown in the inset. The error bars are standard error of the mean. The diagonal dashed line represents symmetrical discounting of gains and losses. Dotted vertical and horizontal lines demarcate logged h value (0) when decisions are made based on expected value. Top right panel: Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (white triangles) or loss

(black triangles). For each task, data are split by order as indicated by the graph titles. Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Bottom right panel: AUC for losses plotted against AUC for gains for each individual. Order split and group means are plotted as described for the left panel graph.

For consistency with the analysis of Studies 1 and 2, the analyses were repeated with AUC as the dependent measure despite the good fits of the hyperbolic model. While analysis using AUC confirmed the presence of a significant interaction, it also showed a significant effect of condition contrary to the log(h) results. A mixed measures ANOVA showed a significant main effect of condition (F(1, 85) = 8.91, p = .004, η_p^2 = .095), non-significant main effect of order (F(1, 85) = 0.06, p = .805, η_p^2 = .001) and a significant interaction (F(1, 85) = 5.72, p = .019, η_p^2 = .063). When losses were presented first, gains (M = 0.33, SE = 0.02) were discounted significantly less steeply than losses (M = 0.24, SE = 0.02; p < .001), and when gains were presented first, gains (M = 0.29, SE = 0.02) were discounted significantly more steeply than losses (M = 0.28, SE = 0.02, p = .674; Bonferroni correction applied).

Gains and losses were significantly negatively correlated when using log(h) (r(85) = -.22, p = .040) but not when using AUC (r(85) = -.05, p = .641). Overall, experience of ACT appears to contribute to the stronger effect of condition in Study 2.

Discussion

Participants' choices on the probability discounting task in Study 2 showed a stronger effect of condition consistent with the predictions of the reversed reflection effect and a weaker interaction between condition and order. We conducted two follow-up experiments in order to determine the likely contributing factors to this novel pattern. The combination of Study 3, Experiments 1 and 2 was to determine if experience of the ACT produced the gain-loss asymmetry via providing participants with more experience of probabilistic outcomes and moving behaviour closer towards decision from experience.

We observed no significant effect of condition and replicated the order effect of Study 1 when participants discounted with a reset balance in Study 3, Experiment 1. Controlling for the variation in balance at the start of each condition did not affect the direction or presence of the order effect, suggesting that this was not a contributing factor to the observed order effects in Study 1. Furthermore, the absence of a significant effect of condition indicated that this was not the likely cause of novel results in Study 2. Overall, Study 3, Experiment 1 replicated Study 1, Experiment 3 and showed that resetting the balance was not the likely cause of the gain-loss asymmetry as seen in Study 2.

Study 3, Experiment 2 showed that the inclusion of the experience of probabilistic outcomes, in the form of a shortened ACT, was sufficient to show an effect of condition, but not alone sufficient to produce a difference between discounting of gains and losses of the size and consistency observed in Study 2. Results based on log(h) showed a significant interaction between order and condition and no significant effect of condition, while analysis using AUC showed an additional significant effect of condition consistent with the reversed reflection effect (Hertwig & Erev, 2009). The hyperbolic model showed a good fit to the data, so unreliability of the derived h parameter as the source of this discrepancy seems unlikely. Comparison of effect sizes in Study 3 showed a reduced strength of the interaction effect in Experiment 2 relative to Experiment 1, while the effect size for the main effect of condition showed an increase from Experiment 1 to 2. Furthermore, the presence of significant negative correlation in the log(h), but not AUC data also suggested reduced carry-over effects of discounting in one condition on discounting in the other. On balance, the experience of ACT appeared to reduce the interaction effect and increase the effect of condition in a manner comparable to, but not of the same extent as seen in Study 2. Having the two conditions a day apart appears to have been an important additional feature that further reduced, but not eliminated the carry-over effects (see Matusiewicz et al., 2013 and Ohmura et al., 2006 for similar effects).

General Discussion

We assessed whether gain-loss asymmetry occurred in a series of experiential contexts and examined whether choice patterns accorded with the predictions of the reversed reflection effect in choice from experience. We synthesised the predictions of the reflection effect and the reversed reflection effect. We then tested the resulting predictions in two experiential choice procedures: probability discounting and concurrent schedules tasks. It was not previously known whether gain-loss asymmetry occurred in either experiential context. Specifically, in each experiential task we aimed to determine 1) whether there was a gain-loss asymmetry, and 2) whether this asymmetry was consistent with a reversed reflection effect. In both probability discounting and concurrent schedules tasks, we observed several consistencies with the predictions of the reversed reflection effect. We also observed several inconsistencies with these predictions, the probable causes of which were discussed in the relevant studies above. Here, we wish to focus on the features of a gain-loss asymmetry in choice from experience that have emerged across the two procedures, and, in the context of these findings, to end with a discussion of the challenges in human experiential research.

Gain-loss asymmetry in experiential choice

Throughout the experiential tasks used, participants showed an asymmetry in how they responded to changes in the probability of a gain or a loss. In the probability discounting tasks, when the participants chose between pairs of certain and risky outcomes that were described, experienced, and preceded by a recent exposure to other probabilistic outcomes, the rate at which a loss outcome lost its value with decreasing probability was steeper than that of a gain outcome for most participants. In the ACT, when the participants chose between probabilistic outcomes that were not described, but were experienced and choices were made after sampling the options, they were more sensitive to changes in the frequency of gains than the frequency of losses. Both of these asymmetries were consistent with predictions of Prospect Theory (Kahneman & Tversky, 1979), in that gains and losses were not valued symmetrically. Furthermore, the direction of this asymmetry was consistent with the reversed reflection effect of the description-experience gap literature (Hertwig & Erev, 2009).

However, despite the general consistency in the direction of a gain-loss asymmetry, when we compared participants’ choice patterns to choice patterns that would allow maximization, what we found was not consistent with our predictions. In probability discounting, choice according to the expected value of the outcomes would maximize gains and minimize losses with repeated choice. In concurrent schedules tasks fitted with the GML, response allocations that strictly match the distribution of gains and losses between the two alternatives maximize net profit by the end of the session. The S-shaped probability weighting function (Figure 1.3), that we used to model our predictions for choice from experience, predicted that participants would not strictly follow a pattern of responses that would maximize net profit. We expected underweighting of low probability and overweighing of high probability events, which was more extreme for decisions involving gains than losses. In probability discounting, we reasoned that according to this S-shaped function, discounting rates ought to be shallower than choice based on expected value (given that most of our chosen probabilities were in the high range), and in concurrent schedules, for response distributions to overmatch relative to perfect matching (given that participants ought to overweight high probabilities and underweight low probabilities).

In the discounting task, despite making the task experiential and adding exposure to probabilistic outcomes (via the ACT), discounting remained steeper than choice based on expected value and did not show the predicted pattern of shallower discounting. In the ACT, a similar pattern was observed with losses, where the majority of participants undermatched relative to ideal matching. Our prediction of overmatching was supported in only the gain frequency condition of the ACT.

One explanation for the lack of shallower discounting on the probability discounting task and of overmatching in the loss frequency condition of the ACT is that behaviour was not fully informed by experience of the outcomes. We return here to the issue of capturing behaviour in transition discussed in Study 1. We attributed the consistent effect of condition order on discounting to the task capturing choice in transition between being informed by description only to being informed by both experience and description. Experience on a task that delivered chance outcomes, the ACT, did not fully transition probability discounting rates to those expected when choice is informed by both experience and description in Studies 2 and 3 (evidenced by significant order effects). It might have partially produced this transition, however, as more participants showed steeper discounting of losses than gains after completing the ACT than in Study 1 when they had not completed the ACT. Furthermore, across all of our discounting experiments that were sufficiently powered to examine the effect of order, discounting shifted closer to choice based on expected value with each condition experienced. Average rates of discounting in the experiential tasks were never significantly shallower than choice based on expected value, but as Table 5.1 shows, when we compare mean log(h) to zero, the effect size reliably decreased from first to second condition across all of the experiments. The magnitude of the difference between a log(h) of zero, corresponding to choice based on expected value, and the average log(h) for each condition generally decreased with more experience of the task.

Table 5.1

Summary of effect sizes from one sample t tests in Studies 1-3 comparing discounting to that predicted by the expected value.

	Effect size (Cohen's d_z)			
	Study 1 E3	Study 2	Study 3 E1	Study 3 E2
1st condition experienced				
Gains	0.86*	0.77*	0.97*	0.81*
Losses	0.96*	1.32*	1.24*	1.00*
2nd condition experienced				
Gains	0.37*	0.51*	0.42*	0.25
Losses	0.54*	0.73*	0.45*	0.58*
1st condition experienced after 1 task repetition				
Gains	0.50*			
Losses	0.58*			
2nd condition experienced after 1 task repetition				
Gains	0.32			
Losses	0.60*			

Notes. * = significantly steeper than choice based on expected value.

We noted in our discussion of the order effects in probability discounting that more experience of probabilistic outcomes might shift behaviour closer to our predictions. Examining a gain-loss asymmetry on a concurrent schedules task was proposed to ameliorate this issue. Our data on sensitivity to changes in gain frequency supported this, as the more gains were experienced by the participant, the higher their sensitivity was to differences between the two alternatives. Most participants reached overmatching by the completion of each component. While Bull et al. (2015) observed the same effect with losses, where increasing experience of the outcomes resulted in overmatching by the end of each component, we did not, with average sensitivity showing undermatching. In our Study 2 introduction, we suggested that undermatching is consistent with the predictions of the Prospect Theory for choice from description, and we can further parallel this to generally steeper discounting than predicted by expected value. Consistent with the notion that the difference between our patterns and Bull et al.'s data was that our participants had less experience with the scheduled probabilistic outcomes, we observed that the rate of incomplete components in the ACT, where our participants did not experience all of the scheduled outcomes, was higher in the loss as opposed to the gain conditions. Thus, it is possible that participants did not have sufficient experience of probabilistic outcomes on the loss frequency task to result in the predicted overmatching. However, even among participants who reached the final block and had thus experienced most or all scheduled losses, sensitivity remained at undermatching, suggesting that differing experience of probabilistic outcomes was not the sole explanation.

This brings us to another consistency across the tasks: data from both the discounting tasks and the ACT demonstrated that more experience with probabilistic outcomes did not have an equivalent effect on choices about gains and losses. In the discounting tasks, the rate at which gains were discounted varied more with contextual changes, such as condition order and task repetition, while discounting of losses did not vary to the same extent. Furthermore, comparison of effect sizes in Table 5.1 for Study 1, Experiment 3 showed that repeating the discounting task twice did not result in any further shift in the group means towards choice based on the expected value, but the discounting of gains showed more variability. Those participants that completed their second gain condition as the final condition in the task showed the shallowest average rate of discounting. As we noted above, a similar pattern was observed in the ACT, where experience of more losses within each component did not result in increasing sensitivity to the underlying contingencies, unlike experience of more gains.

However, changes in sensitivity across blocks and changes in discounting with condition order and task repetition are related, but not necessarily equivalent measures of the effect of experience. In the extended level of analysis, sensitivity at each block is derived across the four components, and these components could have occurred in any order within the condition. That is, data at block 2 does not correspond to choice based on little experience of outcomes in the task, but rather to choice based on little experience of each component. A more relevant test might be to examine local effects of gains and losses early versus late in the task. We did not conduct such an analysis partially due to the relatively short task in which further subdivisions in data would lose reliability and partially due to the fact that we counterbalanced the order of the four conditions, as well as the four components within each condition. Future experiments could, for example, administer a single gain frequency task with a longer component duration that would allow for a comparison between corrected preference pulses at the start versus by the end of the task for the differential pattern seen at the extended level of analysis. This would also speak to whether the local effects that we observed constituted transient and unlearnt responses to gains and losses that did not change with task experience or whether they too, like sensitivity and discounting, were affected by task experience (Davison & Baum, 2000; 2002).

This relatively unchanging sensitivity to losses and the greater shift in behaviour in response to gains was not only consistent with our discounting data, but also with discounting literature that observed an effect of magnitude with discounting of gains, but not of losses (e.g. Estle et al., 2006; Green et al., 2014). However, our analysis of the effect of magnitude in the ACT was inconsistent with the effect of magnitude in the discounting literature. That is, a pattern in local effects consistent with the magnitude effect in discounting literature would be a greater impact of changing gain magnitude than loss magnitude on subsequent responses, which was not what we observed. Aggregate data showed no difference in sensitivity to gain and loss magnitude, and local analysis showed transient effects where changes in loss magnitude had a greater impact than changes in gain magnitude. More importantly, the pattern in local effects was consistent with the predictions of the value function of the Prospect Theory (Kahneman & Tversky, 1979), which predicted that a change in loss magnitude ought to have a greater impact on subjective value than an equivalent change in gain magnitude. Notably, this comparison of local effects and discounting remains speculative as discounting includes variation in both probability and amount, hence subject to the joint predictions of the value and probability weighting functions. A more direct test of whether the effect of magnitude on local effects directly parallels the effect of magnitude on discounting would be to conduct gain and loss frequency conditions which measure sensitivity at one magnitude (e.g. $50 as in ACT) and compare this to a second set of gain and loss frequency conditions which measure sensitivity at a larger magnitude (e.g. $5000). If there is a greater difference between the $50 and $5000 conditions for gains than losses then this would indicate a consistency of effect across the two task contexts.

Challenges in human experiential tasks

We have observed a gain-loss asymmetry in both of the tasks used, and while this asymmetry was present in the majority of the participants in each sample, a substantial minority remained that did not show the predicted pattern of behaviour. On the one hand, it echoes research that has observed sub-samples that do not conform to the predictions of the differences in the functional value of gains and losses (e.g. Hershey & Schoemaker, 1980; Schneider & Lopes, 1986; Experiment 2 of Critchfield et al., 2003). On the other hand, the extent to which gain-loss asymmetry deviates in choice from experience is highly dependent on procedure used to measure it (Wulff et al., 2018), and we have encountered several challenges in our experiential tasks that may have contributed to the inconsistencies in the observed gain-loss asymmetry. We focus on two such procedural challenges: determining sufficient experience of probabilistic outcomes and ensuring delivery of scheduled outcomes.

We noted that the changes in the rate of discounting with task repetition and after exposure to recent probabilistic outcomes suggested that we captured choice in transition from decision from description to decision from experience (Jessup et al., (2008); Lejarraga & Gonzalez, 2011; Wulff et al., 2018). Behaviour in this transition state is informative in itself, as it predicts how individuals would behave when starting with a full description of the outcomes and, subsequently, how increasing experience of the frequency of the described events would change their behaviour (Shafran, 2011). To return back to the example of taking daily medication, we might expect an individual's initial choice to conform to models that predict behaviour based on description only. Subsequently, with enough experience, choice would transition to being better described by models based on experience. During this transition, our data suggested that behaviour might shift more rapidly in response to increasing experience of gains, rather than losses. Furthermore, changes in the magnitude of a gain would not be expected to affect subsequent choice as much as changes in the magnitude of a loss. Note that we are not referring to the effects of framing the same health outcome as a gain or loss (e.g. Akl et al., 2011); here we are referring to how much we expect choices between gains or choices between losses to shift in response to context or extent of experience of these events.

An alternative is to view the pattern of behaviour we observed in the discounting task combining description and experience not as a transition from one to the other, but as behaviour that is expected in this context. In other words, a combination of description and experience ought to be treated as a separate paradigm. For one, such contexts are certainly analogous to many everyday situations that are not purely descriptive or purely experiential (e.g. taking daily medicine example; living in an earthquake-prone city where descriptions of earthquake risks are widely available). This combination has also received limited attention in literature, where the majority of research focuses on either choice purely from description or purely from experience (although see Shlomi, 2014; Fantino & Navarro, 2012; Lejarraga & Muller-Trede, 2016 for recent examples of investigations of choice from a combination of description and experience).

Barron, Leider and Stack (2008) presented two groups of participants with a choice between a smaller gain for certain ($0.10) and a risky prospect that had a high chance of a larger gain (0.999 probability of gaining $0.13) and small chance of a large loss (0.001 probability of losing $15). The "early" group made a 100 choices between these options that were both described and experienced from the start, akin to our discounting procedure. The "late" group made the first 50 choices without the description and with experience only, and the options were described at the start of the second 50 trials. Participants in the "early" group who relied on both description and experience had a lower proportion of risky choices than participants who relied on experience alone. The introduction of description to the experience only group resulted in a decrease in the proportion of risky choices, but this decrease was not sufficient for risk preference to converge with levels in the description and experience group. The "late" group remained more risk seeking than the "early" group even when both had complete and identical information from both description and experience. Barron et al. concluded that the experience of generally positive outcomes from choosing the risky option could not be fully overridden by a description of the possible negative outcomes after a period of exposure. Hence, and perhaps in contrast to Jessup et al. (2009) and Lejarraga and Gonzalez (2011), another factor that could explain why behaviour in our task did not quite resemble some of our predictions is that the combination of description and experience itself produces a novel pattern that would not necessarily converge with choices purely from experience. In this regard, determining sufficient exposure to probabilistic outcomes is more a matter of deciding what specific decision-making context the laboratory task is an analogue of. Furthermore, our data suggested that decisions about losses informed by both description and experience were particularly resistant to the overriding effects of experience. It would be of interest to repeat Barron et al.'s procedure when the choice is between losses; based on our data, we would predict that the "late" group would show even less convergence with the "early" group.

This brings us to the second issue, in that having determined what is sufficient exposure to probabilistic outcomes, the next challenge is to ensure that this is delivered as intended to the participant. In the ACT, we assumed that a concurrent schedules task would provide sufficient experience with probabilistic outcomes that might have been lacking in our discounting procedure. Indeed, this largely proved to be the case in conditions that delivered frequent gains, but less so in tasks that delivered frequent losses. Although we observed a gain-loss asymmetry in sensitivity in the last experienced block of outcomes, most of these participants did not experience the full set of scheduled outcomes. In the Study 2 discussion, we explored some differences between the procedure Bull et al. (2015) used and the procedure we used that might explain their higher rate of component completion. These were the fact that their participants were paid contingent on performance and completed longer task sessions. However, when considered in the context of other studies using rapid-acquisition tasks (Lie et al., 2009; Krageloh et al., 2010), these features did not provide plausible explanations for the low rate of component completion in our data.

We observed rapid development of exclusive preference in the ACT conditions, more so in the gain rather than loss conditions. Finding a balance between exclusive preference and indifference (50-50 distribution of responses) is not a unique challenge to our task, and has been an issue of debate in both human and animal research using concurrent schedules. On the one hand, a COD was implemented to avoid reinforcing frequent switching behaviour that could lead to indifference (Shull & Pliskoff, 1967). On the other hand, we used dependent scheduling in our procedure, with the aim of decreasing the rate of exclusive responding on one alternative and increasing the rate of switching (Stubbs & Pliskoff, 1969). Dependent scheduling suspends further delivery of reinforcers until the scheduled one is picked up by the participant. While dependent scheduling does not eliminate instances of exclusive preference, with animal subjects that are able to run long sessions an eventual switch usually occurs, the reinforcer scheduling resumes, and the animal begins to experience something closer to the programmed distribution of outcomes. With humans, whose participation duration is often constrained by cost and fatigue, an eventual switch may not occur within a given session and the scheduled outcomes are not delivered by session completion. This was the case in our data, and to a lesser extent in Bull et al. (2015).

Our participants’ responses in the ACT showed an additional complication: when the two sides delivered frequent losses, the rate of exclusive preference was higher than in the context of frequent gains. While exclusive preference for the less frequently punishing side is suboptimal, in the sense that this would also suspend the delivery of both gains and losses and thus minimize net profit, it is optimal from the perspective of preventing experience of any further losses. Participants who developed this exclusive preference failed to maximize gains across the component, but they did succeed in minimizing losses. As Klapes et al. (2018) noted, since gains and losses are delivered based on the same behaviour, which is being reinforced by occurrence of a gain and punished by occurrence of a loss, ensuring that all scheduled outcomes are delivered would require changes to the experimental design.

One option for increasing the likelihood that participants experience all programmed outcomes is introducing a stimulus to signal overall success in the task, i.e. whether the participant is close to maximizing net reinforcement rate. The ACT included hints at the end of each component that prompted sampling of both sides and had a bar tracking overall net total of the points gained and lost. However, this did not indicate what the possible outcomes were should the participant change the pattern of their behaviour during the component. The advantages of a stimulus signalling this information seems particularly salient in tasks delivering frequent losses, as the benefits of switching do not seem to be apparent to the participants. Warry, Remington and Sonuga-Barke (1999) showed that information indicating what the favourable strategy was in a given task acted as a reinforcer for choosing, at times, what was the locally sub-optimal choice, but that resulted in an overall more optimal net profit by the session completion. While such stimuli may not be necessary with animal subjects, the study of experiential choice with humans may necessitate such procedural changes to increase the discriminability of the schedules in operation and reduce exclusive preference (Kollins et al., 1997; Horne & Lowe, 1993; Bradshaw et al., 1976; Madden & Perone, 1999; Warry et al., 1999).

Conclusion

We applied the predictions of the reversed reflection effect of the description-experience gap literature to probability discounting, and conducted a novel synthesis to the same effect with behaviour on concurrent schedules tasks. Our examination of gain-loss asymmetry included an analysis of discounting rates, sensitivity to frequency and magnitude, and local effects of gains and losses on behaviour. We add to a growing body of literature that uses experiential probability discounting tasks (e.g. Scheres et al., 2006; Greenhow et al., 2015; Hinvest & Anderson, 2010) and rapid-acquisition concurrent schedules tasks with humans (e.g. Lie et al., 2009; Krageloh et al., 2010; Bull et al., 2015). In both procedures, we found support for the predictions of a reversed reflection effect in that patterns in choice between gains was not symmetrical to patterns in choice between losses. In both cases, a gain held its value better with decreased probability than a loss. In the case of magnitude, we found support for the predictions of the value function of Prospect Theory, where responses were affected by magnitude of losses more than by magnitude of gains. We found less support for our predictions in the patterns of discounting and matching relative to expected value and ideal sensitivity respectively, although whether this was indicative of a difference in behavioural processes or procedural challenges was inconclusive. We interpret this as a matter of refining procedure in experiential tasks with humans before a more consistent gain-loss asymmetry may be observed, which requires careful consideration of the extent of exposure to chance outcomes and ensuring delivery of scheduled events.

References

Abdellaoui, M., L'Haridon, O., & Paraschiv, C. (2011). Experienced vs. described uncertainty: do we need two Prospect Theory specifications. *Management Science, 57*, 1879-1895. http://dx.doi.org/10.1287/mnsc.1110.1368

Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In E. Parzen, K. Tanabe, & G. Kitagawa (Eds.), *Selected Papers of Hirotugu Akaike* (pp. 199-213). New York: Springer.

Akl, E. A., Oxman, A. D., Herrin, J., Vist, G. E., Terrenato, I., Sperati, F., & Schunemann, H. (2011). Framing of health information messages. *Cochrane Database of Systematic Reviews, 12*, 1–82. https://doi.org/10.1002/14651858.CD006777.pub2

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine [Rational man's behaviour in face of risk: Critique of the American School's postulates and axioms]. *Econometrica, 21,* 503–546. http://dx.doi.org/10.2307/1907921

Andreoni, J., & Sprenger, C. (2011). *Uncertainty equivalents: Testing the limits of the independence axiom*. Cambridge: National Bureau of Economic Research, Inc. https://doi.org/10.3386/w17342

Asgarova, R., Macaskill, A. C., Robinson, B. J., & Hunt, M. (2017). Probability discounting and cardiovascular risk: The effect of side-effect severity and framing. *The Psychological Record, 67*, 169-179. https://doi.org/10.1007/s40732-017-0243-2

Barnett, V., & Lewis, T. (1978). *Outliers in Statistical Data*. Chichester: Wiley.

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. http://dx.doi.org/10.1002/bdm.443

Barron, G., Leider, S., & Stack, J. (2008). The effect of safe experience on a warnings' impact: Sex, drugs, and rock-n-roll. *Organizational Behavior and Human Decision Processes, 106,* 125-142. http://dx.doi.org/10.1016/j.obhdp.2007.11.002

Baucells, M., & Villasis, A. (2010). Stability of risk preferences and the reflection effect of prospect theory. *Theory and Decision, 68*, 193-211. https://doi.org/10.1007/s11238-009-9153-3

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behaviour*, *22*, 231-242. https://doi.org/10.1901/jeab.1974.22-231

Baum, W. M. (2002). From molecular to molar: a paradigm shift in behaviour analysis. *Journal of the Experimental Analysis of Behaviour, 78*, 95–116. https://doi.org/10.1901/jeab.2002.78-95

Baum, W. M. (2003). The molar view of behaviour and its usefulness in behaviour analysis. *The Behaviour Analyst Today, 4*, 78–81. http://dx.doi.org/10.1037/h0100009

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica, 22*, 23-36. https://doi.org/10.2307/1909829

Blampied, N. M. (2017). Analysing therapeutic change using modified Brinley plots: history, construction, and interpretation. *Behaviour Therapy, 48*, 115-127. https://doi.org/10.1016/j.beth.2016.09.002

Bleichrodt, H., Abellan-Perpiñan, J., Pinto-Prades, J., & Mendez-Martinez, I. (2007). Resolving Inconsistencies in Utility Measurement under Risk: Tests of Generalizations of Expected Utility. *Management Science, 53*, 469-482. https://doi.org/10.1287/mnsc.1060.0647

Bowen, H. J., & Kensinger, E. A. (2017). Cash or Credit? Compensation in Psychology Studies: Motivation Matters. *Collabra: Psychology, 3*, 12. http://doi.org/10.1525/collabra.77

Bradshaw, C. M., Szabadi, E., & Bevan, P. (1976). Behavior of humans in variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 26*, 135-141. https://doi.org/10.1901/jeab.1976.26-135

Brown, G. S., & White, K. G. (2005). The optimal correction for estimating extreme discriminability. *Behaviour Research Methods, 37*, 436-449. https://doi.org/10.3758/BF03192712

Budescu, D. V., & Weiss, W. (1987). Reflection of transitive and intransitive preferences: A test of prospect theory. *Organizational Behaviour and Human Decision Processes, 39*, 184-202. https://doi.org/10.1016/0749-5978(87)90037-9

Bull, P. N. (2013). *The role of sensitivity to reward and sensitivity to punishment in decision making on the Iowa Gambling Task* (Master's thesis). The University of Auckland, Auckland, New Zealand). Retrieved from https://researchspace.auckland.ac.nz/handle/2292/20125

Bull, P. N., Tippett, L. J., & Addis, D. R. (2015). Decision making in healthy participants on the Iowa Gambling Task: new insights from an operant approach. *Frontiers in Psychology, 6,* 391. https://doi.org/10.3389/fpsyg.2015.00391

Camilleri, A. R., & Newell, B. R. (2013). The long and short of it: Closing the description-experience "gap" by taking the long-run view. *Cognition, 126,* 54–71. http://dx.doi.org/10.1016/j.cognition.2012.09.001

Cohen, J. (1988). *Statistical power analysis for the behavioural sciences*. New York: Routledge. https://doi.org/10.4324/9780203771587

Cook, R. D., & Weisberg, S. (1982). *Residuals and influence in regression*. New York: Chapman & Hall.

Critchfield, T. S., Paletz, E. M., MacAleese, K. R., & Newland, M. C. (2003). Punishment in human choice: Direct or competitive suppression? *Journal of the Experimental Analysis of Behavior, 80*, 1-27. https://doi.org/10.1901/jeab.2003.80-1

Dai, Z., Grace, R. C., & Kemp, S. (2009). Reward contrast in delay and probability discounting. *Learning & Behavior, 37*, 281-288. https://doi.org/10.3758/LB.37.3.281

Davison, M., & Baum, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior, 74*, 1–24. https://doi.org/10.1901/jeab.2000.74-1

Davison, M., & Baum, W. M. (2002). Choice in a variable environment: Effects of blackout duration and extinction between components. *Journal of the Experimental Analysis of Behavior, 77*, 65–89. https://doi.org/10.1901/jeab.2002.77-65

Davison, M., & McCarthy, D. (1988). *The matching law: A research review*. Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Deluty, M. Z. (1976). Choice and the rate of punishment in concurrent schedules. *Journal of the Experimental Analysis of Behavior, 25,* 75–80. https://doi.org/10.1901/jeab.1976.25-75 de Villiers, P. A. (1980). Toward a quantitative theory of punishment. *Journal of the Experimental Analysis of Behavior, 33,* 15–25. https://doi.org/10.1901/jeab.1980.33-15

Elliffe, D., Davison, M., & Landon, J. (2008). Relative reinforcer rates and magnitudes do not control concurrent choice independently. *Journal of the Experimental Analysis of Behaviour, 90*, 169-185. https://doi.org/10.1901/jeab.2008.90-169

Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition, 34*, 914–928. https://doi.org/10.3758/BF03193437

Fantino, E., & Navarro, A. (2012). Description-experience gaps: Assessments in other choice paradigms. *Journal of Behavioral Decision Making, 25*, 303-314. https://doi.org/10.1002/bdm.737

Farley, J. (1980). Reinforcement and punishment effects in concurrent schedules: A test of two models. *Journal of the Experimental Analysis of Behavior, 33*, 311–315. https://doi.org/10.1901/jeab.1980.33-311

Field, A. P. (2009). *Discovering statistics using SPSS: And sex and drugs and rock 'n' roll* (3rd ed.). London: Sage.

Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General, 145,* 486–508. http://dx.doi.org/10.1037/a0040103

Gomes-Ng, S., Elliffe, D., & Cowie, S., (2017). How do reinforcers affect choice? Preference pulses after responses and reinforcers. *Journal of the Experimental Analysis of Behaviour, 108*, 17–38. http://dx.doi.org/10.1002/jeab.260.

Gomes-Ng, S., Landon, J., Elliffe, D., Bensemann, J., & Cowie, S. (2018). The effects of changeover delays on local choice. *Behavioural Processes, 150,* 36-46. https://doi.org/10.1016/j.beproc.2018.02.019

Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology, 38*, 129–166. https://doi.org/10.1006/cogp.1998.0710

Green, R. M., & Lawyer, S. R. (2014). Steeper delay and probability discounting of potentially real versus hypothetical cigarettes (but not money) among smokers. *Behavioural Processes, 108,* 50-56. https://doi.org/10.1016/j.beproc.2014.09.008

Green, L., & Myerson. J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, *130*, 769-792. https://doi.org/10.1037/0033-2909.130.5.769

Green, L., Myerson, J., Oliveira, L., & Chang, S. E. (2014). Discounting of delayed and probabilistic losses over a wide range of amounts. *Journal of the Experimental Analysis of Behavior*, *101*, 186–200. https://doi.org/10.1901/jeab.2014.101-186

Green, L., Myerson, J., & Ostaszewski, P. (1999). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *25*, 418-427. https://doi.org/10.1037/0278-7393.25.2.418

Greenhow, A. K., Hunt, M. J., Macaskill, A. C., & Harper, D. N. (2015). The effect of reinforcer magnitude on probability and delay discounting of experienced outcomes in a computer game task in humans. *Journal of the Experimental Analysis of Behaviour, 104*, 186-197. https://doi.org/10.1002/jeab.166

Hachiga, Y., Sakagami, T., & Silberberg, A. (2015). Preference pulses and the win–stay, fix-and-sample model of choice. *Journal of the Experimental Analysis of Behaviour*, *104*, 274-295. https://doi.org/10.1002/jeab.170

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: the role of sample size and experienced probabilities. *Journal of Behavioural Decision Making, 21,* 493–518. https://doi.org/10.1002/bdm.598

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of *d'*. *Behavior Research Methods, Instruments, & Computers*, *27*, 46-51. https://doi.org/10.3758/BF03203619

Herrnstein, R. J. (1961). Relative and absolute strength of response as function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, *4*, 267–272. https://doi.org/10.1901/jeab.1961.4-267

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13*, 243-266. https://doi.org/10.1901/jeab.1970.13-243

Hershey, J. C., & Schoemaker, P. J. H. (1980). Prospect Theory's reflection hypothesis: a critical examination. *Organizational Behaviour and Human Performance, 25*, 395-418. https://doi.org/10.1016/0030-5073(80)90037-9

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

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

Hinvest, N. S., & Anderson, I. M. (2010). The effects of real versus hypothetical reward on delay and probability discounting. *The Quarterly Journal of Experimental Psychology, 63*, 1072-1084. https://doi.org/10.1080/17470210903276350

Hoaglin, D., & Welsch, R. (1978). The hat matrix in regression and ANOVA. *American Statistician, 32*, 17–22. https://doi.org/10.1080/00031305.1978.10479237

Holt, D. D., Green, L., & Myerson, J. (2003). Is discounting impulsive?: Evidence from temporal and probability discounting in gambling and non-gambling college students. *Behavioural Processes, 64*, 355-367. https://doi.org/10.1016/S0376-6357(03)00141-4

Horne, P. J., & Lowe, C. F. (1993). Determinants of human performance on concurrent schedules. *Journal of the Experimental Analysis of Behavior, 59,* 29–60. https://doi.org/10.1901/jeab.1993.59-29

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science, 19*, 1015-1022. https://doi.org/10.1111/j.1467-9280.2008.02193.x

Johnson, M. W., & Bickel, W. K. (2008). An algorithm for identifying nonsystematic delay discounting data. *Experimental and Clinical Psychopharmacology, 16*, 264–274. https://doi.org/10.1037/1064-1297.16.3.264

Johnstone, V., & Alsop, B. (2000). Reinforcer control and human signal-detection performance. *Journal of the Experimental Analysis of Behavior, 73*, 275–290. https://doi.org/10.1901/jeab.2000.73-275

Kacelnik, A., & Bateson, M., (1996). Risky theories - the effects of variance on foraging decisions, *Integrative and Comparative Biology*, *36*, 402-434. https://doi.org/10.1093/icb/36.4.402

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–292. http://dx.doi.org/10.2307/1914185

Kalenscher, T., & van Wingerden, M. (2011). Why we should use animals to study economic decision making – a perspective. *Frontiers in Neuroscience, 5,* 1-11. https://doi.org/10.3389/fnins.2011.00082

Klapes, B., Riley, S., & McDowell, J. J. (2018). Toward a contemporary quantitative model of punishment. *Journal of the Experimental Analysis of Behaviour, 109,* 336-348. https://doi.org/10.1002/jeab.317

Kollins, S. H., Newland, M. C., & Critchfield, T. S. (1997). Human sensitivity to reinforcement in operant choice: How much do consequences matter? *Psychonomic Bulletin & Review*, *4*, 208-220. https://doi.org/10.3758/BF03209395

Krägeloh, C. U., & Davison, M. (2003). Concurrent schedule performance in transition: Changeover delays and signalled reinforcer ratios. *Journal of the Experimental Analysis of Behaviour, 79*, 87–109. https://doi.org/10.1901/jeab.2003.79-87

Krägeloh, C. U., Zapanta, A. E., Shepherd, D., & Landon, J. (2010). Human choice behaviour in a frequently changing environment. *Behavioural Processes, 83*, 119-126. https://doi.org/10.1016/j.beproc.2009.11.005

Kubanek, J., Snyder, L. H., & Abrams, R. A. (2015). Reward and punishment act as distinct factors in guiding behavior. *Cognition*, *139,* 154–167. https://doi.org/10.1016/j.cognition.2015.03.005

Kudryavtsev, A., & Pavlodsky, J. (2012). Description-based and experience-based decisions: Individual analysis. *Judgment and Decision Making, 7*, 316-331.

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for *t*-tests and ANOVAs. *Frontiers in Psychology, 4*(863), 1-12. https://doi.org/10.3389/fpsyg.2013.00863

Landon, J., Davison, M., & Elliffe, D. (2003). Concurrent schedules: Reinforcer magnitude effects. *Journal of the Experimental Analysis of Behaviour, 79,* 351–365. https://doi.org/10.1901/jeab.2003.79-351

Lawyer, S. R., Schoepflin, F., Green, R., & Jenks, C. (2011). Discounting of hypothetical and potentially real outcomes in nicotine-dependent and nondependent samples. *Experimental and Clinical Psychopharmacology, 19*, 263-274. https://doi.org/10.1037/a0024141

Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behaviour and Human Decision Processes, 116*, 286-295. https://doi.org/10.1016/j.obhdp.2011.05.001

Lejarraga, T., & Muller-Trede, J. (2016). When experience meets description: how dyads integrate experiential and descriptive information in risky decisions. *Management Science, 63,* 1953-1971. https://doi.org/10.1287/mnsc.2016.2428

Lie, C. (2010). Human conditional-discrimination performance: Local reinforcer and punisher effects. *Experimental Analysis of Human Behavior Bulletin, 28*, 1-6.

Lie, C., & Alsop, B. (2009). Effects of point-loss punishers on human signal-detection performance. *Journal of the Experimental Analysis of Behavior, 92*, 17-39. https://doi.org/10.1901/jeab.2009.92-17

Lie, C., Harper, D. N., & Hunt, M. (2009). Human performance on a two-alternative rapid-acquisition choice task. *Behavioural Processes, 81*, 244–249. https://doi.org/10.1016/j.beproc.2008.10.008

Lobb, B., & Davison, M. C. (1975). Performance in concurrent interval schedules: A systematic replication. *Journal of the Experimental Analysis of Behavior, 24,* 191-197. https://doi.org/10.1901/jeab.1975.24-191

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. http://dx.doi.org/10.1371/journal.pone.0020262

MacDonald, D. H., Kagel, J. H. & Battalio, R. C. (1991). Animals' choices over uncertain outcomes: Further experimental results. *The Economic Journal, 101*, 1065–1084.

Madden, G., & Perone, M. (1999). Human sensitivity to concurrent schedules of reinforcement: Effects of observing schedule-correlated stimuli. *Journal of the Experimental Analysis of Behavior, 71*, 303-318. https://doi.org/10.1901/jeab.1999.71-303

Magoon, M.A., & Critchfield, T. S. (2008). Concurrent schedules of positive and negative reinforcement: Differential-impact and differential-outcomes effects. *Journal of the Experimental Analysis of Behavior, 90,* 1–22. https://doi.org/10.1901/jeab.2008.90-1

Matusiewicz, A. K., Carter, A. E., Landes, R. D., & Yi, R. (2013). Statistical equivalence and test-retest reliability of delay and probability discounting using real and hypothetical rewards. *Behavioural Processes, 100*, 116-122. https://doi.org/10.1016/j.beproc.2013.07.019

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behaviour: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, New Jersey: Erlbaum.

McCord, M., & de Neufville, R. (1986). Lottery equivalents: Reduction of the certainty effect problem in utility assessment. *Management Science*, *32*. 56-60. https://doi.org/10.1287/mnsc.32.1.56

McDowell, J. J. (1989). Two modern developments in matching theory. *The Behavior Analyst, 12,* 153–166. https://doi.org/10.1007/BF03392492

McKerchar, T. L., Green, L., & Myerson, J. (2010). On the scaling interpretation of exponents in hyperboloid models of delay and probability discounting. *Behavioural processes*, *84,* 440–444. https://doi.org/10.1016/j.beproc.2010.01.003

McLean, A. P., Grace, R. C., Pitts, R. C., & Hughes, C. E. (2014). Preference pulses without reinforcers. *Journal of the Experimental Analysis of Behavior*, *101*, 317–336. https://doi.org/10.1002/jeab.84

Mitchell, S. H., & Wilson, V. B. (2010). The subjective value of delayed and probabilistic outcomes: Outcome size matters for gains but not for losses. *Behavioural Processes, 83,* 36–40. https://doi.org/10.1016/j.beproc.2009.09.003

Mitchell, S. H., Wilson, V. B., & Karalunas, S. L. (2015). Comparing hyperbolic, delay-amount sensitivity and present-bias models of delay discounting. *Behavioural processes*, *114*, 52–62. https://doi.org/10.1016/j.beproc.2015.03.006

Myerson, J., Green, L., Hanson, J. S., Holt, D. D., & Estle, S. J. (2003). Discounting delayed and probabilistic rewards: Processes and traits. *Journal of Economic Psychology, 24*, 619-635. https://doi.org/10.1016/S0167-4870(03)00005-9

Myerson, J., Green, L., & Warusawitharana, M. (2001). Area under the curve as a measure of discounting. *Journal of the Experimental Analysis of Behavior, 76*, 235-243. https://doi.org/10.1901/jeab.2001.76-235

Ohmura, Y., Takahashi, T., Kitamura, N., & Wehr, P. (2006). Three-month stability of delay and probability discounting measures. *Experimental and Clinical Psychopharmacology, 14*, 318–328. https://doi.org/10.1037/1064-1297.14.3.318

Peters, J., & Büchel, C. (2009). Overlapping and distinct neural systems code for subjective value during intertemporal and risky decision making. *The Journal of Neuroscience, 29*, 15727–15734. https://doi.org/10.1523/JNEUROSCI.3489-09.2009

Rachlin, H. C. (2006). Notes on discounting. *Journal of the Experimental Analysis of Behavior, 85*, 425-435. https://doi.org/10.1901/jeab.2006.85-05

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behaviour, 5*5, 233–244. https://doi.org/10.1901/jeab.1991.55-233

Rasmussen, E. B. & Newland, M. C. (2008). Asymmetry of reinforcement and punishment in human choice. *Journal of the Experimental Analysis of Behavior, 89*, 157-167. https://doi.org/10.1901/jeab.2008.89-157

Real, L. (1996). Paradox, performance, and the architecture of decision-making in animals. *American Zoologist*, *36*, 518-529. http://dx.doi.org/10.1093/icb/36.4.518

Richards, J. B., Mitchell, S. H., de Wit, H., & Seiden, L. S. (1997). Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the Experimental Analysis of Behaviour, 67*, 353-366. https://doi.org/10.1901/jeab.1997.67-353

Scheres, A., Dijkstra, M., Ainslie, E., Balkan, J., Reynolds, B., Sonuga-Barke, E., & Castellanos, F. X. (2006). Temporal and probabilistic discounting of rewards in children and adolescents: Effects of age and ADHD symptoms. *Neuropsychologia, 44*, 2092-2103. https://doi.org/10.1016/j.neuropsychologia.2005.10.012

Shlomi, Y. (2014). Subjective integration of probabilistic information from experience and description. *Judgment & Decision Making, 9*, 491–499. http://hdl.handle.net/1903/9922

Schmitt D. R. (1974). Effects of reinforcement rate and reinforcer magnitude on choice behavior of humans. *Journal of the Experimental Analysis of Behavior*, *21,* 409–419. https://doi.org/10.1901/jeab.1974.21-409

Schneider, S. L., & Lopes, L. L. (1986). Reflection in preferences under risk: Who and when may suggest why. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 535-548. http://dx.doi.org/10.1037/0096-1523.12.4.535

Schoemaker, P. J. H. (1990). Are risk-attitudes related across domains and response-modes. *Management Science*, *36,* 1451–1463. https://doi.org/10.1287/mnsc.36.12.1451

Shafir, A., Reich, T., Tsur, E., Erev, I., & Lotem, A. (2008). Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature, 453,* 917-920. https://doi.org/10.1038/nature06841

Shafran, A. (2011). Self-protection against repeated low probability risks. *Journal of Risk and Uncertainty, 42*, 263-285. https://doi.org/10.1007/s11166-011-9116-2

Shead, N. W., Callan, M. J., & Hodgins, D. C. (2008). Probability discounting among gamblers: Differences across problem gambling severity and affect-regulation expectancies. *Personality and Individual Differences, 45*, 536–541. https://doi.org/10.1016/j.paid.2008.06.008

Shead, N. W., & Hodgins, D. C. (2009). Probability discounting of gains and losses: Implications for risk attitudes and impulsivity. *Journal of the Experimental Analysis of Behaviour, 92*, 1-16. https://doi.org/10.1901/jeab.2009.92-1

Shull, R. L., & Pliskoff, S. S. (1967). Changeover delay and concurrent schedules: some effects on relative performance measures. *Journal of the Experimental Analysis of Behavior, 10,* 517–527. https://doi.org/10.1901/jeab.1967.10-517

Smith, K. R., Lawyer, S. R., & Swift, J. K. (2018). A meta-analysis of nonsystematic responding in delay and probability reward discounting. *Experimental and Clinical Psychopharmacology, 26*, 94-107. https://doi.org/10.1037/pha0000167

Starmer, C. (2000). Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature, 28*, 332–382. https://doi.org/10.1257/jel.38.2.332

Stubbs, D. A., & Pliskoff, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior 12*, 887–895. https://doi.org/10.1901/jeab.1969.12-887

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5,* 297–323. http://dx.doi.org/10.1007/BF00122574

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behaviour*. Princeton, NJ, US: Princeton University Press.

Warry, C. J., Remington, B., Sonuga-Barke, E. J. S. (1999). When more means less: factors affecting human self-control in a local versus global choice paradigm. *Learning and Motivation*, *30*, 53-73. https://doi.org/10.1006/lmot.1998.1018

Weatherly, J. N., & Derenne, A. (2013). Probability and delay discounting of gains and losses using the multiple-choice method. *The Psychological Record*, *63*, 563-582. https://doi.org/10.11133/j.tpr.2013.63.3.011

Weatherly, J. N., & Terrell, H. K. (2014). Magnitude effects in delay and probability discounting when monetary and medical treatment outcomes are discounted. *Psychological Record*, *64*, 433–440. https://doi.org/10.1007/s40732-014-0052-9

Wu, G., & Gonzalez, R. (1996). Curvature of the probability weighting function. *Management Science*, *42*, 1676-1690. https://doi.org/10.1287/mnsc.42.12.1676

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin, 144*, 140-176. http://dx.doi.org/10.1037/bul000011

Wurster, R. M., & Griffiths, R. R. (1979). Human concurrent performances: Variation of reinforcer magnitude and rate of reinforcement. *The Psychological Record*, *29*, 341-354. https://doi.org/10.1007/BF03394622

Yi, R., de la Piedad, X., & Bickel, W. K. (2006). The combined effects of delay and probability in discounting. *Behavioural Processes*, *73*, 149-155. https://doi.org/10.1016/j.beproc.2006.05.001

Yi, R., Stuppy-Sullivan, A., Pickover, A., & Landes, R. D. (2017). Impact of construal level manipulations on delay discounting. *PLoS ONE*, *12*, 1-18. https://doi.org/10.1371/journal.pone.0177240

Young, M. E. (2017). Discounting: a practical guide to multilevel analysis of indifference data. *Journal of the Experimental Analysis of Behaviour, 108*, 97-112. https://doi.org/10.1002/jeab.265

Appendix A: Group median discounting curves reproduced from Experiment 1 and 2 of my honours thesis

Subjective value (indifference points as a proportion of the larger, uncertain amount) as a function of increasing odds against occurrence of a gain (black triangles) or loss (black squares). Dashed (gains) and solid (losses) curves are the best-fitting hyperbolic functions. The dotted curve is a hyperbolic function from decisions made based on expected value (h = 1). Top row are data from Experiment 1 money (left) and ski (right) tasks. Bottom row are data from Experiment 2 money (left) and ski (right) tasks.

Non-experiential safe-risky money

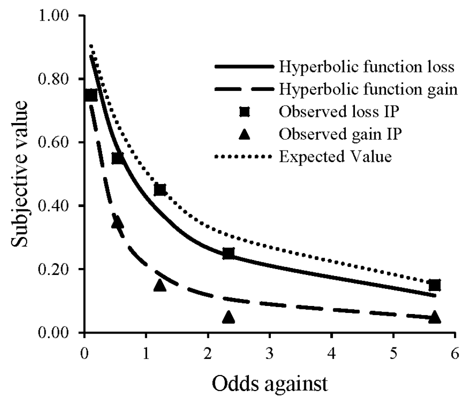

Experiential safe-risky ski

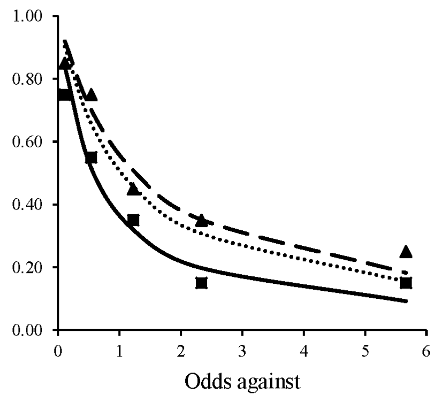

Experiential safe-risky money

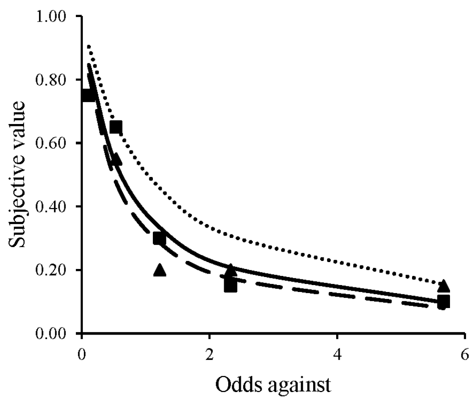

Experiential safe-risky ski

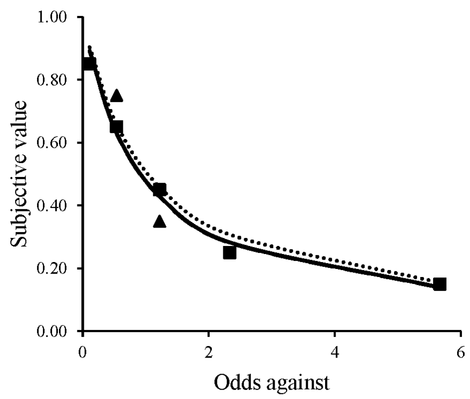

Appendix B: Analysis of mean residuals for Study 1, Experiments 2-3

Mean residuals calculated from individual data for each task as a function of odds against in the gain (white circles) and loss (black circles) conditions. Error bars are standard error of the mean.

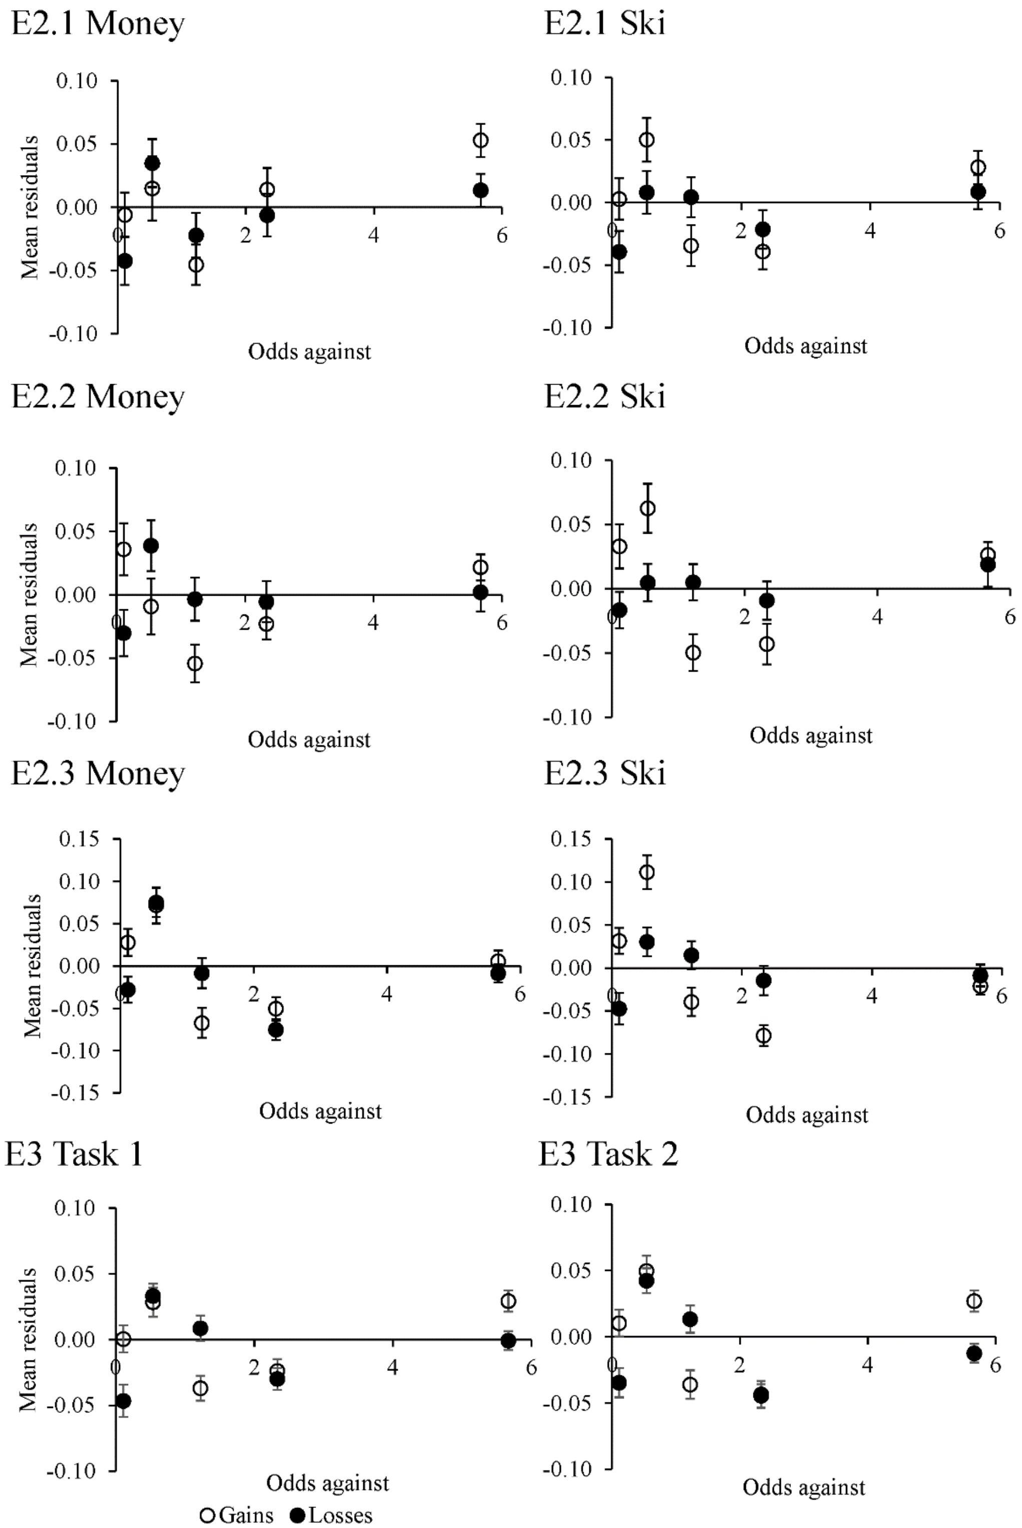

Appendix C: ANOVA results for Study 1, Experiment 2 using AUCs

	F	df	p	η_p^2	F	df	p	η_p^2
	Money				Ski			
E2.1	Safe-risky experiential				Safe-risky experiential			
Condition	0.33	1, 32	.567	.010	0.003	1, 33	.957	<.001
Condition order	<0.001		.985	<.001	0.02		.892	.001
Task order	0.47		.500	.014	1.13		.295	.033
Condition by Condition order	**10.25**		**.003**	**.243**	**8.03**		**.008**	**.196**
Condition by Task order	0.04		.846	.001	0.62		.437	.018
Condition order by Task order	0.55		.466	.017	3.16		.085	.087
Condition by Condition order by Task order	1.92		.175	.057	0.55		.463	.016
E2.2	Risky-risky non-experiential				Risky-risky experiential			
Condition	3.48	1, 42	.069	.076	0.03	1, 44	.855	.001
Condition order	**5.37**		**.025^**	**.113**	1.19		.282	.026
Task order	0.01		.944	<.001	0.16		.693	.004
Condition by Condition order	0.06		.803	.001	0.03		.855	.001
Condition by Task order	0.06		.801	.002	0.83		.369	.018
Condition order by Task order	3.09		.086^	.069	0.01		.920	<.001
Condition by Condition order by Task order	<0.001		1.00	<.001	0.10		.751	.002
E2.3	Risky-risky experiential				Risky-risky experiential			
Condition	0.64	1, 47	.428	.013	0.81	1, 47	.372	.017
Condition order	0.58		.450	.012	0.43		.517	.009
Task order	**6.54**		**.014**	**.122**	2.27		.139	.046
Condition by Condition order	**27.76**		**<.001**	**.371**	0.79		.379	.017
Condition by Task order	0.03		.862	.001	1.66		.204	.034
Condition order by Task order	0.20		.660	.004	0.002		.962	<.001
Condition by Condition order by Task order	0.39		.534	.008	0.38		.539	.008

Note. Bold emphasis added to significant results. ^ indicates results that differed when using log(h) (Table 2.8).

Appendix D: ANOVA results for Study 1, Experiment 3 using AUCs

	F	df	p	η_p^2	F	df	p	η_p^2
		Task 1				Task 2 (Task 2 order)		
Condition	2.88	1, 114	.093	.025	2.92	1, 114	.090	.025
Condition order	0.27		.602	.002	0.14		.707	.001
Condition by Condition order	**20.33**		**<.001**	**.151**	0.25		.617	.002
						Task 2 (Task 1 order)		
Condition					2.90	1, 114	.092	.025
Condition order					0.16		.690	.001
Condition by Condition order					2.63		.108	.023

Note. Bold emphasis added to significant results

Appendix E: Instructions used in Auckland Card Task

Starting instruction screen:

You will be presented with two virtual decks of playing cards. Each deck contains hundreds of cards. Shuffled into each deck are some WINNING cards, which will add money to your overall score, and some LOSING cards, which will subtract money from your score. One deck will always be better than the other, and your goal is to maximize your winnings by learning which deck is better.

Each game has four rounds, and you'll be given a chance to rest between each round. The good deck may change from round to round, so always attend closely to the winning and losing cards you receive. These cards will vary a lot, so at the start of each round you'll need to sample the cards from each deck until you can figure out which deck is better.

In each game you will be given a hint (which is repeated during rest breaks). It's important that you read and understand the hint, as it will provide you with a strategy to maximize your winnings. Please press spacebar to continue.

Instructions/hint for gain frequency condition:

Winning cards can be found in both decks, but one deck has MORE winning cards than the other. Both decks also contain an equal number of losing cards. In each round, to maximize your score in the time given, you'll first need to figure out which deck has more winning cards in it, then choose more often from that deck.

Instructions/hint for loss frequency condition:

High winning cards can be found in both decks, but there are also many losing cards. One deck has MORE losing cards than the other. In each round, to maximize your score in the time given, you'll first need to figure out which deck has more losing cards in it, then choose less often from that deck, only occasionally checking it for winning cards.

Instructions/hint for gain magnitude condition:

Winning cards can be found in both decks, but one deck has HIGHER winning dollar amounts (on average) than the other. Both decks also contain an equal number of losing cards. In each round, to maximize your score in the time given, you'll first need to figure out which deck has the higher winning cards (but note the amounts vary a lot), then choose more often from that deck.

Instructions/hint for loss magnitude condition:

High winning cards can be found in both decks, but there are also many losing cards. One deck has HIGHER losing dollar amounts (on average) than the other. In each round, to maximize your score in the time given, you'll first need to figure out which deck has the higher losing cards (but note the amounts vary a lot), then choose less often from that deck, only occasionally checking it for winning cards.

Instructions for how to respond:

During the game, press the 'Caps Lock' key on the keyboard to choose a card from the left deck, or the 'Enter' key to choose from the right deck. You are free to switch from one deck to the other at any time. Please use only your dominant hand to respond - do NOT use both hands. If you have any questions, feel free to ask the researcher now. Please press spacebar to start the game.

Appendix F: Two-stage hierarchical regression analysis using AUC for losses in Study 2

The hierarchical regression showed that adding loss frequency and magnitude sensitivity did not account for significantly more variance in discounting than condition order alone. Both Models 1 ($F(1, 90) = 6.95$, $p = .010$) and 2 ($F(3, 88) = 2.78$, $p = .046$) were significantly better at predicting discounting than the mean. Model 1 showed that condition order was a significant predictor, accounting for 6.10% of the variation in AUC for losses. In Model 2, adding loss frequency and loss magnitude sensitivities did not explain significantly more variance in AUC than Model 1.

Correlations between AUC for losses, condition order, and sensitivities to loss frequency and magnitude

	AUC loss	Order	LF	LM
AUC loss	-	-0.27*	-0.02	-0.12
Order	-	-	-0.07	0.02
LF	-	-	-	-0.02
LM	-	-	-	-

Note. *$p < .01$.

Results of a two-stage hierarchical regression analysis using AUC for losses

					Change statistics			
	B (SE)	β	*t*	*p*	*F*	*df*	*p*	adj. R^2
Model 1: Order alone					6.95	1, 90	.010	.06
Order	0.05 (0.02)	0.27	2.64	.010				
Model 2: Order and sensitivity					0.71	2, 88	.494	.06
Order	0.06 (0.02)	0.27	2.65	.010				
LF	0 (0.01)	0.001	0.01	.994				
LM	-0.02 (0.01)	-0.12	-1.19	.237				

Appendix G: Two-stage hierarchical regression analysis using AUC for gains in Study 2

The hierarchical regression showed that neither order alone (*F*(1, 92) = 0.16, *p* = .690) nor the inclusion of sensitivities (*F*(3, 90) = 0.17, *p* = .918) were significantly better at predicting discounting than the mean. Table X shows that none of the predictors significantly explained variance in the discounting of gains.

Correlations between AUC for gains, condition order, and sensitivities to gain frequency and magnitude

	AUC gain	Order	LF	LM
AUC gain	-	-0.04	-0.03	0.06
Order	-	-	0.09	-0.19
LF	-	-	-	0.22*
LM	-	-	-	-

Note. *$p < .05$.

Results of a two-stage hierarchical regression analysis using AUC for gains

					Change statistics			
	B (SE)	β	*t*	*p*	*F*	*df*	*p*	adj. R^2
Model 1: Order alone					0.16	1, 92	.690	-.01
Order	-0.01 (0.02)	-0.04	-0.40	.690				
Model 2: Order and sensitivity					0.17	2, 90	.842	-.03
Order	-0.01 (0.02)	-0.03	-0.25	.801				
LF	-0.003 (0.01)	-0.04	-0.33	.739				
LM	0.01 (0.01)	0.06	0.55	.583				

Appendix H: Analysis of mean residuals for Studies 2-3

Mean residuals calculated from individual data for each task as a function of odds against in the gain (white circles) and loss (black circles) conditions. Error bars are standard error of the mean.