What are the chances? How descriptions of probability affect risky decisions.

By

Jared Pickett

A thesis submitted to Victoria University of Wellington in fulfilment of the
requirements for the degree of Master of Science

Victoria University of Wellington

2018

Abstract

People make different decisions when they know the odds of an event occurring, (e.g. told 10% chance of an earthquake that year) than when they draw on only their own experience (e.g. living in a city with, on average, one earthquake every 10 years). It may be that when we make decisions based on our past experience (decisions from experience) we are more likely to choose a risky option when it can lead to the biggest win and avoid it when it can lead to the biggest loss, this effect is called the Extreme-Outcome rule. Across three Experiments we tested the Extreme-Outcome rule by having participants make repeated choices between either safe or risky options which had the same expected value. In each experiment, we varied the magnitude of the reinforcer's participants could win in both an Experience condition and a condition that had both description and experience information.

In Experiment 1 where we had two reinforcer sizes (small and large) we found an Extreme-Outcome effect in the Experience condition, but not the Description-Experience condition. In Experiment 2 we tested a prediction of the Extreme-Outcome rule that participants would be sensitive to the best and worst outcome by adding another reinforcer size (reinforcers were small, medium and large) and therefore on some trials neither alternative included an extreme outcome. We also removed zero as a potential outcome to investigate whether zero aversion might be driving the effect of reinforcer magnitude in the Experience condition. We did not find response patterns consistent with an Extreme-Outcome rule in the Experience condition. Instead, participants were least risk seeking when the reinforcer was small, but there was no difference in levels of risk seeking between the medium and large reinforcer trials. In other words, there was an effect of the low-extreme outcome but not the high-extreme outcome. Like Experiment 1, in the Description-Experience condition risk preference was not influenced by reinforcer size, but the absolute levels were higher. To investigate whether this increase in risk preference was due to removing the zero, in Experiment 3 we manipulated

whether zero was present or absent. When zero was absent, risk preference was not influenced by the size of the reinforcer in the Description-Experience condition, but there was an effect of the low-extreme outcome when zero was present. We also found an effect of the low extreme outcome in the Experience condition regardless of whether zero was present or absent. Overall, these findings suggest the Extreme-Outcome rule needs to be modified to take into account the effect of the low extreme but not the high extreme outcome.

Acknowledgments

I owe an enormous amount of gratitude to the team of mentors, friends and family who have helped made this thesis happen. Firstly, thank you to my supervisors, Anne and Maree. Thank you for all the opportunities you have given me, from tutoring, giving presentations at conferences, speaking in lectures, collecting data and the great feedback I have received. I have learned a lot, all the way from how to give presentations to how “it’s possible to not know how to use an apostrophe and still be a good person”. Thank goodness. If I made one correct choice during my time at university, it was my supervisors. Not everyone is as lucky to have such amazing mentors. I hope I can one day pay it forward in my own career.

Chance plays a large part in the people you meet in postgrad. Here, I have been lucky again. I’ve met a fantastic bunch of people and they have made this time more entertaining and worthwhile. Thank you for making the Honours room a chat room rather than a work room, for the BYOs, Asbestos lunches and 2 for 1 pizzas (which we will need now more than ever, since we are unemployed).

Two things have derailed this project at multiple times throughout the year: Robust ANOVAs and Tessa Burgess. It’s been a deeply satisfying relationship, character building and wonderfully enjoyable. While at times challenging and confusing, it has made this project and this year all the better for it. And Tessa you’re okay too. ;) But to be fair, thank you for your help over the year, for your fresh perspective on life, introducing me to new words like Hygge, and for teaching me that the New World on the waterfront is called Chaffers (who knew?). We all need a Tessa in our lives. So thank you for being in mine.

Thank you to my lab mates, Rana, Rebecca and Geraldine for not only a friendly and supportive lab but also for giving me free food and letting me know about emails that have the phrase “refreshments will be provided” in it. I probably would have starved without you.

Thanks to my flatmates, Tom and Andrew, whose constant puns make my insides crawl. They really bug me. But thanks for the great flat, it’s really intense. Much like the Mongolians.

Thank you to Sophia, for all the lunch’s throughout the year they have invaluable. And for being the definition of a good friend.

To my family, Mum, Dad and Mark, thank you for your help over the years and the frequent phone calls and tolerance for my delay in getting back to you. I wouldn’t have done this without you.

So thank you all.

Table of Contents

Acknowledgments.....	4
List of Figures	7
List of Tables	8
Introduction.....	Error! Bookmark not defined.
Experiment 1	26
Method	28
Results	31
Discussion	42
Experiment 2	45
Method	45
Results	49
Discussion	58
Experiment 3	63
Method	63
Results	64
Discussion	75
General Discussion	80
References.....	92

List of Figures

<i>Figure 1:</i> Three different procedures to study decisions from experience	14
<i>Figure 2:</i> Screen shot of the task in Experiment 1	29
<i>Figure 3:</i> Risk preference in Experiment 1	34
<i>Figure 4:</i> Modified brinley plot of risk preference in the Experience condition for Experiment 1	35
<i>Figure 5:</i> Modified brinley plot for risk preference description-experience condition in Experiment 1	38
<i>Figure 6:</i> Modified brinley plot showing risk preference in the Experience condition.	38
<i>Figure 7:</i> Risk preference over time in the Experience condition.	41
<i>Figure 8:</i> Risk preference over time in the Description-Experience condition	42
<i>Figure 9:</i> Risk preference in Experiment 2	49
<i>Figure 10:</i> Modified brinley plots of risk preference in Experiment 2	52
<i>Figure 11:</i> Modified brinley plot of recency effect in Experiment 2	54
<i>Figure 12:</i> Recency Effect score in Experiment 2	55
<i>Figure 13:</i> Risk Preference changed over time in Experiment 2	57
<i>Figure 14:</i> Risk Preference in Experiment 3.	68
<i>Figure 15:</i> Modified Brinley plots of risk preference in Experiment 3	69
<i>Figure 16:</i> Recency effect score in Experiment 3.	71
<i>Figure 17:</i> Modified Brinley plots showing recency effect in Experiment 3	72
<i>Figure 18:</i> Risk preference over time in Experiment 3 in the Description-Experience condition	74
<i>Figure 19:</i> Risk preference over time in Experiment 3 in the Experience condition.	74
<i>Figure 20:</i> Objective difference and perceived difference of diminishing sensitivity	85

List of Tables

<i>Table 1:</i> Comparing decisions from description to decisions to experience	18
<i>Table 2:</i> Trial types in Experiment1	29
<i>Table 3:</i> Post Hoc tests between groups in the Experience condition	50
<i>Table 4:</i> Points used for when zero is present or absent in Experiment 3	64
<i>Table 5:</i> Post Hoc results comparing the influence the reinforcer size and the presence of zero had on risk preference in each condition	66

What are the chances? How descriptions of probability affect risky decisions

“The pattern is that every 2000 to 3000 years there is a monstrous eruption (of Mount Vesuvius). And it has now been about 2000 years.” – Michael Sheridan (Vergano, 2006)

As the quote above alludes to, around 2000 years ago in one of the most destructive eruptions in European history, Mount Vesuvius erupted and destroyed the Roman city of Pompeii (Scandone, Giacomelli, & Gasparini, 2000). Today, Mount Vesuvius remains one of the most dangerous volcanoes in the world, in part because the city of Naples (with a population of around 3 million people) lies so close. Experts on Mount Vesuvius estimate that in the next 80 years there is a 10% chance of a large-scale eruption, which would have a catastrophic effect on human life. One of these experts, Dr Mastrolorenzo, argues the whole city of Naples should be evacuated as it is too dangerous to live there (Mastrolorenzo et al, 2006). So it is surprising that the people you would think would be most concerned about Mount Vesuvius, the people of Naples, don't seem that concerned. In fact, to get people to relocate the Italian Government is willing to pay just under \$50,000 NZ dollars to anyone who will move away from Naples, yet hardly anyone has taken them up on the offer. When we look at this situation, why are different groups of people making different decisions based on the same risky event (the possibility of Mount Vesuvius erupting)? In recent years studies have shown that *how* we learn about a risky event influences the choices we make. In this situation, the experts are making their decision largely off a probability (10% chance of an eruption). These types of decisions, where we know the chances of an event occurring are called decisions from description. On the other hand, while the public might have known the probability too, they also had access to another form of information: their personal experience of living with Mount Vesuvius on their door step. When we make decisions using our own personal experience we call these decisions from experience. It is an important to make the

distinction between these two types of decisions because (as the Mount Vesuvius example shows) people make different choices when they make decisions from description compared to when they make decisions from experience, this effect is called the description-experience gap (DE gap) (Hertwig & Frey, 2017).

Why should we study decisions from experience?

Largely, decision-making research has focused on how we make decisions from description. In these types of studies, participants are generally asked question such as, *and would you prefer \$50 or a 50% chance of \$100?* (Tversky & Kahneman, 1992) or, *if a disease breaks out, would you save 200 people or have a one in three chance of saving 600?* (Tversky & Kahneman, 1981). Studies on decisions from description have been useful in understanding how we make choices involving risk. However, the implications we can draw from this body of research are limited, because often when we make decisions we do not know the chances of an uncertain event occurring. For example, should I back up my hard drive? Is it safe to cross the street? Do I need to move away from Naples? To make these decisions, rather than rely on knowing the probability of an event occurring (which we often do not know) we instead rely on our past experience of making the same or similar choice. Therefore, to have a more complete picture of how people make decisions under uncertainty, we need a better understanding of how we make decisions from experience.

To illustrate why studying decisions from experience helps us better understand decision making consider what prospect theory, the most influential theory on decision making, would have predicted about how people would have behaved to the risks of Mount Vesuvius. According to prospect theory people overweigh rare events (Kahneman & Tversky, 1979b) - a rare event is an outcome that has a 20% chance or less of occurring (Hertwig et al., 2004). In other words, when there is a 10% chance an event will occur,

people act as if the event is more likely to happen than it actually is. Thus, prospect theory says people should be more sensitive to the risk of Mount Vesuvius than objective probability suggests they should be, much like the experts. However, prospect theory fails to explain why the people of Naples are less sensitive to the dangerous of Mount Vesuvius than objective probability would suggest. This is not a criticism of prospect theory as it has always been a theory about how we make decisions from description and doesn't pretend to explain decisions from experience. However, what this does illustrate is the importance of understanding decisions from experience and how much our current understanding is lacking. A flurry of studies in the last decade show that there is a consistent difference between decisions from description and decisions from experience. We find a DE gap in non-humans (Heilbrunner & Hayden, 2016), when people are making medical decisions (Armstrong & Spaniol, 2017), when choices are made under time pressure (Camilleri & Newell, 2013b), in both between- and within- subjects designs (Camilleri & Newell, 2011; Ludvig & Spetch, 2011) and in field studies (Yechiam, Erev, & Barron, 2006).

How do we study decisions from experience?

Studying decisions from description is straightforward: tell participants the probability of each outcome occurring, for example, *would you prefer a 100% chance of \$3 or an 80% chance of \$4?*, and have them tell you which one they want. However, if you want to study how people make decisions from experience you need to find another way to represent the same choice alternatives without descriptions. Figure 1 shows the three main experimental procedures used for studying decisions from experience: the sampling method, partial feedback method and full feedback method. In all three procedures, rather than explicitly stating the chances of an event occurring, participants are shown different options on a screen (e.g. different coloured buttons), each option is associated with a different payoff distribution. Going back to our example earlier, one button might represent a 100% chance

of \$3 and another represents an 80% chance of \$4. We now turn to each procedure in more detail.

The Sampling Procedure

The sampling procedure has two parts: a sampling phase and a choice phase. In the sampling phase participants can choose each outcome as many times as they like in order to learn the payoff distribution of the two options. Each choice in the sampling phase does not affect their balance (the total number of points they have). Rather, participants get feedback about how many points each option would have given them (Noguchi & Hills, 2016). Once a participant has decided they have sampled enough from each option they can decide to move into the choice phase where they have one final choice between each option. Only points won in this choice phase are added to their overall balance (Yoon, Vo, & Venkatraman, 2017). Think of the sampling procedure as like a wine tasting event; you try different types of wine and then when you are happy with the amount that you have sampled from you can choose which bottle to buy.

The Partial and Full Feedback Procedure

Unlike the sampling method the partial feedback procedure has no sampling phase and every decision is consequential (Madan, Ludvig, & Spetch, 2016). Therefore, participants have to juggle earning as many points as possible with exploring other outcomes for a potentially higher paying alternative. This juggling act is known as the exploration-exploitation tension. Many decisions in life resemble the partial feedback procedure, for example, deciding which restaurant to eat at: you get feedback on the option you choose (whether you liked the food or not) but you won't know what would have happened had you chose the other restaurant. The full feedback procedure is the similar to the partial feedback procedure except participants also get feedback about what would have happened had they

chosen the other option (Camilleri & Newell, 2011) . Using the example from Figure 1, say you chose the button on the left and received \$3, you will also find out that, had you picked the right button, you would have received \$4. For a real-life example, consider the decision of whether to invest in a stock or not. Regardless of your choice, you will know what happened had you chose the other option. For example, if you decided not to invest in Coca-Cola, you can see what happened to the value of that stock and thus how much you would have gained or lost.

Participants are shown two squares on a screen and each trial are asked to choose one

Figure 1: The difference between the three procedures used to study decisions from experience. Each row represent successive choice trials presented to participants. Grey circles indicate the alternative the participant chose. Each square box represents an option the participants can choose. Before making a choice, both boxes are blank. Any numbers seen are feedback after a choice has been made. Squares with dashed lines represent choices that have no effect on balance, solid lines do effect balance. Grey writing represents the outcome the participant would have got if they chose the other option

Looking at the differences and similarities between people's choices in the sampling, partial and full feedback procedures can reveal clues about what factors influence the description-experience gap. For example, as outlined above, in the partial feedback procedure, participants must strike a balance between earning as many points as they can and yet exploring other options to see whether there might be a higher paying alternative. This exploration-exploitation trade-off is unique to the partial feedback procedure, therefore, comparing people's choices between the partial feedback, sampling and full feedback procedures allows us to access how important the exploration-exploitation trade-off is in decisions from experience. However, a DE gap has been found using all three of these experience procedures (Erev et al., 2010). Thus, the exploration-exploitation trade-off is unlikely to be the main driving force behind the DE gap.

Comparing the three experience procedures

Despite the fact a DE gap has been found with all three procedures, how this gap manifests can differ depending on which experimental task is used. For example, Camilleri and Newell (2011) had participants make several choices between a risky and a safe option. There were four conditions, a description condition, and the three experience conditions: sampling, partial feedback and full feedback. As expected when they compared each of the three experience conditions to the description condition they found a DE gap. However, while there was no difference between the partial and full feedback conditions, when they compared both the feedback conditions to the sampling procedure, they found a significant difference between the groups.

Participants in the feedback conditions tended to underweight rare events (e.g. behaving as though a 10% chance of an event occurring was less likely to occur than it is), while those in the sampling procedure tended to be risk averse. To illustrate, consider

one choice presented in their study: one alternative produced 9 points for sure and the other 90% chance of 10 points. In this case, the rare event is choosing the risky option and not getting any points. If you underweight rare events you act as if the rare event is less likely to occur than it is, thus people should tend to choose the risky option (which those in the two feedback conditions did). Whereas if you are risk averse, (like those in the sampling condition) you will be more likely to take the 100% chance of 9 points. The difference between participants' choices in the sampling condition and those in the description condition was those in the description condition were more risk averse than those in the sampling condition. Therefore, when you compare the description condition to each of the experience conditions a DE gap emerges even though this gap looks a bit different depending on the experience group.

What explains the difference between the feedback and sampling conditions?

Camilleri and Newell (2011) argued the repeated and consequential decisions that characterize the feedback conditions but not the sampling condition may drive the difference between the paradigms. According to Camilleri and Newell repeated, consequential, choices encouraged participants to choose the option they chose last regardless of the outcome. However, other studies suggest whether the last outcome is favourable or not does influence the participant's next choice. For example, when people choose a risky option and win, they are likely to choose the same risky option in the next trial, whereas if they lose they are likely to switch to the other option (Madan, Ludvig, & Spetch, 2014). This type of recency effect is known as a win-stay/lose-shift pattern. However, the study conducted by Camilleri and Newell does suggest that consequential decisions may drive the difference between the experimental procedures. To summarise, people make decisions differently when they make them from experience compared to

description, but they may also make them differently depending on how they learn about uncertain events from experience.

How are decisions from description different to decisions from experience?

Before we explore what we know about how decisions from experience differ to decisions from description it's important to point out that the definition of the DE gap is vague. Most definitions will say essentially that the DE gap is where people make different choices when they make decisions from description compared to decisions from experience. Researchers often don't say how specifically they are different (Yoon et al., 2017). Or if they do, they will only talk about what happens when small probabilities are involved (Noguchi & Hills, 2016; Regenwetter & Robinson, 2017). This is because (as we will see) there is no one way to explain how decisions from description differ to decisions from experience. That is because how they differ depends on many factors, such as the probability of the uncertain event, whether the choice is a gain or loss etc. A common measure used to discuss the difference in choices between decisions from description and decisions from experience is risk preference, which is the proportion of times the risky option is chosen (Yechiam, Rakow, & Newell, 2015). The higher a participant's risk preference, the riskier a participant is. Under some conditions decisions from description show higher risk preference to decisions from experience, and, in others, the reverse occurs.

Various factors affect what the DE gap looks like¹. As a framework for understanding how decisions from description differ to decisions from experience we are going to use the fourfold pattern of risk preference (Scholten & Read, 2014). The fourfold pattern of risk

¹ While other factors influence the DE gap, from decisions made under time pressure (Madan, Spetch, & Ludvig, 2015), to making decisions for other people (Martin, Gonzalez, Juvina, & Lebiere, 2014) this thesis will focus on the effect of probability and reinforcer size.

preference describes how people tend to behave when they make decisions from description, and how their decisions are influenced by both the chances of an event occurring and whether choices are gains or losses. Table 1 shows this fourfold pattern of risk preference from Scholten and Read (2014) combined with a table from Hertwig (2012) that shows the same situations but how people behave when making decisions from experience.

Table 1: Typical behaviour when people make decisions from description or decisions from experience for both gains and losses when the probability is either high or low. When the probability is high consider the example: would you prefer to win (lose) \$9499 or a 95% chance of to win (loss) \$10,000. When the probability is low consider: would you prefer to win (lose) \$501 or a 5% chance to win (lose) \$10,000. Where risk averse is defined as choosing the sure thing and risk seeking is taking the gamble.

Probability	Gain		Loss	
	decisions from description	decisions from experience	decisions from description	decisions from experience
High	<i>Risk Averse</i>	Risk Seeking	Risk Seeking	<i>Risk Averse</i>
Low	Risk Seeking	<i>Risk Averse</i>	<i>Risk Averse</i>	Risk Seeking

In decisions from experience, Expected Utility theory claimed that when faced with a gamble people tend to be risk averse. However, years later prospect theory showed that people can be risk averse in some choice contexts, and risk seeking in other choice contexts (Kahneman & Tversky, 1979a). Table 1 shows that when people make decisions from description in gain trials and when the probability of an event occurring is high, they are generally risk averse, yet when the probability is low people tend to be risk seeking (Jessup, Bishara, & Busemeyer, 2008). However, the opposite occurs when people make decisions from experience, when the probability is high they are generally risk seeking yet, become risk averse when the probability is low (Hertwig et al., 2004).

For a more concrete example of how decisions from description and decisions from experience can differ, consider the famous Asian Disease problem posed by Kahneman and

Tversky (1979). In this study participants are told a deadly disease is going to break out. To fight the disease participants can choose between two courses of action. In the gain condition they are told if they choose option A, 200 people will be saved. Whereas choosing option B will mean there is a $\frac{1}{3}$ chance 600 people will be saved and a $\frac{2}{3}$ chance that no one will be saved. Those in the loss condition were given the same problem but instead asked to choose between option A: where 400 people will die, and option B, which has a $\frac{1}{3}$ chance no one will die, and $\frac{2}{3}$ chance 600 people will die. Both the gain and loss condition have identical average outcomes. In both cases option A provides a certainty that 200 people will live, and option B provides a gamble that either all or none will survive. The only difference between the conditions is whether the two alternatives are framed as either a gain or loss. People who make decisions from description tend to be risk averse in the gain condition (they tend to choose to save 200 people) and risk seeking in the loss condition (they tend to choose a $\frac{1}{3}$ chance no one will die). Yet, you find the opposite pattern when people make decisions from experience. For example, Vallée-Tourangeau, Vallée-Tourangeau and Ramasubramanian (2016) had two decks of cards (one deck for option A the other for option B). By picking up a card from either deck, participants could see the outcome of their choice. Sometime the outcome was presented as a gain and sometimes as a loss. For example, choosing a card from one deck always resulted in 200 people being saved or 400 people dying. Whereas if you choose a card from the second deck, sometimes all would be saved and other times no one will be saved. In contrast to participants who made decisions from description in Tversky and Kahneman's study, participants tended to choose the risky option in the gain condition and the certain option in the loss condition.

While Table 1 shows decisions from description and decisions from experience differ depending on whether the probability is high or low, the literature is less clear on how decisions from description and decisions from experience differ when an event is equally

likely to occur or not, that is a 50% chance. There is reason to believe decisions from experience would not differ from decisions from description when the probability of an event occurring is equally likely to the probability that it will not occur. When people make decisions from description we have a good understanding of how people make decisions when a risky event is equally likely to occur or not occur: they tend to be risk averse. To illustrate, if I offered you the choice between \$20 or I tossed a coin, heads I give you \$40, tails \$0, what would you choose? If you are like most people, you would be risk averse and take the \$20 (Erev et al., 2010). But how would people behave if they were making decisions from experience? While few studies address this question directly, the little research we have suggests this may be one area where decisions from description and decisions from experience are alike. However, as we will see further on, how people behave when they make decisions from experience might depend on the context that they make that choice.

In the studies we have looked at thus far, you could argue that the DE gap might be because when people make decisions from experience they make repeated choices whereas those who make decisions from description often make a one-off choice. We know that repeated decisions can influence decision making (Chu & Chu, 1990; Koehler, 1996). However, when you control for repeated choices by having participants make repeated decisions from description a DE gap remains (Hertwig et al., 2004). This suggests it is not repeated choices that drive the DE gap rather the gap is more likely due to having direct experience of the outcomes. However, what happens when people have access to both description and experience information?

When people have both Description and Experience

“No one ever did anything because of a number” – Daniel Kahneman (Lewis, 2017)

When people make decisions when they have access to both information from description and experience, experience tends to have greater influence over their choices. For example, Jessup, Bishara and Busemeyer (2008) compared decisions between a description condition and a hybrid description-and-experience condition. In the description condition, participants made repeated decisions where they were told the probability of winning (e.g. *would you prefer 3 cents for sure or an 80% chance of 4 cents?*). The hybrid condition required participants to make exactly the same decisions as the description condition, but they received feedback after each choice. As expected, there was a significant difference between the conditions. Those in the hybrid condition tended to choose the risky option, an 80% chance of 4 cents, whereas those in the description condition were more likely to choose 3 cents for sure. According to the authors, this suggests when we have both description and experience information, experience has more influence over our decisions. From this study though, it is hard to tease apart whether descriptions or experience had more influence because to do so you would need to compare decisions in the hybrid condition to an experience condition as well as a description condition.

Lejarraga and Gonzalez (2011) used the same design as Jessup, Bishara and Busemeyer (2008) but added an experience condition. They found no difference between the hybrid description-experience conditions and the experience condition, whereas the hybrid condition was significantly different to the description condition. This study provides stronger support for the idea that when we have access to both description and experience information, experience is dominant while description information is mostly ignored. This also helps to explain why it is so hard to communicate risk to the general public, for example, providing people with probabilities of the chances of a natural disaster and the need to be prepared often falls on deaf ears, because people value their personal experience above descriptions of risks. In this thesis we looked at how risk preference changed over the course

of the experiment to better understand how decisions are influenced as people gain more experience with the outcomes.

What mechanism is behind decisions from experience? A dive into recency effects and the Extreme-Outcome Rule

To better understand the DE gap, we need to understand how we make decisions from experience. Two of the candidates are a recency effect and the Extreme-Outcome Rule; either one, or both, may help explain decisions from experience.

Recency Effect

According to the recency effect, outcomes experienced more recently have greater influence over our next choice than outcomes experienced earlier (Gilovich, Thomas, Vallone, Robert, Tversky, 1985). For example, if you are deciding whether or not to buy a lotto ticket, the fact you won last week influences your decision more than the fact you lost a month ago. On the surface, it makes sense a recency effect would influence decisions from experience. When you make a decisions from experience, you are relying on your memory and we know people have better memories for more recent events (Brodie & Murdock, 1977). Furthermore, the gamblers fallacy suggests after a few wins we tend to think we are on a lucky streak and we keep gambling (Gilovich, Vallone, Tversky, 1985). The idea is that a recency effect may be driving decisions from experience, but not decisions from description (perhaps not surprising as feedback is generally not given to description conditions).

However, the evidence a recency effect is driving decisions from experience is mixed. On one hand, several studies have found a recency effect when participants made decisions from experience. For example, Hertwig et al (2004) found more recently sampled outcomes better-predicted participants' decisions compared to earlier choices. Moreover, recency effects have also been found using both within and between subject designs (Camilleri & Newell, 2011; Barron & Barron, 2014). While these studies focused on rare events, Madan et

al (2014) found a recency effect when there were no rare events. In their task, participants clicked on doors, some of which gave them a guaranteed number of points (e.g. 20 points) and others produced risky outcomes (e.g. 50% chance of 40 or 0 points). They found that people adopted a win-stay/lose-shift strategy (a type of recency effect). In other words, if people chose the risky option and won they were more likely to choose the risky option on the next trial, whereas when they lost, they generally chose the other option in the following trial. However, the DE gap has also been found without any recency effects (Hau, Pleskac & Hertwig, 2010). These mixed findings suggest that while a recency effect may contribute to the DE gap it is unlikely to be the main driver. As this is an inconsistent finding with no obvious procedural difference between the studies that have and have not found recency effects we will look for a recency effect in our study.

The Extreme-Outcome Rule

While a recency effect may not be the main cause of the DE gap, Ludvig et al (2011) argue another memory bias, the Extreme-Outcome Rule, may drive the DE gap. The Extreme-Outcome Rule was inspired by the peak-end rule, which says when we think about a past event we have a tendency to remember two parts of the experience in particular, a) the most emotionally intense part and b) how it ended (Kahneman et al., 1993). Ludvig and Spetch (2011) argued if people have a bias to remember the most emotionally intense moments of an experience, this could influence not only how people remember an event but also the decisions they make. Therefore, the Extreme-Outcome Rule says when we use our past experience to make decisions we are more sensitive to the best and the worst outcome we have experienced (the extreme outcomes).

To investigate the Extreme-Outcome rule, Ludvig and Spetch (2011) had participants play a computer game where they could either win (gain trials) or lose (loss trials) points. When people made decisions from description they were asked *would you prefer 20 points*

(or -20 points) or a 50% chance of 40 points else 0 (or -40 points or 0). Whereas in decisions from experience trials participants were shown pictures of doors, each door was associated with a different number of points, by clicking on each door they received feedback on how many points they received. They found when participants made decisions from description they were more risk seeking for losses compared to gains. However, the opposite occurred when people made decisions from experience, people were more risk seeking for gains and risk averse for losses. How does the Extreme-Outcome Rule make sense of this finding? To answer this question, we first need to identify the extreme-outcomes. In their study the maximum number of points participants could win on any single trial was 40 points (biggest win). The other extreme-outcome (the biggest loss) was -40 points. In decisions from experience trials when faced with a choice between 20 points or a 50% chance of 0 or 40 points, Ludvig et al (2011) argued participants were more sensitive to the 40 points (best outcome) and thus were more likely to choose to gamble. Whereas in loss trials people are more sensitive to the -40 (worst outcome) and as a result tend to choose the sure option of -20 points. In other words, this study supports the idea for the Extreme-Outcome Rule, as people were more likely to choose to gamble when a risky option could lead to the biggest win, whereas if the risky option could lead to the biggest loss the sure thing was preferred.

In addition, people also have a better memory for the extreme outcomes compared to the non-extreme outcomes. Madan et al (2014) repeated Ludvig et al.'s (2011) study, but after the task gave participants a series of memory tests. First, participants were shown the four doors used in the study and asked to report the first number that came to mind. Secondly, they were shown each outcome (e.g. 40 points) and asked how often each outcome had occurred. Not only were participants more likely to recall an extreme outcome in the first memory test but they also reported that the extreme outcomes occurred more often than they did compared to the non-extreme outcomes. Moreover, participants' memory for the extreme-

outcome was correlated with their risky decisions. That is, participants who reported the extreme outcomes more in the memory test tended to gamble more compared to participants who did not show a bias for the extreme-outcomes in the memory test. Therefore, according to the Extreme-Outcome Rule people's bias to remember the extreme outcomes may drive the difference between decisions from experience and decisions from description, and hence the DE gap. Other studies have also found a similar pattern of results, that people have better memories for the extreme-outcomes and their memory is related to the choices they make (Madan & To, 2017; Ludvig, Madan, & Spetch, 2014; Madan et al., 2016).

Let's summarise the key points we have looked at so far.

1. People make different decisions when they make choices from description compared to experience (DE gap). However most research on decision making has focused on decisions from description, there is very little on decisions from experience.
2. When people have both description and experience information experience tends to be more influential on their choices.
3. One idea to explain how people make decisions from experience is the Extreme-Outcome rule, which says people give more weight to the best outcome and the worst outcome they have experienced when choosing between alternatives.

This thesis will build on our understanding of how we make decisions from experience. We will do this by testing the predictions the Extreme-Outcome rule makes. Throughout all studies we will also compare decisions from experience to a hybrid description and experience condition to investigate how having both description and experience information influences decision making when the risky alternative produces two

equally likely outcomes. To my knowledge, no study has investigated whether this hybrid condition differs from the Experience condition and whether response patterns are consistent with the predictions made by the Extreme-Outcome Rule.

All experiments were based off Ludvig et al.'s (2011) study. We choose this procedure as the basis for the experiments for a few reasons. Firstly, as research on the Extreme-Outcome Rule is in its infancy, using a similar design to Ludvig et al (2011) makes it easier to compare our results to other work on the Extreme-Outcome rule. Moreover, keeping the expected value between the safe and risky doors the same allows us to consider the decisions people are making without needing to take into account the 'optimal decision' because there is none. Also, as we have discussed, one of the main gaps in the decisions from experience literature is that we don't have much research on decisions from experience when the uncertain event has a 50% chance of occurring. Using the partial feedback procedure in particular allows us to understand how people's tolerance for risk changes over the course of the experiment.

Experiment 1

In Experiment 1 participants earned points by clicking on pictures of doors. There were two conditions, an Experience condition and a hybrid description and experience condition where participants got descriptive information (how many points they could win and the odds of winning) as well as experience information (after choosing an option they got feedback about how many points they received). While our study was similar to Ludvig et al (2011) we extended it in a few ways. Firstly, instead of a pure description condition we had a condition that had both description and experience information. As we have previously discussed, when we have access to both description and experience information, experience tends to have greater influence on our choices compared to descriptive information. However, to my knowledge, no studies have investigated this when the event is equally likely, I would

hypothesize that participants' response patterns would be consistent with an Extreme-Outcome Rule in the Description-Experience condition because research suggests that when making decisions we rely more on our experiences than descriptions.

Secondly, according to the Extreme-Outcome rule we are more sensitive to the worst outcome we have experienced, this should be true even when that outcome is not a loss (as it was in the study by Ludvig et al., (2011)). Therefore Experiment 1 had no losses, and the worst outcome in our study was zero points. If we find support for the Extreme-Outcome Rule we should still find people are more sensitive to the worst outcome even when it is not a loss. Ludvig et al (2011) had two types of trials they were interested in, gain trials (20 points or 50% chance of 40 or 0 points) and loss trials (-20 points or 50% chance of -40 or 0 points). Whereas the two trials we are interested in for Experiment 1 are when the reinforcer is small (20 points or 50% chance of 40 points or 0) and when the reinforcer is large (60 points or a 50% chance of 80 points or 40). If we find support for the Extreme-Outcome Rule people will be more risk seeking when the reinforcer is large compared to when it is small, as when the reinforcer is large the gamble can lead to the best outcome (80 points), but when the reinforcer is small the gamble can lead to the worst outcome (0 points).

To summarise, Experiment 1 has four research questions.

- 1) Do we get an Extreme-Outcome rule in the Experience condition? I.e. more risk seeking when the reinforcer is large compared to when it's small.
- 2) Is there a difference between the Experience and Description-Experience condition?
- 3) Does the Extreme-Outcome rule predict how people's choices in the description-and-experience condition.

- 4) Do we find a recency effect in both the Experience and Description-Experience conditions?

Method

Design: Experiment 1 was a mixed 2 (Source of information: Experience or Description-Experience) x 2 (reinforcer size: small, large). Source of information was the between-subjects factor and reinforcer size the within-subjects factor.

Participants: 156 third-year Psychology students completed the experiment during one of their tutorials. This study was approved by the School of Psychology Human Ethics Committee.

Materials: Participants completed the experiment on a Dell computer, the screen dimensions were 1920 x 1080. The task was programmed on Visual Studio 2017.

Procedure: The study was completed in a room with up to 16 students during their scheduled tutorial time. Participants were seated individually at desk with a computer with screens on either side to ensure privacy. They read a consent form on the computer, then if they agreed to complete the experiment the task began. Participants also had the option to complete the task as part of their lab work without their data being recorded. After completing the consent procedure, students, they were shown instructions on the screen which told them they would be shown pictures of doors, either one or two at a time. By clicking on a door they could earn points - each door was a different colour (the colours used were red, green, white and purple) and was associated with a different number of points. Participants were asked to try to earn as many points as possible. There were four doors in total. Below are the points associated with each door and the chances of each outcome occurring.

1. 100% chance of 20 points (20)
2. 50% chance of 40 points or 0 points (40/0)
3. 100% chance of 60 points (60)

4. 50% chance of 80 points or 40 points (80/40)

As shown by Table 1 there were three types of trials. In decision trials participants chose between two doors with the same expected value. Forced-choice trials were where participants were presented with only one door, these trials ensured that participants experienced all the different contingencies associated with each door. Catch trials involved decisions where one door has a greater expected value than the other (0/40 vs 40/80). Catch trials also allowed us to determine whether participants in the experience condition had learnt what the colour of doors indicated about the outcomes; participants who did not meet an accuracy criterion on the catch trials were excluded (see below). Each set of 24 trials contained 12 decision trials, 8 catch trials and 4 forced trials presented in random order.

Table 2: Trial Types for Experiment 1. LEO indicates the alternative that includes the lowest extreme outcome and HEO indicates the alternative that includes the highest extreme outcome.

Left Door	Right Door				
	20	40/0(LEO)	60	80(HEO)/40	None
20		Decision	Catch	Catch	Forced
40/0(LEO)	Decision		Catch	Catch	Forced
60	Catch	Catch		Decision	Forced
80(HEO)/40	Catch	Catch	Decision		Forced
None	Forced	Forced	Forced	Forced

Figure 2. Screen shot of the difference between the sources of information conditions in Experiment 1. The Description-Experience condition is on the left and the Experience on the right

Participants were either in the Description-Experience condition or the Experience condition. As shown by Figure 2, those in the Description-Experience condition had the number of points they could win and the probability of winning the points written below the doors (e.g. 50% chance to win 40 points and 50% chance to win 80 points). Whereas those in the Experience condition did not have any written information provided. After participants decided which door to click feedback on how many points they won was presented for 1.2 seconds, during this time if a participant won points, a picture of a pot of gold with the number of points won was displayed. If no points were won then a black square with “0” was displayed. The total number of points that a participant had won throughout the task was shown in the top centre of the screen throughout the course of the experiment. The Experiment lasted 12 minutes.

Participants completed six blocks of trials, a block was two minutes. Each block contains approximately 12 decision trials, 8 catch trials and 4 forced trials. Only the choices in the decisions trials were used in the analysis. In the Description-Experience condition participants completed on average 153.17 trials (max = 193, min = 111). Those in the Experience condition completed on average 188.46 trials (max = 212, min = 162). At the end of the study, participants were asked whether they had read the article by Ludvig et al. (2011) which was on their course reading list for a later point in the semester. Participants who indicated that they had were excluded as they might no longer be naïve to the purpose of the experiment.

Exclusions

Nine participants were excluded from the analysis. There were three reasons participants data were not analysed. Firstly, if a participant in either condition scored less than 60% in the catch trials (trials where one outcome has a higher expected value than the

other e.g. choice between 20 points or 40 points) across the whole experiment they were excluded from the analysis. This exclusion criteria comes from similar studies (Christopher, Madan & To, 2017; Ludvig, Madan, Pisklak, et al., 2014; Madan, Ludvig, & Spetch, 2016). Three from the Experience condition (4.11 %) were excluded for this reason. Secondly, participants were also left out of the analysis if they reported they had read the Ludvig et al (2011) article or if they had discussed the experiment with anyone else. For this reason three from the Description-Experience condition (3.3%) and one from the Experience condition (1.34%) were excluded. The final dataset contained 69 participants in the Experience condition and 87 participants in the Description-Experience condition.

Results

Data Analysis

Our dependent variable was risk preference, which is a measure of how often a participant choose the risky option. Risk preferences was calculated as the percentage of times the risky option was chosen, thus the higher a participant's risk preference score the more often they choose the risk option. Each participant had two risk preference scores, one when the reinforcer was small, the other when the reinforcer was large.

Risk preference was significantly non-normal, $W = 0.10$, $p < .001$ This non-normality as well as the unequal group size will bias the F-ratio in an ANOVA (Field, Jeremy, Field, 2012). Therefore, we analysed the results using a robust mixed ANOVA (Wilcox, 2012). The robust ANOVA we used combines a bootstrapping procedure and uses trimmed means to estimate key statistics, in this case the mean. Bootstrapping is based on the idea of estimating the shape of the sampling distribution. The first step is to take a bootstrap sample which is the same size as the sample data by taking a single value from the sample and replacing it, and then drawing another sample and so on. To illustrate, if you are measuring depression on a 7

point scale and you have 5 scores: 1, 2,3,4,5. Each bootstrap sample taken will also have five scores. The first sample might randomly select 3 from the sample, as the sample with replacement method is used, the 3 is returned to the sample and available for the next selection. Therefore, you could end up with a bootstrap sample that has repeats, for example, 1,2,2,3,4, not that 2 has been selected twice whereas 5 has not been selected at all. Once the bootstrap sample is collected, the statistic of interested is calculated (e.g. the mean), this number is stored, then the process repeats and another bootstrap sample is calculated. This process is repeated a large number of times (in our case 2000) which is then used to estimate the sampling distribution. Field and Wilcox (2017) argued that a bootstrapping procedure is the best way to deal with non-normal distributions because the p -value is empirically derived from the data. We set the number of bootstrapping samples to 2000 as Field (2013) suggested that this was more than enough to estimate the distribution). The bootstrapping procedure only provides a p value, rather than F-statistic.

A trimmed mean is a mean that is calculated after a certain percentage of data has been removed from both the extreme ends of the distribution to reduce skew (Field & Wilcox, 2017). While other manual trims are used to remove outliers (e.g. based on standard deviations) percentage-based trims provide good estimates of relevant statistic such as the mean (Wilcox, 2010). While some robust tests use either a bootstrapping procedure or trimmed means, we choose to use both because Wilcox (2012) recommended using a bootstrapping procedure as well as trimmed means if trimming less than 20% of the data. We trimmed 5 % as this is a commonly used (Bishop & Hsu, 2015).

Therefore, in Experiment 1 we analysed the data using a 2 (source of information: Experience, Description-Experience) x2 (reinforcer size: small, large) robust mixed ANOVA. We set the number of bootstrapping repetitions to 2000, used an alpha level of 0.05 and used

a trimmed mean of 5%. We ran the analysis using the R function `bwtrimbt` (Wilcox, 2012).

In the between-subjects factor of source of information condition there were two levels (description and description-experience) and two levels of the within-subjects factor of reinforcer size (small and large). Results were analysed using the statistical software package R, version 3.4.1 (R Core Team, 2017)

I analysed the data from the final two blocks of the experiment (last four minutes)². This gave participants in the experience condition time to learn the payoff distributions. I specifically looked at the last two blocks because other studies have also looked at the final two blocks, this makes it easier to compare our findings to theirs, crucially, among those studies is Ludvig et al (2011).

Effects of Source of Information and Reinforcer Magnitude on Risk Preference

There was no main effect of source of information condition ($p = 0.321$), there was a main effect of reinforcer size ($p < 0.01$) these main effect need to be interpreted in light of a significant interaction between source of information and reinforcer size ($p = 0.004$). Follow up tests using the R function `mcbb20` reveal that in the Experience condition on average, participants were more risk seeking when the reinforcer was large ($M = 54.72$, $SD = 31.93$) compared to when the reinforcer was small ($M = 34.53$, $SD = 27.94$). This pattern of results is consistent with an Extreme-Outcome rule in the Experience condition. This is also seen in Figure 3.

As we have seen, participants had higher levels of risk preference when the reinforcer was large compared to when it was small (evidence of an Extreme-Outcome rule), but only in

² While Ludvig et al (2011) used the last two last blocks for their analysis as we have done here, subsequent studies have often analysed the last three blocks (Ludvig, Spetch, 2017). Carrying out the current analyses using the last three blocks did not influence the pattern of results.

the Experience condition. Does this pattern hold when we look at the individual data?

Figures 4 and 5 suggest the answer is yes by showing us how consistent the effect is across individuals. Looking at Figure 4 we see in the Experience condition the majority of participants (72.46%) are above the line. That is, participants tended to be more risk seeking when the reinforcer was large compared to when it was small. This is consistent with the Extreme-Outcome rule in the Experience condition.

Figure 3: Risk preference in both the Experience and Description-Experience condition for smaller and larger reinforcers during the final two blocks of the experiment. Higher bars on the y axis represent a greater tendency to choose the risky option. Error bars are the standard error of the mean

Unlike the Experience condition, in the Description-Experience condition, risk preference was similar whether the reinforcer was small ($M = 39.59$, $SD = 36.16$) or large (M

$= 42.02$ $SD = 33.61$). Therefore, this pattern of results is not consistent with an Extreme-Outcome rule in the Description-Experience condition.³

Figure 4: Modified Brinley plot showing each participants average risk preference for both small and large reinforcers in the Experience conditions. Each dot in the graph represents a single participant. The black diagonal line represents matching risk preference for smaller and larger reinforcer size. For example, if a participant's average risk preference for the smaller reinforcer was 40% then their average risk preference is also 40% for the larger reinforcer. The dots above the diagonal line are participants who have higher risk preference when the reinforcer was large.

Figure 5: Modified Brinley plot showing each participant's average risk preference for both small and large reinforcers in the Description-Experience condition. Each dot in the graph represents a single participant. The black diagonal line represents identical risk preference for smaller and larger reinforcer size. For example, if a participant's average risk preference for the smaller reinforcer was 40% then their average risk preference is also 40% for the larger reinforcer. The dots above the diagonal line are participants who have higher risk preference when the reinforcer was large.

³ The default level for trimming means is 20%, as we didn't want to get rid of that much data we used 5% trimmed means. However, even when a 20% trimmed mean is used the pattern of results remains the same.

Figure 5 shows the individual data, we can see that in the Description-Experience condition there is a fairly even divide between people who are above and below the reference line (44.83% of participants are above the line). This suggest the size of the reinforcer had no systematic influence over individuals risk preference. Overall, the results from the ANOVA and the individual data suggest that the response pattern in the Description-Experience condition was not consistent with the Extreme-Outcome rule (Research Question 3), and, therefore that the pattern of risk preference was different in the Experience and Description-Experience-conditions (Research Question 2)

Recency Effect

To find out whether participants decisions were influenced by the outcome of the previous gamble we calculated recency scores. Also note, that we categorized each score as belonging to either the small, medium or large reinforcer based on the current trial. For example, if a participant gambled when the current trial was a small reinforcer, and they had just won a gamble on the previous trial when the reinforcer was large, this would be included in the small recency effect score. Therefore, we calculated three recency scores for each participant, one each for when the reinforcer was small and large. To get the recency effect score we first calculated the probability of each participant choosing the risky option if they won a gamble on the previous trial. This was done using the equation below:

$$\begin{aligned}
 & \text{probability of gambling after a win} \\
 = & \frac{\text{number of times they chose to gamble in a decision trial} \\
 & \text{when won a gamble on the previous trial}}{\text{number of decision trials won} \\
 & \text{a gamble on the previous trial}}
 \end{aligned}$$

This probability ranges from 0 = never gamble after a win, to 1 = always gamble after a win. We then calculated the probability of each participant choosing the risky option if they lost a gamble on the previous trial using the equation below

$$\text{probability of gambling after a loss} = \frac{\text{number of times they chose to gamble in a decision trial when lost a gamble on the previous trial}}{\text{number of decision trials lost a gamble on the previous trial}}$$

This probability ranges from 0 = never gamble after a loss, to 1 = always gamble after a loss. Using these two numbers allows us to look at the individual data for each participant

and Figure 6 presents the probability of choosing the risky option on trials immediately after a win plotted against the probability of choosing the risky option immediately after a by the reinforcer size and source of information condition. We see a clear pattern sin both the Description-Experience and Experience condition and regardless of reinforcer size. Participants were more risk seeking after they had won a gamble on the previous trial compared to if they lost the previous gamble. This is reflected in the fact that the majority of the participants' data points are above the reference line for both **Error!** **Reference source not found.** and Figure 6.

Figure 5: Modified Brinley plot showing participants in the Description- Experience condition's probability of choosing the risky option if they had gambled on the previous trial and won compared to the probability of them choosing the risky option if they had gambled on the previous trial and lost. Each dot represents a single participant. The back-diagonal line is a reference line, dots above the line show participants who were more likely to gamble if they won in the previous gamble compared if they lost.

Figure 6: Modified Brinley plot showing participants in the Experience condition's probability of choosing the risky option if they had gambled on the previous trial and won compared to the probability of them choosing the risky option they had gambled on the previous trial and lost. Each dot represents a single participant. The back-diagonal line is a

reference line, the dots above the line show participants who were more likely to gamble if they won in the previous gamble compared if they lost.

Figure 7: Recency Effect score for both conditions depending on the size of the reinforcer. The Y axis shows the recency effect score (reversed to make it easier to interpret). Higher negative numbers represent participants were more likely to gamble after winning on the previous trial compared to if they lost in the previous trial.

Using the formula below, we calculated the recency effect score

Recency Effect Score

$$\log \frac{(\text{Probability of choosing the gamble if lost on the previous trial})}{\text{Probability of choosing the gamble if won on the previous trial}).$$

Where either probability was zero we added 0.1 to both probabilities. This meant in the Experience condition we added 64 0.1's when the reinforcer was small and 55 when the reinforcer was large. In the Description-Experience condition we added 100 0.1's when the reinforcer was small and 88 when the reinforcer was large.

We ran a 2 (source of information: Description-Experience, Experience) x 2 (reinforcer size: small, large) robust ANOVA on the recency scores using the R function `bwtrimbt` (Wilcox, 2012). There was no main effect of source of information, ($p < 0.461$) and no main effect of reinforcer size ($p < 0.763$). There was a significant interaction between source of information and reinforcer size ($p < 0.01$). We investigated this interaction using the R function `mcpp20` (Wilcox, 2012). This tests revealed that in the Experience condition, participants were more likely to show a win-stay/lose-shift pattern when the reinforcer was small compared to when it was large. Whereas, the opposite was found in the Description-Experience condition, here participants were more likely to show a win-stay/lose-shift pattern when the reinforcer was large compared to when it was small.

Risk preference across the session

As part of an exploratory analyses we wanted to understand how risk preference changed as the experiment progressed. To answer this question, we calculated individuals' risk preference in each of the six blocks, and then averaged the scores for each block. We calculated the average risk preference for Block 1, then Block 2 and so on. We then ran a two 2 x 6 (reinforcer size x blocks) robust mixed ANOVA, one for each condition. In the Description-Experience condition there was no main effect of reinforcer size ($p = 0.467$), no main effect of block ($p = 0.432$) and no interaction ($p = 0.109$). In other words, risk preference did not change over the course of the experiment regardless of whether the reinforcer was small or large. We can see this visually in Figure 8 as risk preference remained at similar levels throughout the course of the experiment. However, the Experience condition tells a different story. In this case we see a main effect of reinforcer size ($p < 0.01$) and no main effect of block ($p = 0.326$) which need to be interpreted in light of a significant interaction ($p < 0.01$). Indeed, when we look at Figure 7 we see that in the Experience condition risk preference changes as the experiment went on. As such, when the reinforcer

was large, risk preference increased as the experiment progressed, whereas when the reinforcer was small risk preference decreased the further participants were in the experiment. This difference between the two reinforcer sizes grows over the third and fourth Block, after which the risk preference plateaus, for both the small and large reinforcers, for remainder of the experiment.

Figure 7: Risk preference over the course of Experiment 1 in the Experience condition. Higher lines on the Y axis represents a greater tendency to choose the risky option. The black line shows the risk preference when the reinforcer was small, the grey line shows the risk preference when the reinforcer was large.

Figure 8: Risk preference over the course of Experiment 1 in the Description-Experience condition. Higher lines on the Y axis represents a greater tendency to choose the risky option. The black line shows risk preference when the reinforcer was small, the grey line shows risk preference when the reinforcer was large.

Discussion

Experiment 1 replicated Ludvig et al's (2011) finding as we found support for the Extreme-Outcome rule in the Experience condition. Here, participants were more risk seeking when the gamble could lead to the biggest win (80 points) compared to when the gamble could lead to the worst outcome (0 points). This finding is also consistent with other research that has found evidence of an Extreme-Outcome rule. Madan, Ludvig and Spetch (2014) for example, used a similar task to ours where participants had to click doors to earn points. However, while the extreme outcomes in our study were 0 and 80, their extreme outcomes extended into negatives, from -40 to +40. Yet, as we found, participants were more risk seeking when faced with a gamble that could lead to the best outcome (+40) compared to when the gamble could lead to the worst outcome (-40). Therefore, our results suggest you can still get an Extreme-Outcome rule even the worst outcome is not a loss. This is consistent with other studies that

find an Extreme-Outcome Rule can emerge in situations involving only gains (Ludvig, Madan, & Spetch, 2014b).

Possible alternative explanations to the Extreme-Outcome Rule

The Extreme-Outcome rule is not the only explanation for the fact that participants were more risk seeking when the reinforcer was large compared to when it was small in the Experience condition. Rather, this pattern might have occurred because only the small reinforcer condition included a possible outcome of zero points; there are two explanations as to why having a zero as an outcome could be an alternative explanation.

1. Participants were zero averse

Shampanier, Mazar and Ariely (2007) argue that zeroes can act as ‘special’ numbers which people treat differently than positive and negative numbers. In our study for example, rather than people being more sensitive to the worst outcome (zero) because it was an extreme value in the context of Experiment 1, perhaps participants were more sensitive to it because it was zero. To illustrate, when the reinforcer was small and participants lost the gamble, they earned zero points. Whereas when the reinforcer is large, even if they lost the gamble they still won points. Therefore, people may be more likely to choose the risky option when the reinforcer is large because they can gamble knowing if they lose they still get points. So, something is better than nothing, and nothing is aversive.

2. Delay discounting makes the zero less appealing

We know people tend to value a reinforcer more when they get it for certain compared to getting the same reinforcer probabilistically (McKerchar & Renda, 2012). For example, we value \$100 more than a 50% chance of \$100. Green, Myerson and Calvert (2010) found that when a reinforcer has a probability of occurring less than 1.0 (i.e. sometimes we get points, sometimes we don’t) the decrease in value of that reinforcer can be equivalent to experiencing a delay. To relate this to Experiment 1, when a participant chose

the risky option when the reinforcer was small and lost, they experienced a delay until they next earned points (following a later trial). Whereas if they chose the risky option when the reinforcer was large and lost there was no delay in obtaining points because they still earned some. In fact, in all four possible outcomes, the only time a participant has to wait after a response to receive points is if they lose when they gamble, and the reinforcer was small. Therefore, it is possible that when the reinforcer was small, the risky option lost value because of that delay to get points (have to wait until the next trial). Thus making the safe option more appealing. Therefore, participants might have been more likely to choose the safe option in the small reinforcer condition not so much because they were avoiding the worst outcome but because it was the only alternative where the risky choice signalled a possible delay to the next reinforcer.

The Description-Experience Condition

We did not find evidence of an Extreme-Outcome rule in the Description-Experience condition. In fact, the size of the reinforcer had no influence on risk preference. This was surprising as research suggests when we make decisions when we have both description and experience, experience tends to be more dominant, and in some cases descriptions are ignored (Jessup et al., 2008; Lejarraga & Gonzalez, 2011; Lejarraga, Woike, & Hertwig, 2016). Therefore, we expected no difference between the two conditions. What might explain this pattern of results?

Recall from the introduction that we mentioned how some scholars argue a recency effect is sometimes found when people make decisions from experience (Barron & Barron, 2014; Hertwig et al., 2004), although it is unlikely to be the main driver of the DE gap, perhaps it is driving the difference between the two conditions in Experiment 1. However, although we did find a strong recency effect in the Experience condition, a similarly strong effect was also

found in the Description-Experience condition. Of note, while both conditions showed a recency effect only the Experience condition showed a change over time in levels of risk preference. The lack of a change over time in the Description-Experience condition suggests that, in this study, participants used the description, not their experience to make their choices. We will return to reasons why this difference occurred in the general discussion.

The results of Experiment 1 pose multiple possible roads to go down in follow up experiments. While further investigating why the Description-Experience and Experience condition differed and testing the ideas mentioned above would be interesting, we choose to further explore the Extreme-Outcome rule as this one of the main focuses of this thesis.

Experiment 2

After the findings of Experiment 1, we used the same task but we made three main changes for Experiment 2. Firstly, we increased all points by five to remove the zero. Secondly, we added another pair of doors to change the extreme outcomes, and finally, we made the task a within-subjects design. All changes are explained below.

Method

Within Subject Design

In Experiment 1 participants experienced either the Description-Experience or Experience condition. However, a DE gap emerges even when the same participant completes both the Description-Experience and Experience conditions in the same task (Ludvig, Madan, & Spetch, 2014a; Ludvig & Spetch, 2011; Madan et al., 2014). Using a within-subjects design for Experiment 2 allowed us to increase the power of the experiment meaning that fewer participants were required. The task was 24 minutes long, participants spent the first 12 minutes completing either the Description-Experience or Experience

condition, immediately after they then completed the opposite condition. The order conditions were completed was counterbalanced across participants. Although the points that could be won and the probability of winning did not change over the course of the whole experiment, when the experiment entered the second 12 minutes and thus the condition changed, we replaced the six doors that were used in the first condition with a new set of six differently coloured doors for the second. We did this for two reasons. Firstly, because the points to be won and the probability of winning them were the same as the first half we wanted to reduce the chances the first condition would influence decisions in the second. For example, if you completed the first half of the task in the Description-Experience condition when you entered the Experience condition and the labels were removed you might carry over or simply remember the probabilities from the labels. Secondly, changing the doors would clearly indicate to the participant that the condition had changed.

Removing the zero

In Experiment 1 we found support for the Extreme-Outcome rule in the Experience condition in that participants were more risk seeking when the gamble could lead to the biggest win (80 points) compared. Why? One idea is that, rather than participants being more sensitive to the extreme outcome, participants were more risk seeking when the reinforcer was large because if they gambled and lost they still earned some points (40 points). Whereas if they gambled when the reinforcer was small and lost they got no points. So perhaps participants were more likely to gamble when the reinforcer was large because winning some points was better than no points. Moreover, zero is neither an absolute gain nor loss and therefore zero may be treated differently to positive and negative numbers (Shampanier et al., 2007). Therefore, in Experiment 2 we added five points to all outcomes (e.g. 0 became 5, 20 became 25 and so on). Therefore, the worst outcome was not a zero anymore, rather it was a five. Adding five to every outcome as opposed to just the zero also kept the expected value of

the two alternatives presented on each decision trial the same. If we find support for the Extreme-Outcome rule in Experiment 2 participants should again be more risk seeking when the reinforcer was large compared to when it is small. If we don't find this, it suggests the presence of a zero was driving the effect in Experiment 1 rather than the best and worst possible outcome.

Adding extra decision trial

The two decision trials in Experiment 1 both contained an extreme outcome. Thus if we add a more extreme outcome – by adding a door that can give more than 80 points, if the Extreme-Outcome rule describes participants' response patterns, we should find that participants will be more sensitive to the new extreme-outcome rather than the previous extreme-outcomes. Therefore, in Experiment 2, we added a third set of doors. One of which had a 100% chance of 105 points and the risky door had a 50% chance of 125 points or 85 points (this pair of doors is hereafter referred to as 85/125). Hence, Experiment 2 had three sets of decision trials, 5/45, 45/85 and 85/125. If we find support for the Extreme- Outcome rule then we should see more risk seeking in the 85/125 decision trials compared to the 5/45 and 45/85 decision trials. And there should be a significant difference between the medium reinforcer and the small and large reinforcers (Ludvig, personal communication, 2018).

To summarise, there were six doors used in Experiment 2:

1. 100% chance of 25 points (25)	}	Small Reinforcer
2. 50% chance of 45 points or 5 points (45/5)		
3. 100% chance of 65 points (65)	}	Medium Reinforcer
4. 50% chance of 85 points or 45 points (85/45)		
5. 100% chance of 105 points (105)	}	Large Reinforcer
6. 50% chance of 125 or 85 points (125/85)		

Participants: 25 first year psychology students completed the task as part of a course requirement.

Design: Experiment 2 used a within subjects design. There were two independent variables: source of information (description-experience and experience), and the size of the reinforcer (small, medium and large). Risk preference was measured to examine how the source of information condition and the size of the reinforcer influenced participant's choice.

Procedure: For the first 12 minutes participants completed one source of information condition (Experience or Description-Experience) then in the second 12 minutes they completed the other condition. The order they completed the conditions was randomized. When the conditions changed in the second half of the task, the colours of the doors also changed – this was to avoid any learning that might carry over from the first half of the experiment. For example, in the first half a purple door had a 100% chance of 60 points, but in the second half a grey door gave a 100% chance of 60 points. Participants completed six blocks of trials, a single block was 2 minutes long. Each block contained approximately 24 decision trials, 24 catch trials and 8 forced trials. All other aspects of the procedure were as for Experiment 1. In the Description-Experience condition participants completed on average 149.40 trials (max = 207, min = 110). When participants completed the Experience condition, on average they completed on 179.84 trials (max = 218, min = 131).

Exclusions: One participant (4%) was excluded from the analysis due to failing the catch trials. This left 24 participants in the study.

Results

Data Analysis

To determine whether reinforcer size, source of information or an interaction of the two affected risk preferences we ran an ANOVA. Risk preference was significantly non-normal, $W = 0.11$, $p < 0.01$. Therefore, we used a robust 2 (source of information condition: Description-Experience, Experience) x 3 (reinforcer size: small, medium, large) ANOVA using the R function `wwtrimbt` (Wilcox, 2012), the number of bootstrap repetitions was set to 2000 with an alpha level of 0.05 and using a 5% trimmed mean. The data was analysed using the statistical software package R (R Core Team, 2017)

Figure 9: Mean risk preference by reinforcer size in both the Description-Experience and Experience conditions. Higher scores on the Y axis indicate that the risky option was chosen more often. Error bars are standard error of the mean.

There was a significant main effect of condition ($p = 0.019$), no main effect of reinforcer size ($p = 0.107$), there was a significant interaction between condition and

reinforcer size ($p = 0.024$).⁴ To further investigate this interaction, we ran two separate ANOVAs, one on each condition using the R function `t1waybt`, a trimmed mean of 5%, an alpha level of 5% and by setting the number of bootstraps to 2000 (unlike some of the other robust methods we have seen, `t1waybt` produced both an F statistic as well as a p value). In the Description-Experience condition there was no significant difference in risk preference regardless of the reinforcer size, $F_t = 0.07$, $p = 0.941$. However, in the Experience condition there was a significant difference in risk preference depending on the reinforcer size $F_t = 5.94$, $p < 0.01$. To investigate where this difference lay, we ran a series of multiple comparisons using the R function `mcppb20` (Wilcox, 2012). This function uses a bootstrapping method (set to 2000 repetitions), the results are presented in Table 3. There was a significant difference in risk preference between the small and large reinforcers and between the small and medium reinforcers in the Experience condition, but not between the medium and large.

Table 3: Post Hoc tests between groups in the Experience condition

Condition	Mean Difference Ψ	Lower 95% CI	Upper 95 % CI	P value ⁵
Experience				
Small vs Medium * ⁶	-21.15	-35.43	-5.88	0.020
Small vs Large *	-30.73	-45.30	-15.91	<0.01
Medium vs Large	-9.58	-26.23	6.26	0.315

⁴We also found no order effect, that is, the order participants completed each source of information condition did not influence the pattern of results for any analysis on the Experiment 2 data.

⁵Both the p values and confidence intervals are adjusted to control for the probability of at least 1 type one error.

⁶ * Indicates a significant difference between the groups

Although participants had higher levels of risk preference when the reinforcer was large ($M = 65.23$, $SD = 31.30$) compared to when it was small ($M = 37.12$, $SD = 27.40$) there was no significant difference between the medium and large reinforcers. This pattern of results is not consistent with the Extreme-Outcome rule. In addition, as found in Experiment 1, there was no Extreme-Outcome rule in the Description-Experience condition, as the size of the reinforcer had no influence on risk preference. Risk preference was similar regardless of whether it was small ($M = 64.97$, $SD = 32.03$), medium ($M = 64.71$, $SD = 29.55$), or large ($M = 61.80$, $SD = 30.53$).

We see the same pattern of results shown by the ANOVA when we look at the individual data. Figure 10 shows that for participants in the Description-Experience condition, the size of the reinforcer did not have a systematic influence over their risk preference, as participants are well scattered both above and below the reference line. The exception is when we compare the medium and large reinforcers, where instead participants tend to hug the reference line. In the Experience condition on the other hand, compared to the small reinforcer participants tend to be more risk seeking when the reinforcer was medium and large.

Description-Experience Condition

Experience Condition

Figure 10: Modified Brinley plots showing how risk preference was influenced by the size of the reinforcer, each dot represents a single participant. The higher the risk preference score the more often a participant choose the risky option. The black line running through each graph is a reference line with a slope of 1. Dots above the line show participants who were more risk seeking when the magnitude of the reinforcer was larger.

Recency Effect

To find out whether participants decisions were influenced by the outcome of the previous gamble we calculated the probability of choosing the risky option after both winning and losing a gamble on the previous trial. We did so using the same method from Experiment 1.

Using these two numbers allows us to look at the individual data. For each participant Figure 11 presents the probability of choosing the risky option on trials immediately after a win plotted against the probability of choosing the risky option immediately after a loss by the reinforcer size and source of information condition. In the Description-Experience condition participants were more risk seeking after they had won a gamble on the previous trial compared to if they lost the previous gamble. This is reflected in the fact that most of the participants' data points are above the reference line in all three of the Description-Experience graphs. This is also seen for the medium reinforcer in the Experience condition, whereas there is no systematic pattern when the reinforcer is small and large in the Experience condition.

Figure 11: Modified Brinley plots showing participants' tendency to choose the risky option after either winning or losing a gamble (recency effect) in both the Description-Experience and Experience condition and by the reinforcer size. The black line running diagonally through each graph is a reference line, each dot represents a single participant. Any dots above the line show participants who were more likely to choosing the risky option after winning a gamble compared to if they lost a gamble the last time they choose the risky option

Where either probability was zero we added 0.1 to both probabilities. Positive recency effect scores indicate that participants were more likely to choose the risky option risk after a loss while negative scores indicate that they were more likely to choose the risky option after a win (note that the axis in the figure below is reversed to make it visually easier to match to the individual data). Figure 12 shows that in the Description-Experience condition the probability of choosing the risky option was larger following a win than it was following a loss, in other words a win-stay/lose-shift pattern. We see this pattern to a lesser extent for the medium reinforcer in the Experience condition, yet not when the reinforcer was small or large. Instead, the recency effect here is close to zero (the bar for the small reinforcer is so tiny it is not visible).

Figure 12: Bar graph showing the mean recency score for each condition by the reinforcer size. A negative recency score indicates that a participant that was more likely to gamble if they won a gamble on the last trial compared to trials when they lost that gamble. Whereas, a positive recency score shows a participant who was more risk seeking when they lost a gamble on the previous trial compared to if they won it.

To compare how the condition and reinforcer sizes were influenced by the outcome of the previous gamble we used the recency score and ran a 2 (source of information: Description-Experience, Experience) $\times$ 3 (reinforcer size: small, medium, large) robust ANOVA using the R function `wwtrimbt` (Wilcox, 2012). There was no interaction between condition and reinforcer size ($p=0.705$), no main effect of reinforcer size ($p=0.701$), there was a significant main effect of condition ($p<0.01$). Meaning participants were more likely to choose the risky option after they won a gamble on the previous trial in the Description-Experience condition ($M = -0.20$, $SD = 0.90$) compared to the Experience condition ($M = 2.93$, $SD = 0.96$).

Risk preference across the session

How did Risk Preference change over time? To answer this question, we calculated each participant's risk preference for each block of the experiment, then we averaged all the scores for Block 1, then Block 2 etc. As Figure 13 shows, in the Experience condition the first few blocks show little difference in risk preference between the three reinforcer sizes. However, as the experiment progressed, the difference between the three reinforcer sizes increased. On the other hand, in the Description-Experience condition, risk preference is at a similar level

Figure 13: Line graph showing how risk preference changed over the course of Experiment 2 by the size of the reinforcer. The order participants completed the Description-Experience and Experience condition were collapsed, so although this graph shows the Description-Experience condition first followed by the Experience condition, some participants completed it in the opposite order. The black vertical line shows the point where the conditions changed. The first six blocks show the Description-Experience condition, the second show the Experience condition

To investigate whether the effect of reinforcer size on risk preference changed over time we ran a 2 (source of information: Description-Experience, Experience) x 3 (reinforcer size: small, medium, large) x 6 (blocks) robust ANOVA using the R function `wwwtrimbt` (Wilcox, 2012).

There was a main effect of source of information condition ($p < 0.01$), no main effect of reinforcer size ($p = 0.99$), and no main effect of block ($p = 0.944$). This main effect must be interpreted in light of a significant interaction between reinforcer and block ($p = 0.038$). No other two or three way interactions were significant. We investigated the interaction between reinforcer size and block using the R function `mcpp20` (Wilcox, 2012). In the Description-Experience condition only in Block 2 was there a difference in risk preference between the three reinforcer sizes. During this block, risk preference was higher for the small reinforcer compared to the medium and large reinforcers. In the other five blocks there was no difference in risk preference between the reinforcer sizes. In the Experience condition, risk preference stabilises by Block four as the same pattern emerges that we see in the final two blocks and continues to exist until the end of the condition. This suggests using the final two blocks to analyse the effect of reinforcer on risk preference was appropriate as by then risk preference was stable, suggesting participants had learnt the payoff distributions.

Discussion

In Experiment 2 we tested an assumption made by the Extreme-Outcome rule, that participants are more sensitive to both the best and worst outcome compared to less extreme outcomes. According to the Extreme-Outcome rule, risk preference would not only be higher when the reinforcer is large compared to when it is small, but risk preference for the medium reinforcer should be significantly different from both the small and large reinforcers. We also investigated whether the presence of zero might have provided an alternative explanation for the Extreme-Outcome effect found in the Experience condition of Experiment 1 by removing zero as a potential outcome. In Experiment 2 we found that in the Experience condition, risk preference was higher when the reinforcer was large compared to when it was small. This suggests the zero aversion, or the temporal discounting explanation does not provide an alternative explanation for the Extreme-Outcome effect found in Experiment 1. However, at the same time we did not find support for the Extreme-Outcome rule, as in the Experience

condition, there was no difference between the medium and large reinforcers. This indicates an effect of the low extreme outcome (worst outcome) but not the high extreme (best outcome), and therefore does not fully support the prediction made by the Extreme-Outcome rule. In other words, this suggests when we make decisions from experience we are more likely to avoid a gamble when it can lead to the worst outcome, but when the gamble can lead to the best outcome we are no more likely to gamble compared to a non-extreme outcome. Experiment 2 is consistent with the idea that participants were sensitive to the worst outcome, rather than showing zero aversion (Madan, Ludvig, Spetch, 2017; Ludvig, Madan, & Spetch, 2014).

The lack of effect for the best extreme outcome contrasts with Ludvig, Madan, McMillan and Spetch (2018 Experiment 2a). They used a similar task but they varied how close the second-best outcome was to the best outcome. In one condition, the best outcome was 75 points and the second best was 46, in the other the second best was 74. They found participants were more sensitive to the best outcome regardless of how close in points the second-best outcome was. However, Experiments 2 results suggest the line between the best outcome and second-best is more blurred. What explains this conflicting result? One procedural difference between our study and Ludvig et al (2018b) was their study had both gains and losses, whereas all the outcomes in our task were gains. However, Ludvig, Madan, McMillan and Spetch (2018 Experiment 2b), used a task with all gains like our Experiment 2, they found participants were more sensitive to the best outcomes regardless of how close in value the next option was. One possible reason for our not finding a pattern of risk preference consistent with the Extreme-Outcome rule for the best option may be a lack of power. As can be seen in Figure 11 we did have a non-significant difference in the same direction to that observed by Ludvig et al. We had 24 participants while Ludvig et al. had a minimum of 64 in their group. This pattern of results is also not what we would expect when we have both

description and experience information. Rather, as we have seen, when people have both, Experience tends to be dominant (Jessup et al., 2008). We will return to possible reasons why in this difference occurred in the general discussion.

The Description-Experience condition, mirrored the pattern of results found in Experiment 1, risk preference was not influenced by the size of the reinforcer. However, this pattern manifests differently in Experiment 2. In Experiment 1, regardless of the reinforcer size participants in the Description-Experience condition were risk averse (risk preference was below 50%). However, in Experiment 2, participants were risk seeking (risk preference was above 50%). What explains this increase in risk preference? The likely answer is that the changes we made from Experiment 1 to 2 affected participants' mean levels of risk preference. Therefore, there are at least three potential explanations for the increase in risk preference observed in Experiment 2: removing the zero outcomes, making the information source conditions a within-subjects factor, and adding an additional reinforcer size.

It seems unlikely that adding the extra reinforcer size is responsible for this effect. Although, there are few studies on how the size of a reinforcer influence decisions when participants have both description and experience, however we can make inferences based on decision research more generally. If adding the extra set of doors pushed up risk preference in the Description-Experience condition, this suggest the context was influencing their decisions, albeit, in a different way to the Experience condition. If this is true, we should find in other studies in which decisions with a high reinforcer size increase risk preference for other decisions with lower reinforcer size, but, that is not what we find. Mahoney, Buboltz, Levin, Doverspike and Svyantek (2011) for example, gave participants various versions of the Asian Disease problem where they varied how many people the disease would affect. While they found that the magnitude of the disease impact (how many would die) did

influence participants' tendency to choose the risky option, it affected each decision differently. That is, in some problems participants were more likely to choose the risky option than in other problems. Therefore, decisions with a higher reinforcer size didn't lead to equally high levels of risk preference among other decisions.

Moreover, if this increase in risk preference from Experiment 1 to Experiment 2 was an effect of the change in context in our study then we would expect to see this pattern emerge over time. Instead, there is no difference between Block 1 and the Block 6, only in Block 2 does risk preference for the reinforcer sizes differ. Otherwise there is no significant difference between the reinforcer sizes at any other stage of the task.

Could the fact that participants in Experiment 2 experienced both the Experience condition and the Description-Experience condition explain the increase in risk preference from Experiment 1 to Experiment 2? I also think this is unlikely. If this did cause the change, then we would expect to see this increase in risk preference only for participants who completed the Description-Experience condition second. We did not find this, instead there was no order effect. This suggests making the study a within-subject design did not lead to an increase in risk preference in the Description-Experience condition. This still leaves the question of whether the increase in risk preference in Experiment 2 was because of removing zero as a potential outcome. In Experiment 3 we investigated whether the presence of zero had any influence on the level of risk preference in the Description-Experience condition. In Experiment 3, the source of information was a between subjects factor, if we don't find zero influenced risk preference in the Description-Experience condition, this supports the idea to investigate if the within subject change influenced the results. Although as we mentioned above, this is unlikely.

Recency Effect

There was a recency effect in the Description-Experience condition, more specifically participants showed a win-stay/lose-shift pattern for all three reinforcer sizes. To my knowledge this is the first experiment to investigate a recency effect for a hybrid Description-Experience condition. In the Experience condition this pattern was only found in the Experience condition when the reinforcer was medium. Madan, Ludvig, and Spetch (2014), also found a win-stay/lose-shift pattern in their Experience condition. However, Experiment 2 suggests the presence of a recency effect depends not only on the condition, but also on the reinforcer size. Recall from the introduction, a recency effect is unlikely to be the main driver behind decisions from experience (Camilleri & Newell, 2013a). However, the presence of a recency effect might make people more sensitive to the best outcome. For example, in Experiment 1 when the reinforcer was large, and participants make decisions from experience, participants were more likely to choose the risky option after winning the previous gamble compared to when they lost the previous gamble. This could increase the chances participants will be more risk seeking when the reinforcer is large compared to when the reinforcer is small and thus could contribute to an Extreme-Outcome rule. So while a recency effect may not be the main driver of the Extreme-Outcome rule, it could be a cog in the engine that helps explain part of the effect.

However, this idea is not supported by Experiment 2 because the recency effect doesn't explain why participants were also more risk seeking for the medium reinforcer compared to the small reinforcer. Or, why there was no difference between the medium and large reinforcers. Furthermore, the recency effect explanation doesn't explain why there was a win-stay/lose-shift pattern in the Description-Experience condition but no Extreme-Outcome effect. To conclude, the local effect of a recency bias, is unlikely to explain the

global effect of the Extreme-Outcome rule. Nevertheless, Experiment 2 suggest that in the Description-Experience condition a recency effect had more influence over participant's decisions than the Experience condition. The first potential explanation is to do with description part of the condition, but Mata and Rieskamp (2012) argue that there is no recency effect when participants make decisions from description. Therefore, we shall run the same analysis in Experiment 3 to see whether this pattern of recency effect results replicates.

Experiment 3

To investigate whether removing zero as a potential outcome lead to the increase in risk preference in the Description-Experience condition from Experiment 1 to Experiment 2, in the next experiment we manipulated whether zero was present or absent.

Method

Participants: 64 first year psychology students completed the task as part of a course requirement. Participants were randomly assigned into one of two conditions. One participant from the experience condition was excluded from the analysis (1.56%) because they failed to meet the catch trial criterion. This left us with 28 participants in the Experience condition and 35 in the Description-Experience condition.

Design: We used a 2 (information-source condition: Description-Experience, Experience) x 3 (reinforcer size (small, medium, large) x 2 (zero: present, absent) mixed subject design. Information source was the between-subjects factor and reinforcer size and presence of zero were within-subjects factors.

Procedure: The general procedure was as for Experiments 1 and 2. This Experiment had two halves, in one half zero was a potential outcome, in the other zero was absent. (The points used in each half of the Experiment are shown in Table 4. The order participants encountered each phase was randomized across participants. When the second half of the

experiment started, the colours of the doors changed to reduce the chance of the first half of the experiment influencing behaviour in the second half. The Experiment was 24 minutes long, with 12 minutes for each half. Only the choices in the decisions trials were used in the analysis. In the Description-Experience condition participants completed on average 315.6 trials (max = 364, min = 244). Those in the Experience condition completed on average 353.87 trials (max = 395, min = 287).

Table 4: Points used for when zero is either present or absent in Experiment 3

Decision Trial	Zero Present		Zero Absent	
	Safe	Risky	Safe	Risky
Small Reinforcement	20	40/0	25	45/5
Medium Reinforcement	60	80/40	65	85/45
Large Reinforcement	100	80/120	105	85/125

Results

Data Analysis

Risk preference for Experiment 3 was significantly non-normal $W = 0.93, p < .001$. Therefore, to look at the influence the reinforcer size and the presence of zero had on risk preference we analysed the data using a three-way mixed robust ANOVA with 5% trimmed means, an alpha level of 0.05 and we set the number of bootstraps samples to 2000. The robust ANOVA was run using the R function `bwwtrimbt` (Wilcox, 2012), and the data were analysed using the statistical software package R (R Core Team, 2017)⁷.

We analysed data from the final two blocks of the experiment, to give participants enough time to learn the payoff distributions⁸. We found no main effect of condition ($p =$

⁷ There was no order effect. Therefore, whether participants experienced the zero absent or zero present phase first had no influence on the results in Experiment 3.

⁸ Analysing the final three blocks had no influence on the pattern of results.

0.948), there no main effect of presence of zero ($p = 0.107$). There was no interaction between condition and reinforcer size ($p = 0.636$). There was also no interaction between condition and presence of zero ($p = 0.595$), no interaction between reinforcer size and presence of zero, ($p = 0.168$), but there was a main effect of reinforcer size ($p < 0.01$). There was a significant three-way interaction between condition, reinforcer size and presence of zero ($p = 0.024$). To tease apart the three-way interaction we ran a series of follow up tests using the R function `mcppb20` (Wilcox, 2012), these are reported in Table 5. We found the three-way interaction was driven by the difference between the small reinforcer in the Description-Experience condition depending on whether the zero was present or absent. Figure 14 shows that when the zero was absent, participants had higher levels of risk preference ($M = 47.78$, $SD = 38.45$), compared to when the zero was present ($M = 27.70$, $SD = 34.13$). This was the only case in Experiment 3 where the presence of zero had an effect on risk preference in either condition

The pattern of results was not consistent with an Extreme-Outcome effect in the Experience condition, for the best outcome regardless of whether zero was present or absent, as there was no difference between the medium and large reinforcers. There was only limited support for an extreme-outcome effect for the worst outcome. Participants were more risk seeking when the reinforcer was medium compared to when it was small. This was true both when zero was present and absent, however, there was only weak evidence for the significant difference between the small and large reinforcers. Note that according to the confidence interval there is a significant difference (as the interval does not contain zero), however the p value is greater than 0.05. For many statistical tests the confidence interval and p value are based off the same model and are therefore consistent with each other. The reason they differ in this test is because the R function `mcppb20` adjusts the critical value, and the p value is not

based on this adjustment and therefore can provide different results compared to the confidence intervals.

Table 5: Post Hoc results comparing the influence the reinforcer size and the presence of zero had on risk preference in each condition

	Mean Difference Ψ	Lower 95% CI	Upper 95 % CI	P value ⁹
Description-Experience Zero Present				
Small vs Medium*	-21.18	-36.39	-6.40	0.019
Small vs Large *	-23.70	-37.49	-8.18	0.004
Medium vs Large	8.01	-6.42	22.48	0.365
Description-Experience Zero Absent				
Small vs Medium	-2.52	-17.10	13.02	0.807
Small vs Large	-7.03	-22.13	9.04	0.475
Medium vs Large	-4.51	-19.65	10.01	0.633
Experience Zero Present				
Small vs Medium *	-18.73	-32.05	-5.35	0.023
Small vs Large *	-14.90	-27.44	-2.55	0.046
Medium vs Large	3.83	-9.67	17.46	0.655
Experience Zero Absent				
Small vs Medium *	-26.27	-38.43	-13.13	0.002
Small vs Large *	-14.46	-27.84	-0.265	0.090
Medium vs Large	11.81	-2.19	25.89	0.168

⁹Both the p values and confidence intervals are adjusted to control for the probability of at least 1 type one error.

Description-Experience Condition

There was evidence of the low extreme in the Description-Experience condition only when zero was present. That is, when zero was present, participants were less risk seeking when the reinforcer was small compared to when it was large or medium. However, when zero was absent the reinforcer size had no influence on risk preference.

The pattern of effect indicated by the group analysis is indicated from a visual analysis of the modified Brinley plots in Figure 15. In this figure, data from the conditions with zeros are plotted with open circles and those without zeros plotted with closed circles. The top and middle rows show risk preference in the small reinforcer condition versus risk preference in the medium and large reinforcer conditions respectively. In these figures, the majority of the datapoints are above the diagonal reference line. However, in the Description-Experience condition shown in the left column this effect is driven by the effect in the zero-present condition (most open circles are above the diagonal reference line while not consistently above).

Figure 14: Bar graphs showing how the presence of zero influenced the mean risk preference by the source of information. The top panel shows performance in the Experience condition, the bottom panel the Description-Experience condition. Higher risk preference scores indicate the risky option was chosen more often.

Figure 15: Modified Brinley plots comparing the mean risk preference for both conditions by the reinforcer size and presence of zero. The black line running diagonally through each graph is a reference line. In a single graph each participant will have two dots: an open dot (when zero was present) and a closed dot (when zero was absent). Dots above the line show participants who were

Recency Effect

We used the same method to calculate the probability of gambling after a win and after a loss as well as the recency effect score used for Experiments 1 and 2. Figure 17 is a combination of graphs which plot the probability of gambling after both a win and loss for each participant by condition and reinforcer size. In the Description-Experience condition when the reinforcer was small, the majority of participants are above the line, meaning most were more likely to gamble after a win compared to a loss (win-stay/lose shift pattern). On the other hand, when the reinforcer was medium or large, there is no systematic pattern as participants are scattered both above and below the line. In the Experience condition, both when the reinforcer is small and large, most data points lie beneath the line. Thus, most participants are more likely to gamble after losing on the previous trial compared to if they won (win-shift/lose-stay pattern). However, note when the reinforcer is small, there are three data points in the top left corner which drags the average closer to a negative score. The medium reinforcer on the other hand, shows no clear systematic pattern.

Figure 16 shows the mean recency effect score for each condition by reinforcer size. When the reinforcer was small, participants in the Description-Experience condition show a stronger win-stay/lose shift pattern compared to those in the Experience condition. When the reinforcer is medium there is little difference between the conditions. Whereas when the

reinforcer was large the Experience condition shows a win-shift/lose-stay pattern compared to the Description-Experience condition.

Figure 16: Bar graph showing the average recency score for each condition by the reinforcer size. A negative recency score indicates participants that were more likely to gamble if they won a gamble on the previous trial compared to trials when they lost that gamble. Whereas, a positive recency score shows a participant who was more risk seeking when they lost a gamble on the previous trial compared to if they won it. Note that negative recency-effect scores are plotted above the x axis for consistent presentation with the previous figure.

Figure 17: Modified Brinley plots showing participant's tendency to choose the risky option after either winning or losing a gamble in the previous trial (recency effect) in both conditions and by the reinforcer size. Each dot represents a single participant. The black line running diagonally through each graph is a reference line. Any dots above the line show participants who were more likely to choosing the risky option after winning a gamble compared to if they lost a gamble

To investigate whether the recency effect score differed depending on the condition and reinforcer size, we ran a 2 (information-source condition: Description-Experience, Experience) x 3 (reinforcer size: small, medium, large) robust ANOVA using the R function `wwtrimbt` (Wilcox, 2012). There was no interaction ($p = 0.375$) and no main effect of reinforcer size ($p = 0.122$). There was a significant main effect of condition ($p < 0.01$). Participants in the Description-Experience condition showed more of a win-stay/lose-shift pattern ($M = -0.20$, $SD = 0.90$) compared to those in the Experience condition ($M = 0.26$, $SD = 0.91$).

Risk preference across the session

How did risk preference change over time?¹⁰ To answer this question, we calculated each participant's risk preference for every block of the experiment, then we averaged all the scores for Block 1, then Block 2 etc. Figure 18 shows that in the Description-Experience condition when zero is present, from Block 3 onwards the small reinforcer has lower risk preference compared to the medium and large reinforcers. Whereas, when zero is absent the three reinforcers have similar levels of risk preference throughout the experiment. Figure 19 shows that in the Experience condition, the general trend for risk preference is similar regardless of whether the zero is present or absent. In that, as the experiment progresses, the small reinforcer becomes more distant from both the medium and large reinforcer.

¹⁰. For both Figure 18 and Figure 19 the order participants completed the zero present/absent part of the task were collapsed, so although these graphs show the no zero phase first followed by the zero phase of the experiment some participants would have completed it in the opposite order.

Figure 18: Line graph showing how risk preference changed over the course of the Experiment in the Description-Experience condition, by the size of the reinforcer. The first six blocks show performance when zero was absent, the second six when zero was present.

Figure 19: Line graph showing how risk preference changed over the course of the Experiment in the Experience condition, by the size of the reinforcer. The first six blocks show performance when zero was absent, the second six when zero was present.

To investigate how risk preference changed over time we ran an $2 \text{ (zero: present, absent)} \times 3 \text{ (reinforcer size: small, medium, large)} \times 6 \text{ (Blocks)}$ for each condition using the R function `wwwtrimbt` (Wilcox, 2012).

For the Description-Experience condition there was no interaction between reinforcer, presence of zero or block ($p=0.999$), there was also no interaction between reinforcer size and block ($p=0.999$), and no interaction between presence of zero and block ($p=0.096$). There were also no main effects of reinforcer size ($p=0.954$), or Block ($p=0.293$), but there was a main effect of presence of zero ($p<0.01$).

For the Experience condition there was no interaction between reinforcer, presence of zero or block ($p=0.635$), there was also no interaction between reinforcer size and block ($p=0.771$), and no interaction between presence of zero and reinforcer size ($p=0.864$). There were also no main effects of block ($p=0.187$) and presence of zero ($p=0.230$), but there was a main effect of presence of zero ($p<0.01$).

Follow up tests using the R function `mcpp20` reveal when zero was present from Block 3 onwards, risk preference was lower when the reinforcer was small compared to when it was medium or large.

We choose to look at the final two blocks rather than the last three to keep in line with Ludvig et al (2011). The change over time graphs for the Experience condition indicate that by the third block, the pattern of results found is the same found in the final two blocks. This suggests it was appropriate to analyse the final two blocks as by then participants had learned the payoff distributions.

Discussion

Experiment 3 investigated whether the increase in risk preference in Experiment 2 compared to Experiment 1 in the Description-Experience condition was because zero was removed as a potential outcome, the results suggest the answer is yes. When zero was absent in the Description-Experience condition, the pattern of results were similar to Experiment 2:

risk preference was not influenced by the size of the reinforcer. Although the absolute level of risk preference was lower than that of Experiment 2, the pattern of results was the same.

The Description-Experience condition

When the reinforcer was small, risk preference was lower when zero was present compared to when zero was absent. However, the presence of zero bore no influence over risk preference when the reinforcer was medium or large. Therefore, the presence of zero only influenced risk preference only when a gamble could lead to zero points (when the reinforcer was small). When zero was present in the context but the reinforcer was medium (i.e. no zero involved in that decision), it had no influence on that choice. When zero was present the results are unlikely to be because of an effect of the low extreme as otherwise we would have expected to find a similar pattern when zero was absent (and 5 becomes the low extreme), yet we did not.

There are at least two possible explanations for this finding, the first is it was a statistical fluke. In support of this idea, in Experiment 1 (where zero was a potential outcome) there was no difference between the small and large reinforcers like there was in Experiment 3 when zero was present. Furthermore, in Experiment 3, only one group was affected by the presence of a zero, the small reinforcer. In no other groups was risk preference influenced by whether zero was present or absent. Therefore, the lower risk preference for the small reinforcer could be an anomaly. However, this is unlikely as if this result was due to a statistical fluke, then you would expect the pattern of results to be less consistent in the change over time graph. Instead, from Block 3 onwards, when zero is present, risk preference is lower when the reinforcer is small compared to the medium and large reinforcers. This is consistent with the other experiments in this thesis which show that the pattern of results

tends to level out from around Block 3. Moreover, risk preference also tends to level out around this time in similar studies (Ludvig, 2014; Madan et al, 2016)

The more likely explanation for this pattern is participants were only zero averse when a gamble could lead to zero points, which only occurred when the reinforcer was small. According to this explanation, when zero was present and the reinforcer was medium or large, neither the safe nor risky option resulted in zero points. Therefore, there was no zero to be averse too. This is also supported by the fact that when zero was absent there was no influence of reinforcer size on risk preference, as there was no zero available. This explanation suggest that in the Description-Experience condition, choices were made by comparing only the two options available in each trial to each other, rather than taking into consideration the range of possible outcomes available in the context. This is in contrast to the idea behind the Extreme-Outcome rule, which says that the context influences decisions. This suggests the presence of the descriptions made participants consider the context less, as the Experience condition was influenced by the context and the only difference between the conditions is the descriptions. This is consistent with research on decisions from description where participants that make many multiple decisions between safe or risky options tend to not be influenced by the other decisions available in the context (Konstantinidis, Taylor, & Newell, 2017)

Follow up experiments can tease apart the two explanations for why risk preference was lower for the small reinforcer when zero was present in the Description-Experience condition. To address the possibility of a statistical fluke the obvious suggestion is to replicate the experiment and see whether the effect persists. To tackle the possibility that participants were zero averse, an experiment could add zero as a potential outcome into the medium and large reinforcers and compare it to a condition where there is no zero for these

reinforcer sizes. If the zero aversion explanation holds, then we should see less risk seeking when zero is present for the medium and large reinforcers compared to the other condition where zero is absent. Whereas if there is no difference between the conditions, this would suggest the zero aversion idea is not behind the pattern of results.

The Experience Condition

In the Experience condition zero had no influence on risk preference, regardless of the reinforcer size. Thus, like the results of Experiment 2, this suggests the pattern of the Extreme-Outcome rule (at least as found in Experiment 1) and the effect of the low extreme is not explained by zero aversion. This is consistent with research by Ludvig, Madan, Spetch (2014, Experiment 2) and Pisklat (2015). As the presence of zero did not influence risk preference any discussion of the results of the Experience condition applies to both when zero was present and absent.

There was no Extreme-Outcome effect as there was no difference between the medium and large reinforcers and weak evidence to suggest a difference between the small and large reinforcers. Therefore, we tentatively suggest there was an effect of the low extreme but not the high extreme, which like Experiment 2 suggests the need to modify the Extreme-Outcome rule, a topic we will return to in the general discussion.

The difference between the Description-Experience and Experience

Although in Experiment 3, when zero was present in the Description-Experience condition, the pattern of results matched the Experience condition, it is likely the mechanisms behind this pattern are different for each condition. Recall from the introduction, we discussed research which suggests when we have both description and experience information, we are more influenced by our experiences than the descriptions. We argued that finding the same pattern of results in both the Description-Experience and Experience

conditions would indicate that participant's decisions were pulled more by their experiences than by the descriptions. However, Experiment 3 suggests this is an unlikely explanation for finding the same pattern of results in both conditions. In both conditions there was less risk preference when the reinforcer was small compared to when it was medium and large, as well as being no difference between the medium and large reinforcers. For the Experience condition we attribute this pattern to an effect of the low extreme, whereas for the Description-Experience condition either a statistical fluke or zero aversion. The main reason for this is because if the pattern of results was due to an effect of the low extreme would have expected to find a similar pattern when zero was absent (and 5 becomes the low extreme), which we didn't. Rather, as discussed, it is likely alternative explanations provide a better account of the results. We are more confident of an effect of the low extreme in the Experience condition as this was present both when the zero was present and absent.

Recency Effect

The recency effect pattern presented itself differently depending on the condition. Those in the Description-Experience condition showed more of a win-stay/lose-shift pattern compared to those in the Experience condition. This is consistent with multiple studies involving decision making (Gutiérrez-Roig, Segura, Duch, & Perelló, 2016; Tamura & Masuda, 2015; Wang, Xu, & Zhou, 2014). However, the opposite pattern occurred in the Experience condition, here participants were more likely to gamble after a loss compared to after a win (win-shift/lose-stay pattern). This finding could be explained by the gamblers' fallacy: the cognitive bias where people think after a series of losses, a win becomes more likely. For example, if a coin lands on heads five times people tend to think that on the next coin toss a tail is more likely than a heads, which of course it isn't. Therefore, according to the gamblers fallacy after losing a gamble on the previous trial a participant might think they are more likely to win a gamble on the current trial and so continue to gamble. In Experiment

3 we found two patterns recency effects: the win-stay/lose-shift pattern and the win-shift/lose-stay pattern. What is not clear, is what causes one pattern to emerge over the other? It is unlikely to be because of the size of the reinforcer, as there was no effect of reinforcer size.

General Discussion

According to the Extreme-Outcome rule, when we make decisions from experience we are more sensitive to the best and worst outcomes available in a context. Across three Experiments we investigated if the Extreme-Outcome rule predicts participants' decisions in both an Experience and a Description-Experience condition.

In Experiment 1, we found that in the Experience condition the pattern of risk preference was consistent with the Extreme-Outcome rule as participants were more risk seeking when the gamble could lead to the best outcome, compared to when the gamble could lead to the worst outcome. However, in the Description-Experience condition the pattern of risk preference was not consistent with the Extreme-Outcome rule. Instead risk preference was not influenced by the size of the reinforcer.

Experiment 2 tested a prediction of the Extreme-Outcome rule, that participants are only sensitive to the best and worst outcomes. To do so we added in an extra pair of doors, therefore one pair of doors did not lead to an Extreme-Outcome. We also removed zero as a potential outcome to investigate whether zero aversion or delay discounting provided an alternative explanation for the Extreme-Outcome effect found in the Experience condition in Experiment 1. In Experiment 2, we did not find support for the predictions of the Extreme-Outcome rule in the Experience condition, as, although participants were more risk seeking when the reinforcer was large compared to when it was small, there was no difference between the medium and large reinforcers. Therefore, there was an effect of the low extreme, but not the high extreme. Moreover, like

Experiment 1, there was no Extreme-Outcome effect in the Description-Experience condition, rather risk preference was not influenced by the reinforcer size. However, in the Description-Experience condition risk preference was higher than it was in Experiment 1. To investigate if this increase in risk preference was due to removing zero as a potential outcome, in Experiment 3 we manipulated whether zero was present or absent.

Experiment 3 suggested this might be the case. In the Description-Experience condition when zero was present and the reinforcer was small, risk preference was lower compared to when zero was absent. In fact, when zero was absent, risk preference was not influenced by the size of the reinforcer. This suggests zero could have had a role in the increase in risk preference in the Description-Experience condition from Experiment 1 to 2. Moreover, a future study could test if the number of decision trials might also influence risk preference in a Description-Experience condition. Experiment 1, had two decision trials whereas Experiment 2 and 3 had three. This study would also be useful to test if this might explain why in Experiment 1, reinforcer size bore no influence on risk preference in the Description-Experience condition, whereas in Experiment 3 when zero was present, risk preference was influenced by the size of the reinforcer. In the Experience condition, the presence of zero bore no influence on risk preference. Rather, we replicated the effect in Experiment 2, where we suggest there was only an effect of the low extreme. Overall, Experiment 3 suggests the presence of zero influenced risk preference in the Description-Experience condition, and this change might explain the increase in risk preference between Experiment 1 and 2.

Across all three experiments, we consistently found that in the Experience condition risk preference was lowest on choices that involved the low extreme outcome.

A second consistent finding, was a significant difference between the Description-Experience and Experience conditions. Although the pattern of results between the sources of information was similar when zero was present, we argued there were likely different mechanisms behind these two results. Overall, risk preference in the Description-Experience condition was similar regardless of the reinforcer size. Whereas, in the Experience condition we tended to find an effect of the low extreme outcome.

Theoretical implications for the Extreme-Outcome Rule

The Extreme-Outcome rule offers a parsimonious explanation of how we make decisions from experience. However, the results of this thesis suggest the predictions it makes do not hold up. Specifically, according to the Extreme-Outcome rule, participants are sensitive to the best and worst outcome in a given context. We did not find support for this, as we only found an effect of the low extreme. Meaning that, compared to the other reinforcer sizes, when a gamble led to the worst outcome, participants were more likely to choose the safe option. Whereas when the gamble could lead to either the best outcome (large decision trials) or the non-extreme outcome (medium decision trials) there was no difference in risk preference. Therefore, we found an effect of the low extreme, but not the high extreme.

To my knowledge, the only other study that has investigated the Extreme-Outcome rule when there was a ‘middle’ condition is Ludvig, Madan and Spetch, (2014, Experiment 3). In contrast to Experiment 2 and 3, they found participants were sensitive to both the low and high extreme outcomes, and there was a significant difference between the middle and high-extreme outcome. A key procedural differences between their study and ours was they used both losses and gains, whereas we just used gains. However, it seems unlikely the Extreme-Outcome rule only occurs when there is a mix of gains and losses. As other studies have found an Extreme-Outcome effect in experiments that have all gains (Ludvig, Madan, & Spetch, 2014, Experiment 4a) and all losses (Ludvig, Madan, & Spetch, 2014, Experiment

4b). An alternative explanation is that the number of options makes the extreme-outcome more salient. In their study they had eight doors (we used 6 in Experiments 2 and 3), perhaps how many decision trials are used makes it easier to distinguish between extreme and non-extreme outcomes. To test this idea, a study could manipulate how many decisions trials are used. For example, are participants more sensitive to the extreme outcomes when there are more than three decision trial types? Furthermore, as we mentioned in the discussion of Experiment 2, Ludvig et al (2017) found participants were sensitive to the best outcome even when the next best was only a single point behind. Although each condition in their study had only two types of decision trials and therefore had no ‘middle trial’ Therefore, we need more research on the Extreme-Outcome rule when there are ‘middle trials’.

Overall, the results of this thesis suggest the Extreme-Outcome needs to be modified to take into consideration the effect of the low extreme but not the high extreme outcome. For example, when we make decisions from experience we are more sensitive to the worst outcome. However, firstly, because research in this area is scarce we need more research to understand of people make decisions from experience. As well as replicating our findings (especially Experiments 2 and 3), it will also be crucial to investigate whether the effect of the low extreme is found in a study that has both gains and losses with a ‘middle trial’. This would extend the current work being done, by further testing the predictions made by the Extreme-Outcome rule and deepening our understanding of how we make decisions from experience. If the pattern of risk preference in conditions with gains, losses and ‘middle trials’ are consistent with the Extreme-Outcome rule, this would indicate the Extreme-Outcome rule predicts decisions from experience when there are both losses and gains available. However, it does not predict behaviour when only gains are used. Whereas, if there was only an effect of the low extreme, this would support the conclusions of this thesis that we need to modify the Extreme-Outcome rule to take this into consideration.

An alternative explanation to the effect of the low extreme outcome

We found that when the gamble could lead to the worst outcome, participants were more likely to choose the safe option compared to the risky outcome. One explanation for this is that the smallest reinforcer size had a bigger impact because it is a type of extreme outcome. An alternative explanation is that risk preference was lower for small reinforcer magnitude trials because these trials were perceived as more different to the medium trials than were the large trials. To illustrate, take the same points used in Experiment 3 where zero was present. When we take the expected value of the three reinforcer sizes, there is an even 40 points between each one. However, even though there is equal distance between them, this does not mean they are perceived that way. According to diminishing sensitivity the subjective impact of a change (in this case of this Thesis, a change in points) decreases the further away from the reference point (where you started from) the change occurs (Kahneman & Tversky, 1979b). For example, if you had \$0 and I gave you \$50, that would seem like a greater subjective impact than if you had \$50 and I gave you another \$50. Even though in both cases you earn an extra \$50. The reference point in this thesis is lowest outcome available (either zero or five). This idea might also apply to the points used in this Thesis. In Experiment 2 and 3 the difference in expected value between the small and medium and medium and large is the same: 40 points. However, as shown by Figure 20 because of diminishing sensitivity, participants might perceive the difference between the small and medium as greater than the difference between the medium and large. Therefore, rather than the results in the Experience condition being due to effect of the low extreme, it could be because participants perceive the medium and large reinforcers to be similar and thus treat them similarly.

To test this idea, a study could manipulate how close the medium reinforcer is to the small reinforcer. For example, in one condition we could use the same points as Experiment 3 when zero was present. In the other condition, the same points are used except for the medium reinforcer, the points of which are moved closer to the small reinforcer. To illustrate the idea imagine we make the medium reinforcer a choice between 30 points or a 50% chance of 45 else 15 (these points serve as an example, other combinations of values could be used). In this case, the difference in expected value between the small and medium reinforcers would be 10 points, whereas the difference between medium and large reinforcers would be 70. If diminishing sensitivity explains why there was no difference between the medium and large then we would expect a significant difference between the medium and the large reinforcers (as they should be perceived as dissimilar to each other). On the other hand, if

Figure 20: difference between objective difference and perceived difference according to the idea of diminishing sensitivity

there was no difference between the two conditions, this would support the idea of an effect of the low extreme and no effect of the high.

Why did the Description-Experience condition differ from the Experience condition?

Recall from the introduction that when people make decisions when they have both description and experience information, the general consensus is experience has more influence over our decisions than descriptions, and descriptions are often ignored (Jessup et al., 2008). If we found support for that idea then we would have found there was no difference between the Description-Experience and Experience conditions. Overall, we did not find this, instead the presence of the descriptions influenced participants' decision making. Across our experiments, the Description-Experience condition differed from the Experience condition not only in mean risk preference levels, but also in how risk preference changed over time, and how recent outcomes affected risk preference.

Why in this thesis, did the Description-Experience and Experience conditions differ? It might have to do with how often the risky outcomes occurred. Studies that have looked at decision making when we have both descriptions and experience have tended to focus on outcomes that are rare (occur with a 20% chance or less of occurring, etc. getting 50 points 10% of the time). Therefore, in these studies the outcome (getting points) doesn't happen that often. Whereas throughout this thesis a participant who choose the risky option would win 50% of the time and thus experienced the outcome more often than studies that have focused on rare events. At first this sounds like a paradox, why would experiencing the outcome more often, lead to descriptions having more influence over our decisions? I think the most likely explanation is because how often participants win points on the risky option tends to align with the descriptions they are given. When studies involve rare events, even if the descriptions are accurate, it can often seem as though they are not reliable because they are unlikely to accurately describe a participants' experience over small numbers of trials. For example, if 1000 participants each sample 20 times from an option that provides points 10% of the time, chances are 285 people would experience the event 10% of the time (which would match the description), 323 would encounter it more than expected, and 392 would not encounter it at all

(Hertwig & Frey, 2012). Thus, when the event is rare, often the description of the event won't match up with what people experience.

Why does it matter if the descriptions don't align with what participants' experience? Because if they don't trust the descriptions then they will trust their experiences instead. Weiss-Cohen, Konstantinidis, Speekenbrink and Harvey (2016) had participants make repeated choices between certain and risky options. There were four conditions, a description condition, an experience condition, and two description-experience conditions. In one of these Description-Experience conditions, the descriptions aligned with participants' experience, for example, they said there is an 80% to win points, and they won points 80% of the time. In the other description-experience condition, the descriptions did not match the actual chance of winning points. For example, the descriptions might say choosing this option will win points 80% of the time, but it was actually 20%. In the Description-Experience condition where the descriptions were accurate, participants' behaviour was no different to a Description condition. Whereas in the conflicting conditions, participants' decisions were not significantly different from an experience condition. In other words, participants trusted the descriptions when they aligned with what they experienced. This effect is also consistent from findings from the rule-governed behaviour literature. For example, Galizio (1979) found that participants tend to follow rules (which are analogous to the descriptions in our study) that they are given, until they find out through experience that these rules do not describe the contingencies associated with the outcomes. Therefore, throughout our three Experiments participants in the Description-Experience condition could have relied more on the descriptions because they matched with what they experienced. To test this idea, we could run a variation of a study in this thesis (say Experiment 1) where in the Description-Experience condition the descriptions of the outcomes do not align with participants' experiences. For instance, the descriptions might say the probability of winning a gamble is

50%, where in reality participants win points only 20% of the time (actual probabilities of winning points would be the same in both the Experience- and Description-Experience condition). If we find support for the idea that in this thesis the Experience and Description-Experience conditions differed because participants trusted the descriptions, then we should find that in this experiment there would be no difference between the conditions. That is when the descriptions are inconsistent with what participants would experience, participants rely on experience and so the experience and description/experience conditions would be the same. Whereas, if there is a difference, this suggests another mechanism is at play.

Furthermore, it would be useful to replicate one of the experiments in this thesis but add in a description condition. This would allow us to understand how much of a difference the descriptions had on decisions. Whereas, in our three studies we can only say descriptions had an impact. Although, we don't know the degree of the impact as we did not have a description condition as a comparison.

Recency Effect

A recency effect was a consistent finding across the three Experiments, yet the way the recency effect manifested varied. That is, participants were influenced by the outcome of the previous gamble, but the way that they were influenced was not consistent. In Experiment 1, in both conditions, winning a gamble on the previous trial lead to a greater chance choosing the risky option on the next trial compared to it they lost the previous gamble. However, the opposite pattern occurred in Experiment 3 in the Experience condition. Here, participants were more likely to gamble after a loss compared to a win. Throughout this thesis the pattern of the recency effect has varied, yet the overall pattern of the mean pattern of risk preference was consistent: An effect of the low extreme in the Experience condition, and no effect of reinforcer size in the Description-Experience condition. Few studies investigating the

Extreme-Outcome rule have looked for recency effects. Those that have tend to find a win-stay/lose-shift pattern (Madan et al., 2014), we found this pattern in Experiment 1 in the Experience condition, but the opposite pattern emerged in Experiment 3. Perhaps the magnitude influences the type of recency effect found. As I am unaware of research that has investigated how the magnitude of points won influences recency effects future studies could investigate this.

Overall, the results suggest this local effect (recency effect) did not drive the global pattern found by the risk preference. In other words, the pattern of risk preference we found were unlikely to be driven solely by participant's response to the immediately preceding outcome. If it were we expect a more consistent recency effect pattern from Experiment 2 and 3. This is consistent with research that finds mixed results for a recency effect in decisions from experience (Hau, Pleskac, & Hertwig, 2010). However, to my knowledge, there are no studies that have investigated a recency effect in a hybrid description-experience condition. Furthermore, research on recency effects in decisions from descriptions is scarce because most of these studies involve one off decisions rather than repeated choices and therefore there is no chance for a recency effect to emerge. Therefore more work needs to be done to gain a better understanding of how descriptions influence recency effects. The results of this thesis suggest descriptions do influence recency effect as there are consistently difference between how a recency effect pattern emerges between the Experience and Description-Experience conditions. Moreover, most of the research on recency effects and the Extreme-Outcome rule focus on the overall experience condition. However, the results of our study find the pattern of recency effect can depend not only on what source of information is present but also on the reinforcer size.

Memory Tests

In future studies it would be useful to have a memory test at the end of the task. This has been quite common among studies investigating the Extreme-Outcome rule. These memory tests have generally involved presenting participants each door used in the Experiment one at a time and asking them to recall the first amount of points that comes to mind. They are then asked how often each outcome occurred (Madan, Ludvig & Spetch, 2017; Madan et al., 2014). Participants tend to be more likely to recall the extreme outcomes and state that they occurred more often than they actually did (Ludvig, Madan, McMillan, Xu, 2017; Madan et al., 2016). It would be interesting to see whether the same results are found in a Description-Experience condition. Would participants be more likely to remember the extreme outcomes in the memory tests but not reflect this bias in their risk preference? If we replicated the effect of the low extreme in the Experience condition but found that both types of extreme outcomes were over weighted in memory, this would go against the idea of that memory bias drives decision making about risky outcomes, or at least the idea that this bias influences decisions from experience. If the memory tests found participants had a better memory for the Extreme-Outcomes compared to the non-extreme outcomes, but the risk preference didn't reflect the Extreme-Outcome rule this would suggest the extreme outcomes were not driving their behaviour.

Conclusion

Across, three experiments, we investigated how well the Extreme-Outcome rule predicted decisions from experience. Our results suggest that people are more sensitive to the worst outcome available in a context rather than the best. We have also shown that when we have both description and experience information and an event is equally likely to occur or not descriptions have more influence over our decisions than we predicted. Our research also highlights some important directions decisions from experience research could take to better

understand how our past experiences influence our future decisions. This thesis started with a story about Mount Vesuvius, and how the people who lived so close to it had vastly different options on how dangerous it was compared to the experts. Stories like these highlight the importance of studying how we make decisions from experience. To do so, perhaps we should focus less on what the chances are, and focus more on people's experiences.

References

- Armstrong, B., & Spaniol, J. (2017). Experienced probabilities increase understanding of diagnostic test results in younger and older adults. *Medical Decision Making*, 37(6), 670-679. <https://doi.org/10.1177/0272989X17691954>
- Barron, G., & Barron, G. (2014). The coexistence of overestimation and underweighting of rare events and the contingent recency effect the coexistence of overestimation and underweighting of rare events and the contingent recency effect. *Judgment and Decision Making*, 4(6), 447-460.
- Bishop, D. V. M., & Hsu, H. J. (2015). The declarative system in children with specific language impairment: A comparison of meaningful and meaningless auditory-visual paired associate learning. *BMC Psychology*, 3(1).<https://doi.org/10.1186/s40359-015-0062-7>
- Brodie, D. A., & Murdock, B. B. (1977). Effect of presentation time on nominal and functional serial-position curves of free recall. *Journal of Memory and Language*, 16(2), 185.
- Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: a direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, 18, 377-384. <https://doi.org/10.3758/s13423-010-0040-2>
- Camilleri, A. R., & Newell, B. R. (2013). Mind the gap? Description, experience, and the continuum of uncertainty in risky choice. In *Progress in brain research* (Vol. 202, pp. 55-71). Elsevier. <https://doi.org/10.1016/B978-0-444-62604-2.00004-6>
- 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(1), 54-71. <https://doi.org/10.1016/j.cognition.2012.09.001>
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *The Quarterly Journal of Experimental Psychology*, 70(10), 2048-2059.<https://doi.org/10.1080/17470218.2016.1220608>
- Chu, Y.-P., & Chu, R.-L. (1990). The Subsidence of Preference Reversals in Simplified and Market-like Experimental Settings: A Note. Source: *The American Economic Review*, 80(4), 902-911. Retrieved from <http://www.jstor.org/stable/2006715>
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., ... Technion, I. (2010). A Choice Prediction Competition: Choices from Experience and from Description. *Journal of Behavioral Decision Making* J. Behav. Dec. Making, 23, 15-47. <https://doi.org/10.1002/bdm.683>
- Feild, Andy, Miles, Jeremy, Feild, Z. (2012). *DISCOVERING STATISTICS Discovering statistics Using R* (1st ed.). London: SAGE Publications Ltd.

- Galizio, M. (1979). Contingency-shaped and rule-governed behavior: Instructional control of human loss avoidance. *Journal of the Experimental Analysis of Behavior*, 31(1), 53-70.
<https://doi.org/10.1901/jeab.1979.31-53>
- Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, 17(3), 295-314. [https://doi.org/10.1016/0010-0285\(85\)90010-6](https://doi.org/10.1016/0010-0285(85)90010-6)
- Green, L., Myerson, J., & Calvert, A. L. (2010). Pigeons' Discounting of Probabilistic and Delayed Reinforcers. *Journal of the Experimental Analysis of Behavior*, 94(2), 113–123.
<https://doi.org/10.1901/jeab.2010.94-113>
- Gutiérrez-Roig, M., Segura, C., Duch, J., & Perelló, J. (2016). Market imitation and win-stay lose-shift strategies emerge as unintended patterns in market direction guesses. *PLoS ONE*, 11(8).
<https://doi.org/10.1371/journal.pone.0159078>
- Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, 23(1), 48-68. <https://doi.org/10.1002/bdm.665>
- Heilbrunner, S. R., & Hayden, B. Y. (2016). The description-experience gap in risky choice in non-human primates. *Psychonomic bulletin & review*, 23(2), 593-600.
<https://doi.org/10.3758/s13423-015-0924-2>
- 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(8), 534-539.
<https://doi.org/10.1111/j.0956-7976.2004.00715.x>
- Hertwig, R., & Frey, R. (2015). The challenge of the description-experience gap to the communication of risks. In *The SAGE handbook of risk communication* (pp. 24-40). Sage.
<https://doi.org/10.4135/9781483387918.n6>
- Hertwig, R., & Frey, R. (2017). How representations of knowledge shape actions. In *Knowledge and Action* (pp. 127-143). Springer, Cham. https://doi.org/10.1007/978-3-319-44588-5_8
- Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, 19(10), 1015-1022.
<https://doi.org/10.1111/j.1467-9280.2008.02193.x>
- Kahneman, D., Fredrickson, B. L., Schreiber, C. A., Donald, A., Kahneman, D., Fredrickson, B. L., ... Redelmeier, D. A. (1993). When More Pain Is Preferred to Less : Adding a Better End Adding a Better End, 4(6), 401–405.
- Kahneman, D., & Tversky, A. (1979a). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263–292. <https://doi.org/10.2307/1914185>

- Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and brain sciences*, 19(1), 1-17.
<https://doi.org/10.1017/S0140525X00041157>
- Konstantinidis, E., Taylor, R. T., & Newell, B. R. (2017). Magnitude and incentives: revisiting the overweighting of extreme events in risky decisions from experience. *Psychonomic Bulletin and Review*, pp. 1–9.
- Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, 116(2), 286-295.
<https://doi.org/10.1016/j.obhdp.2011.05.001>
- Lejarraga, T., Woike, J. K., & Hertwig, R. (2016). Description and experience: How experimental investors learn about booms and busts affects their financial risk taking. *Cognition*, 157, 365-383. <https://doi.org/10.1016/j.cognition.2016.10.001>
- Ludvig, E. A., Madan, C. R., Pisklak, J. M., & Spetch, M. L. (2014). Reward context determines risky choice in pigeons and humans. *Biology Letters*, 10(8), 20140451-.
<https://doi.org/10.1098/rsbl.2014.0451>
- Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014a). Extreme Outcomes Sway Risky Decisions from Experience. *Journal of Behavioral Decision Making*, 27(2), 146–156.
<https://doi.org/10.1002/bdm.1792>
- Ludvig, E. A., Madan, C. R., & Spetch, M. L. (2014b). Extreme Outcomes Sway Risky Decisions from Experience. *Journal of Behavioral Decision Making*, 27(2), 146–156.
<https://doi.org/10.1002/bdm.1792>
- 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? 6(6).
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, 21(3), 629–636. <https://doi.org/10.3758/s13423-013-0542-9>
- Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *The Quarterly Journal of Experimental Psychology*, 70(10), 2048-2059. <https://doi.org/10.1080/17470218.2016.1220608>
- Madan, C. R., Spetch, M. L., & Ludvig, E. A. (2015). Rapid makes risky: Time pressure increases risk seeking in decisions from experience. *Journal of Cognitive Psychology*, 27(8), 921-928.
<https://doi.org/10.1080/20445911.2015.1055274>
- Mahoney, K. T., Buboltz, W., Levin, I. P., Doverspike, D., & Svyantek, D. J. (2011). Individual differences in a within-subjects risky-choice framing study. *Personality and Individual Differences*, 51(3), 248–257. <https://doi.org/10.1016/j.paid.2010.03.035>

- Madan, C. R., Spetch, M. L., & Ludvig, E. A. (2015). Rapid makes risky: Time pressure increases risk seeking in decisions from experience. *Journal of Cognitive Psychology*, 27(8), 921-928. <https://doi.org/10.1080/20445911.2015.1055274>
- McKerchar, T. L., & Renda, C. R. (2012). Delay and Probability Discounting in Humans: An Overview. *The Psychological Record*, 62(4), 817-834. <https://doi.org/10.1007/BF03395837>
- Noguchi, T., & Hills, T. T. (2016). Experience-Based Decisions Favor Riskier Alternatives in Large Sets. *Journal of Behavioral Decision Making*, 29(5), 489-498. <https://doi.org/10.1002/bdm.1893>
- Regenwetter, M., & Robinson, M. M. (2017). The construct-behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review*, 124(5), 533-550. <https://doi.org/10.1037/rev0000067>
- Scandone, R., Giacomelli, L., & Gasparini, P. (1993). Mount Vesuvius: 2000 years of volcanological observations. *Journal of volcanology and geothermal research*, 58(1-4), 5-25. [https://doi.org/10.1016/0377-0273\(93\)90099-D](https://doi.org/10.1016/0377-0273(93)90099-D)
- Scholten, M., & Read, D. (2014). Prospect theory and the "forgotten" fourfold pattern of risk preferences. *Journal of Risk and Uncertainty*, 48(1), 67-83. <https://doi.org/10.1007/s11166-014-9183-2>
- Shampanier, K., Mazar, N., & Ariely, D. (2007). Zero as a Special Price: The True Value of Free Products. *Marketing Science*, 26(6), 742-757. <https://doi.org/10.1287/mksc.1060.0254>
- Tamura, K., & Masuda, N. (2015). Win-stay lose-shift strategy in formation changes in football. *EPJ Data Science*, 4(1), 1-19. <https://doi.org/10.1140/epjds/s13688-015-0045-1>
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453-458. <https://doi.org/10.1126/science.7455683>
- Tversky, A., & Kahneman, D. (1992). Advances in Prospect-Theory - Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297-323. <https://doi.org/10.1007/BF00122574>
- Wang, Z., Xu, B., & Zhou, H. J. (2014). Social cycling and conditional responses in the Rock-Paper-Scissors game. *Scientific reports*, 4, 5830. <https://doi.org/10.1038/srep05830>
- Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, 135, 55-69. <https://doi.org/10.1016/j.obhdp.2016.05.005>
- Wilcox, R. (2012). *Introduction to robust estimation and hypothesis testing* (3rd ed.). Amsterdam: Academic Press.

- Yechiam, E., Erev, I., & Barron, G. (2006). The effect of experience on using a safety device. *Safety science*, 44(6), 515-522. <https://doi.org/10.1016/j.ssci.2005.11.006>
- Yechiam, E., Rakow, T., & Newell, B. R. (2015). Super-underweighting of rare events with repeated descriptive summaries. *Journal of Behavioral Decision Making*, 28(1), 67–75. <https://doi.org/10.1002/bdm.1829>
- Yoon, S., Vo, K., & Venkatraman, V. (2017). Variability in Decision Strategies Across Description-based and Experience-based Decision Making. *Journal of Behavioral Decision Making*, 30(4), 951-963. <https://doi.org/10.1002/bdm.2009>