A Contingency Learning Approach to False Belief Formation

Shankar, Sahana

Awarding institution:
King's College London

KING’S COLLEGE LONDON

INSTITUTE OF PSYCHIATRY, PSYCHOLOGY & NEUROSCIENCE

Department of Psychology

A Contingency Learning Approach to False Belief Formation

Sahana Shankar

Supervisors:

Dr. Tim Rakow, Dr. Nicola Byrom & Dr. Wijnand van Tilburg

A thesis submitted for the degree of Doctor of Philosophy in Psychology

April 2023

Abstract

The prevalence of fake news in the current era, is thought to lead to the formation of false beliefs (Ecker et al., 2022). The false beliefs that arise from fake news are dangerous as they often guide decision around health, politics, and finance (European Commission, 2010). However, attempts to remedy false beliefs by exposing people to facts have largely failed to work (Yarritu & Matute, 2015). As such, a systematic, fundamental understanding of the mechanism through which false belief formation occurs is essential to develop effective countermeasures against fake news. Contingency learning provides one framework to understand the development and maintenance of some false beliefs. Contingency learning paradigms can be utilised to further understand the formation of associations between events when there is no true relationship between the events. In this thesis, I investigate these contingency learning principles and put them in context of false belief formation, especially, but not exclusively, in relation to medicine and disease.

Currently, the majority of contingency learning studies utilize active paradigms. These active paradigms require people to be actively involved in the process of responding to individual stimuli before their perceived control over the outcome is measured. Less attention has been paid to passive contingency paradigms. In these, people instead of directly interacting with the stimuli, simply witness the stimuli and outcomes appear, followed by a measure of their perceived relationship. Despite being less prevalent in research, these passive contingency paradigms are possibly especially interesting to study in context of false belief formation in society, which, for many, may be the results of associations formed through observations (e.g., being exposed to social media information)

rather than actively interacting with this information. I adopted this passive contingency paradigm accordingly.

In Chapter 1, I outline and discuss a range of false beliefs, and then narrow down to false beliefs formed from causal illusions. Chapter 2 discusses the methodological framework of the contingency learning literature in light of the methodology of the seven experiments used in this thesis. This review then led to the design of the seven experiments that forms the empirical backbone of this thesis. The first 5 experiments considered the boundary conditions under which density effects occur in context of judgments of effectiveness of putative treatments for disease. These five studies thus helped to uncover what may prevent or exacerbate false beliefs, and simultaneously addressed methodological expansions of established contingency learning paradigms. These initial five studies focused on beliefs that were newly formed; there was no a priori knowledge about cue and outcome prior to the experiments. The final two experiments, in turn, examined the outcome density effect with regards to the broader context in which they may form: the role of pre-existing dispositional affect, and the existence of prior beliefs. These seven experiments are reported in the chapters that follow.

In Chapter 3 I investigate the influence of outcome density on judgements of effectiveness of medication, in a contingency paradigm, to test if outcome density effects differ between online and lab settings (Experiment 1). Large effects of outcome density were found across both settings, suggesting that participants give higher judgement of causality when the probability of outcome is higher compared to when the probability of the outcome is lower. In Chapter 4, I examined density effects by comparing cue and outcome density effects directly, to see if the same cognitive principles underlie both

(Experiment 2). Effects of both cue and outcome density were found, however outcome density effects were larger than cue density effects.

To further understand these effects, in Chapter 5, Experiments 3 and 4 were conducted in which the order of cue and outcome presentation were manipulated, along with cue and outcome density effects. This extension further compares the density effects and allows for the exploration of whether the effect of cue or outcome density on contingency judgment is altered by reversing the order in which event information is seen. The size of the density effects stayed relatively consistent across the presentation order manipulation, and outcome density effects remained larger than cue density effects. The order of event presentation seems to make little or no difference to outcome and cue density effects.

In Chapter 6 I further examined the outcome density effect across positive, negative and zero contingency conditions within the same study to compare the effect of outcome density (Experiment 5). In addition to this, the effect of scale length was explored by comparing a frequently used unidirectional judgement scale to a bidirectional scale. I found effects of outcome density, contingency, and scale. In addition to this, people naturally seem to be making associations that are more positive than the contingency would indicate.

In chapter 7 and 8 I examine the role of prior beliefs (Experiment 6). The influence of pre-existing prior beliefs on judgements in passive contingency tasks has not been explored thus far, as far as I am aware. In addition to this, I examined whether the outcome density effect occurs in different contexts. Chapter 7 covers a pilot study, and Chapter 8 then builds on that through an experiment. There was some variability in judgement and prior beliefs across contexts. Overall, there was evidence to suggest that context-specific prior beliefs influence contingency judgements.

In Chapter 9, the final empirical chapter of this thesis, the role of affect is examined (Experiment 7). The contingency literature has focused on the depressive realism effect within active contingency paradigms. As such this experiment extended the research by examining a wider range of affect to identify whether naturally occurring differences in emotional state relate to differences in judgment. Affect was found to influence contingency judgements. Specifically, in the low outcome density conditions there was an influence of positive affect, negative affect, and anxiety on contingency judgements. In the high outcome density conditions, there was an influence of anger and stress on contingency judgements.

Finally in Chapter 10 I discuss the findings in terms of their theoretical, empirical, and methodological contributions. In addition to this, I also provide recommendations for future areas of research within this field and discuss strengths and limitations of the findings. Together, the work presented in this thesis has enhanced the understanding of the mechanisms under which false beliefs may be formed from causal illusions and provides potential implications for current interventions on reducing false beliefs as well as future research.

Acknowledgements

I would like to thank my supervisors Dr. Tim Rakow, Dr. Nicola Byrom & Dr. Wijnand van Tilburg who created a rich environment of intellectual curiosity and growth, with a balance of good humour and support. I truly think I found the most supportive supervisors, as the past 4 years of the PhD has been an incredibly positive experience. I have had the time and guidance to learn so many skills, and the opportunities to gain experiences in areas that may not have been possible otherwise.

I also have to give my thanks to King's College London where I have had the pleasure of completing nearly 8 years of education over the course of my undergraduate, masters, and doctoral degrees.

எனது பெற்றோருக்கு நன்றி, உங்கள் முடிவில்லாத ஆதரவுக்காகவும், எனது பல மனநிலைகளைத் தாங்கியதற்காகவும். கல்வியின் மதிப்பை எனக்குக் கற்பித்ததற்காக நான் என்றென்றும் நன்றியுடன் இருப்பேன். நீங்கள் இருவரும் இல்லாமல் என்னால் இதையெல்லாம் சாதித்திருக்க முடியாது. எனக்கு நல்ல வாழ்க்கை அமைய வேண்டும் என்பதற்காக நீங்கள் எதிர்கொண்ட அனைத்து கஷ்டங்களுக்கும், இந்த ஆய்வறிக்கையை உங்கள் இருவருக்கும் சமர்பிக்கிறேன்.

தம்பி குட்டி, you have brought the much-needed humour into my life and have always been so uplifting. You forever humble me, and I will always be grateful for that.

Finally, to my husband Piriyandan—for your love, kindness, and never-ending belief in me throughout this whole journey. Thank you for cheering me on during the high times and encouraging me during the more challenging times of the past few years – I am forever thankful for your support.

Sahana Shankar

April 2023

Statement of Authorship

I, Sahana Shankar, confirm the work in this thesis is my own work. The contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. The work submitted is within the specified word limit (100,000 words).

Covid-19 Impact Statement

Data collection for Experiment 1 of Chapter 3 occurred in March 2020. Due to the lockdown and lab closure, participant recruitment was stopped. The experiment was 2 sessions, as such out of the 105 participants recruited, 84 completed the first session, and only 58 participants were able to complete both sessions. The time between sessions was planned to be 3 days, however due to the impact of lockdown, for some participants there was a longer gap between sessions.

Table of Contents

List of Figures

List of Tables

Chapter 1: Introduction

In the age of the Internet, the prevalence of fake news is supporting an increase in beliefs that are not founded on evidence (Carroll, 2015). Algorithms used by social media and search engines exacerbate the impact of spurious beliefs by favouring information that echoes people's ideas and beliefs. The resulting myths become increasingly difficult to eradicate (Lewandowsky et al., 2012).

As a case in point, on the 5th of November 2016, the fake news headline "FBI Agent Suspected in Hillary Email Leaks Found Dead in Apparent Murder-Suicide" trended among more than half a million Facebook users. Three days later, Hillary Clinton was defeated by Donald Trump in the US Presidential elections. Fake stories such as this one has been linked to this election result (Allcott & Gentzkow, 2017).

This problem is not confined to the USA: Fake news is on the rise in European countries and often guides health, financial, and family decisions (European Commission, 2010)(European Commission, 2010). For example after the 2023 earthquake in Turkey and Syria, pictures, videos and stories from previous disasters in other countries where shared claiming these were currently happening in Turkey and Syria ('Turkey Earthquake', 2023). This instilled panic and worry among many, who were unaware of the false nature of the information.

1.1 Thesis rationale and motivation

To combat fake news, false beliefs must be countered, as research suggests that fake news can lead to the formation of false beliefs (Ecker et al., 2022). Unfortunately, campaigns that simply expose people to facts in the hope that this counters false beliefs have largely failed (Schwarz et al., 2007). Research indicates that showing people facts is not enough to eradicate false beliefs (Yarritu & Matute, 2015). Strikingly, such approaches can even have the opposite effect of strengthening the myth they were designed to combat (Lewandowsky et al., 2012). As such, a systematic, fundamental understanding of false belief formation is essential to develop effective countermeasures against fake news.

Many false beliefs are false causal beliefs, and these are among the most important false beliefs because they may encourage courses of action that do not result in the hoped-for benefits. Causal beliefs arise from individuals identifying causality, the process through which one event causes an effect, and the effect is dependent on the cause to occur. Causal relationships suggest that cause and effect are related, even if it may be temporary, and that the outcome follows the cause and does not precede it. People are predisposed to form causal beliefs and do so from a very early age (Sobel & Kirkham, 2006). However the inferences made about the underlying mechanisms of causality do not always match the causal belief about cause and effect, often described as a cognitive illusion (Sawa, 2009). Hence studying the mechanisms through which causal beliefs and the associated reasonings are made, give further understanding to issues that arise from cognitive illusions such as pseudoscience and superstitious thinking. Contingency learning provides one framework to understand the development and maintenance of some false beliefs.

Contingency theory (Allan, 1980) provides a normative account of contingency belief, stating that the belief should correspond with the difference in the probability that two events co-occur, P (Event 1|Event 2), and the probability that one of these events occurs in isolation, P (Event 1|No Event 2). This difference is quantified as ΔP—the degree of contingency. Unfortunately, people do not adhere to this normative account; contingency beliefs are often biased. People tend to believe that events that co-occur often must somehow be causally related, even if they are not (Allan & Jenkins, 1983; Blanco et al., 2011; Byrom et al., 2015; Van Tilburg & Igou, 2014). People furthermore use irrelevant contextual cues to infer causality (van Tilburg & Igou, 2014) yet fail to consider information actually relevant to the judgement (Matute et al., 2015). These discoveries in contingency learning research offer an ideal basis for understanding why people believe fake news, and what can be done to counter these beliefs.

Take Vitamin C for example. Many people believe that taking Vitamin C cures the common cold (Gorton & Jarvis, 1999). A systematic review concluded that there was no trial evidence for the efficacy of Vitamin C (Hemilä & Chalker, 2013). Taking Vitamin C to cure the common cold was equally as effective as taking no medication. As people tend to recover with or without Vitamin C supplement, they are at risk of incorrectly attributing their recovery to Vitamin C. Nonetheless, the relationship between Vitamin C and the common cold is a relatively low-cost scenario.

However, the situation becomes deeply troubling when considering the unsubstantiated belief that cannabis oil is a natural treatment for cancer (www.CureYourOwnCancer.org). Indeed, Facebook pages echoing this claim are followed by many thousands (CannabisOilCuresCancer, 35,282 followers; TheCureForCancers, 110,186 followers). Here people have inferred a correlation between cancer remission with taking cannabis oil and may choose to use cannabis oil as a treatment for cancer, which could have disastrous health consequences. There are similar claims around “clean” eating curing cancer, whereby some people have fabricated claims of having cancer which were cured through diets and alternative medicine, which many have believed and followed (S. A. Baker & Rojek, 2020).

Clearly, there is real danger in forming false beliefs, which makes understanding people’s inability to differentiate between fake and real relationships a vital objective. Research into false beliefs by examining their roots in contingency learning may provide crucial insight into understanding and subsequently preventing the negative impact of fake news. This is a much-needed development.

My objective is to understand this process in the context of false beliefs. My research addresses the factors that cause people to form false beliefs and how to combat this important problem. In doing so, my PhD fits the global call by governments worldwide, including the UK, to promote a knowledge based society and decrease belief in fake news (Claesson, 2019; Department for Digital, Culture, Media & Sport, 2021; Government Communication Service, 2022; The Digital, Culture, Media and Sport Committee, 2018). To meet this objective, I conducted a series of experiments using lab-based contingency learning paradigms. These are adapted from those used in previous research (e.g. Blanco et al., 2013; Matute et al., 2015). In a typical study of this kind, the participants’ task is to view a series of cue-outcome pairings and then to judge the contingency between the cue and outcome.

1.1.1 Chapter and thesis overview

In this chapter, I review literature on the key concepts used in this thesis. Firstly, I will discuss different forms of false beliefs, and then focus on false beliefs formed under the illusion of causality. I will focus on outcome and cue density effects, using a contingency learning framework. I will outline contingency theory in detail as well two other models of associations, the probabilistic contrast model and the associative perspective. Finally, I will give an overview of the thesis as whole, what each chapter consists of and how each chapter relates to each other.

This chapter of the thesis works to give a good understanding of the theoretical basis and framework underpinning the following empirical chapters. Chapter two will provide details on the general methodological framework of this project. The subsequent seven chapters report experiments that examine whether and how false beliefs about cue-outcome contingencies are influenced by various factors such as outcome density, cue density, outcome and cue presentation order, response scale, context, prior belief, and affect. The final chapter concludes the overall findings of the empirical work and the contributions these have in this thesis.

1.2 Examples of false beliefs and how they are studied in the lab

Causal relationships suggest that cause and effect are related, even if it may be temporary, and that the outcome follows the cause and does not precede it. People are predisposed to form causal beliefs from a very early age (Sobel & Kirkham, 2006). However the inferences made about the underlying mechanisms of causality do not always match the causal belief between cause and effect, often described as a cognitive illusion (Sawa, 2009). Hence studying the mechanisms through which causal beliefs and the associated reasonings are made, give further understanding to issues that arise from cognitive illusions such as pseudoscience and superstitious thinking.

In this section, several types of causal belief are discussed, in particular scenarios where people may form false beliefs are outlined. These false beliefs are about the contingency between events or closely related events. The method through which these are studied in the lab will also be outlined.

Chapter 2 will discuss the methodology of the contingency paradigm in detail. However, in this section the different paradigms of the casual beliefs will be compared to the contingency paradigm. As such an outline of the contingency paradigm is given below.

1.2.1 The typical outcome density paradigm

Contingency learning paradigms are typically fairly similar: participants are exposed to a sequence of trials in which two stimuli—usually a cue and an outcome—may or may not co-occur. This results in four possible combinations of cue and outcome, as presented in Table 1. The contingency between these stimuli is manipulated by changing the frequency of stimuli co-occurrence. In some studies, participants are given training with similar stimuli to those in the main trials, to ensure participants understand the task. Next, participants take part in the main task. Participants also provide their estimation of the degree to which the cue and outcome are related to each other.

1.2.2 Illusory Correlations

Table 1 Visual presentation of the possible cue-outcome pairings in a contingency paradigm

	Outcome	
Cue	Present	Absent
Present	A	B
Absent	C	D

An illusory correlation is the perception of a relationship between variables when there is none. When the distribution of two variables and their attributes are skewed such that one may be infrequent but distinctive (Ernst et al., 2019), then this can result in people's overestimation of the frequency of the infrequent or more distinctive events, and making associations based on that (Weigl et al., 2018). Some cases of contingencies can be similar to illusory correlations.

An example of this is the formation of stereotypes. A woman has her purse stolen by a person on a bike in a trench coat. In the future, she hugs her close purse each time she sees a person riding a bike in a trench coat. The infrequent and distinctive nature of these experiences may lead to stereotypes.

One of the earlier studies conducted in illusory correlations, demonstrated the process by which illusory correlations could result in stereotypic beliefs about minorities. In the study, participants read a series of statements describing positive or negative behaviours of individuals who were either part of a majority group (A) or a minority group (B) (Hamilton & Gifford, 1976). Both groups had the same proportion of positive and negative statements presented, whilst group A had more statements overall. However, there was no real relationship between the statements or the groups (Table 2). Participants were asked to make a judgement about the group membership of the individuals they were shown. The results indicated that people were accurate when making associations between

Table 2 Hamilton and Gifford (1967) information given to participants

Groups	Positive attributes	Negative attributes	Total
Group A (majority)	18 (69%)	8 (30%)	26
Group B (minority)	9 (69%)	4 (30%)	13
Total	27	12	39

group membership and positive behaviours. However, they were more likely to overestimate the number of people in the minority group who exhibit negative behaviours. This example shows how the illusory correlation can result in stereotype formation, due to a biased perception of contingency.

Typically, in illusory correlations paradigms, participants are shown a sequence of statements about the members of two groups, with each statement presenting a negative or positive attribute of members of each group. The negative and positive statements for the groups are skewed, as shown in Table 3, so that one group is the 'majority' group whilst the other is the 'minority' group. The proportion of negative and positive statements are the same. This often leads to a perception of a contingency between the groups and the attributes. The perceived contingency is reflected in a biased stereotype in favour of one group or event relative to the other.

Table 3 Contingency calculation compared to illusory correlations examples

Table A

Visual presentation of the possible outcome in a contingency paradigm

	Outcome	
Cue	Present	Absent
Present	A	B
Absent	C	D

Table B

Visual presentation of the possible outcome in an illusory correlation paradigm

	Outcome	
Category	Present	Absent
X	2A	2B
Y	A	B

Table C *Zero Contingency*

	Outcome		
Action	Present	Absent	
Present	40 (80%)	10 (20%)	P(O/C)=.8
Absent	40 (80%)	10 (20%)	P(O/~C)=.8

ΔP=.8-.8=0

Table D *Illusory Correlation*

	Outcome		
Category	Present	Absent	
X	40 (80%)	10 (20%)	P(O/C)=.8
Y	20 (80%)	5 (20%)	P(O/~C)=.8

ΔP=.8-.8=.0

A subset of illusory correlations are the phenomena of superstitions, which are the belief that unrelated events are causally connected despite the absence of any plausible causal links. In particular these beliefs focus on faith in magic or the fear of the unknown, or the belief that certain things or events will bring good or bad luck. A more everyday example of superstitious beliefs is the existence of 'lucky' or 'unlucky' items or events. A child may score a record number of goals during a football match when wearing yellow socks, and decide the socks are 'lucky' and wear them for future football matches. Alternatively, a student may fail an exam given on a Wednesday so decide that Wednesdays must be unlucky and avoid taking any more exams on Wednesdays. In both these cases, these events until then were infrequent, however the nature of the events made them distinctive, and due to this a perception of a relationship is formed, which then affects future behaviour.

Illusory correlations have been studied extensively and have been found to be a robust effect (Fiedler, 2000; Hamilton & Sherman, 1989; Rodríguez-Ferreiro & Barberia, 2017; Sanbonmatsu et al., 1987). In general, theories on illusory correlations are based around the use of heuristics in information processing, that is, short-cuts that are often used in decision making. The availability heuristic, the ease with which an idea comes to mind, is thought to result in illusory correlation, as some ideas, or pairings of ideas are vividly and easily remembered even though they do not occur frequently (Plous, 1993). These theories suggest that there is an advantage in memory for infrequent events (Weigl et al., 2018).

The role of working memory capacity has been investigated in relation to illusory correlations in stereotype formation. In general, the strength of illusory correlations differed, when memory load in the working memory capacity was varied, with increased memory load, there was an increase in illusory correlations (Eder et al., 2011). Alternative explanations for illusory correlations, focus on ad hoc inferences, as inferences of the frequency of the outcome is contingent on the frequency of the predictor. Hence when making inferences, the co-occurrence of the predictor and outcome are not required, for people to believe there is a relationship between the two, as the marginal frequencies of these variables are used to make these judgements (Ernst et al., 2019; Fiedler et al., 2009, 2013; Fiedler & Freytag, 2004).

Interestingly, it has been found that, the more you have been a participant in studies and have seen a wider range of paradigms, the less likely you are to report illusory correlations (Murphy et al., 2011). The authors suggested that the illusory correlation may be considered an effect of incomplete learning, opposed to the phenomenon being caused by a bias or heuristic.

Illusory correlations and contingency biases share the notion that people mistakenly perceive a relationship when there is none. Similar paradigms are used in both areas of research as participants are often presented with a series of trials that contain outcome information for groups or cues. At the end of the trials, participants are asked to form a causal or associative judgement based on the information presented to them. In addition to this, due to the nature of the paradigms, zero contingencies are possible in both, hence it is possible to create scenarios where the frequency of events can be manipulated to be high or low, whilst ensuring there is no true relationship between the events or groups. Therefore it is possible to create scenarios where the frequency of events can be manipulated to be high or low, whilst ensuring there is no true relationship between the events, or create scenarios where groups can be large or small while ensuring that the target characteristic is associated with group.

However, there are a few differences between illusory correlations and contingency learning paradigms. Specifically, the explicit distinction between cues and outcomes does not play much of a role in illusory correlations, which rather focus on the alleged co-occurrence of positive or negative attributes. In contrast, contingency learning places more emphasis on the role of cues and outcomes, which in turn provide the basis for judgements about association *and* directionality of an effect. In addition to this, in the typical contingency learning paradigm the frequency of outcomes and cues were not skewed such that one group is a 'majority'. Rather, these are kept the same so that the probability of the outcome without the cue is the same as the probability of the outcome with the cue.

Nevertheless, even though differences exist between these two paradigms, these seem to be mostly due to the different terminology used between areas of research. And as such, at this stage it is uncertain whether the effects and methodology from illusory correlations truly differ from contingency learning.

To summarise, illusory correlations are a form of biased perception of contingency, whereby individuals perceive a correlation or relationship between two variables, due to the skewed distribution of the occurrence of those variables. The infrequency and distinctiveness of some variables and their occurrence leads to the illusory correlations. Illusory correlations are often used to explain stereotype formations and superstitious thinking. Whilst the paradigms used in illusory correlation research are similar to those in contingency learning approaches, cues and directionality seem to play a much larger role in contingency learning.

1.2.3 Spurious correlations

Spurious correlations are distorted perceptions of causality between two variables whereby the correlation does not derive from those two variables. In many cases the presumed causal effect is attributable to the role of a third variable or is an incorrect inference about the presence of a relationship in the population based on sampling variability (Fiedler et al., 2002).

A common example given for spurious correlations, is the increase in drowning when ice cream sales increase. It may seem like ice cream sales lead to more drownings; however, it is more likely that when it gets warmer, more people are likely to buy ice cream, and go swimming. Here there is a third variable which can explain the causality seen between ice cream sales and drowning**.** As another example, the years in which there are more people drowning in pools, seems to positively correlate with the number of films Nicolas Cage appears in, in those years (Vigen, 2015). Whilst these events are correlated, it is also merely a coincidence.

For example, in a study on group stereotypes, participants were asked to judge the intelligence of two groups based on the performance the members of the group on an anagram task (Schaller & O'Brien, 1992). The following information was provided about 25 members of each group attempting an anagram; the group membership of the person attempting the anagram, the difficulty of the anagram, and the outcome of the attempt (success or failure). This information was presented as a sequence of statements, Table 4,

Table 4 Stimulus spread in Schaller and O'Brien (1992)

	Easy		Difficult		Overall	
Group	Success	Failure	Success	Failure	Success	Failure
A	5	0	5	15	10	15
B	15	5	0	5	15	10

after which participants provided judgements of both groups' ability of future anagram tasks. The findings of the study suggest that participants did judge a difference in intelligence between the two groups, with group B being the more intelligent group, which was consistent with the overall performance only. The difficulty of the anagrams was not taken into consideration, even though this was shown to participants.

Overall, Group A correctly completed fewer easy anagrams and was the only group to complete any of the difficult ones, albeit at a high failure rate. Whilst Group B did complete more of the easy anagrams, they also did some incorrectly, and completed none of the difficult anagrams. When this is aggregated it seems like Group B was the more successful group, see Table 4, even though Group A was the only group to correctly complete the harder anagrams. As the participants later judged Group B as being the more intelligent group, this shows that a perceived contingency of group and behaviour can result in stereotypes, even though the correlation is spurious. The perceived contingency is reflected in a biased stereotype in favour of one category relative to the other.

Typically, in studies of spurious correlations, participants are shown a sequence of statements about the members of two groups, with each statement presenting a negative or positive attribute or outcome of that group. Participants are told these attributes occurred in different contexts or vary due to a third factor. However often, when people consider the group memberships, they do not consider the effect of the third factor. This often leads to a perception of a contingency between the groups and the outcomes. The perceived contingency is reflected in a biased stereotype in favour of one group relative to the other.

Spurious correlations are similar to errors or biases that arise in contingency learning; people may mistakenly believe two factors are causally related, when they are not. In terms of the methodology, similar paradigms are used in both areas of research, whereby participants are often presented with a series of outcome information for groups or cues. At the end of the trials, participants are asked to form a judgement based on the information presented to them. Due to the nature of these paradigms, even when there is no relationship between the cue or groups and outcomes presented to participants, they often form a biased perception of contingency, which favours a group or cue.

However, the spurious correlation paradigm also differs from the contingency learning one. Spurious correlations rely on the existence of a third factor that could explain the perceived causal relationship between two variables. For example, a hot day may explain the increase in ice cream sales and shark attacks, as opposed to shark attack and ice cream sales having a causal relationship. Whilst this is not the case in contingencies paradigms, as there is no true causal relationship between the variables and should not be explainable by another factor. Secondly, for spurious correlations there may be a correlation, however there is not causation, whilst for contingency learning, there is neither a correlation nor causation.

To summarise, spurious correlations are a biased perception of causation, whereby causality is judged between two events, that may just be correlated or attributable to another factor. Whilst similar paradigms are used in both spurious correlation research and contingency learning, with spurious correlation there may seem to be a relationship between events due to chance or sampling variability, which is not the case with contingency learning.

1.2.4 Pseudocontingencies

Pseudocontingencies are a form of contingency judgement in which individuals judge a correlation between two uncorrelated variables when they are jointly skewed. To demonstrate this phenomenon, take the example of academic and sporting achievement in two different schools. At school A there may be a high rate of academic and sporting achievement in comparison to school B. However, neither academic and sporting achievement can be observed at the same time, hence they are assessed or seen at different times. However, the base rate information for students from school A, may be that most students are high academic achievers and also high achievers in sports, in comparison to school B. This may then be generalised to suggest that success at academics and sports are positively correlated, even though neither were observed at the same time and were only observed in the context of different schools.

Pseudocontingencies have been studied in a number of different fields from philosophy to economics. Skinner's (1948) experiments on conditioning in pigeons is some of the earliest work within this area. Across experiments Skinner induced a state of hunger in pigeons by reducing the weight of the pigeons to 75% of their normal weight. The pigeons were put in cages where a hopper swung into place at regular intervals and gave out food. Usually, the birds were conducting some behaviour and then the hopper would appear, afterwards the birds would repeat this behaviour for more food. Interestingly the birds started displaying a wide range of conditioned behaviours, hence a range of superstitions were induced in the pigeons.

Within psychology there has been a greater focus on stereotype formation through pseudocontingencies and biases in information sampling. For example, participants were presented information about two towns, X and Y, and a series of desirable and undesirable behaviours of members of two different groups in each town (Meiser, 2003; Meiser & Hewstone, 2004). For each description of a behaviour, the related town context and group membership was also given, hence giving base rate information on group, town and desirability. For town X, group A was a majority and desirable behaviours were more frequent, whilst town Y, group B was a majority and undesirable behaviours were more frequent. The correlation of group and behaviour type were zero or negative for each town, when aggregated across towns there was a positive correlation. Participants gave judgements of estimation of desirable and undesirable behaviour frequency across both towns for each group, which showed a more positive evaluation of group A compared to B, despite the zero correlation between group and behaviour within each town (Meiser, 2003; Meiser & Hewstone, 2004).

In the typical paradigms for examining pseudocontingencies, participants are presented with information on two or more different characteristics of two or more groups, either jointly or successively in a context. Afterwards participants make predictions using partial information of a particular group characteristic, to predict what level or value this group member may have. The presentation of variable information can vary, either being presented simultaneously or successively. In simultaneously presented information, variables were presented together, most often summarised in tables. Alternatively, in successive presentation the group characteristics are presented one by one, over several trials.

1.2.5 Biases of sequential dependence

Biases of sequential dependence consider that individuals hold beliefs or assumptions about the underlying mechanisms of a sequence which affects how one interprets said observations and future judgements. The two commonly known examples of this are the hot-hand fallacy and gamblers fallacy. It is thought that these fallacies occur due to manifestations of the representative heuristic, as they are judgements formed about the probability of an event under uncertainty. When people rely on representativeness to make judgments, they are likely to judge wrongly because the fact that something is more representative does not actually make it more or less likely.

The hot hand fallacy occurs when individuals believe in 'streaks' in sequences of hits and misses or belief in positive recency in events (Green & Zwiebel, 2017). This leads to individuals believing that once a random event has occurred, it has a greater probability of occurring in succession. For example, having observed 5 patients in a row for whom drug Z worked, someone may now believe it more likely than usual that it will work for the next patient. Most of the hot hand fallacy is seen and studied in sports. The Hot hand fallacy was first described in relation to the 'hot hands' basketball players were thought to have. It was suggested that when basketball players had made a successful shot, the succeeding shots were more likely to be successful as well (Gilovich et al., 1985). The effect has been found in other contexts such as finance (Hendricks et al., 1993), sports (Clark, 2003; Dorsey-Palmateer & Smith, 2004), betting markets (Camerer, 1989), and prediction tasks like roulette games (Croson & Sundali, 2005).

It has been suggested that the streaks associated with the fallacy do exist in some sports such as darts, golf, and horse shoe pitching where wins and misses can occur in streaks (Ayton & Fischer, 2004; Gilden & Wilson, 1995a, 1995b). However, this seems to be due to skill and confidence, which may have led individuals to mistakenly expect clumps of the same events to occur in other domains of human performance as well (Ayton & Fischer, 2004).

The gamblers fallacy occurs when individuals see patterns where they do not exist, a form of negative recency in events, that is, a belief that the probability of independent events occurring decreases or increases depending on recent events, and that these will self-correct over time (Burns & Corpus, 2004). If they believe an event has happened more infrequently recently, then they will expect that event to happen frequently in the future. For example, if a fair coin is flipped 5 times in a row, with heads appearing all 5 times, the probability of flipping heads the following turn is 1/2. This will not be dependent on what the previous turns resulted in. However, those under the gambler's fallacy may believe that due to heads appearing for 5 turns consecutively, then tails is due to occur on the 6th turn.

Perhaps the most famous example of the gambler's fallacy occurred in a game of roulette at the Monte Carlo Casino in 1913, during which the ball landed on black 26 times in a row, which has a probability of around 1 in 136.8 million, this was an extremely uncommon occurrence. Due to the incorrect reasoning that the streak was causing an imbalance in the randomness of the wheel, hence had to be followed by a streak of red, gamblers lost millions of francs by betting against black (Owen, 2011). The fallacy occurs in a number of settings that range from when people choose lottery numbers (Clotfelter & Cook, 1993) to actual gambling situations such as blackjack or roulettes at the casino to race tracks as well (Croson & Sundali, 2005; Keren & Wagenaar, 1985). It has also been found that gambler's fallacy does occur outside the gambling setting, for example baseball umpires, loan officers, and asylum judges, have been found to employ the gambler's fallacy consistently in their decision making (Chen et al., 2016; Suetens et al., 2016).

It is thought that the gambler's fallacy arises out of a belief in the law of small numbers - the belief that the balance of random outcomes that occurs in very large samples will also be reflected in small samples, streaks must eventually even out to be representative (Burns & Corpus, 2004). Hence suggesting the gambler's fallacy is a form of representative heuristic, whereby people assess the probability of certain events by judging how similar the events are to previous events they have experienced and the similarity of the surroundings of the two events (Tversky & Kahneman, 1971, 1974).

In a typical study for both of these fallacies, participants are asked to view a series of outcomes, and then asked to predict the outcomes of the succeeding events before they occur. However, in contingency paradigms participants are asked to make a judgement of association, based on cue and outcomes, rather than prediction of outcome based on previous outcomes. As such cues play a bigger role in contingency paradigms compared to these fallacies. However as outlined earlier, many contingency learning tasks examine cue effects by asking participants to predict the outcome based on the cue on a trial by trial basis.

1.2.6 Confirmation Bias

Confirmation bias is the tendency for individuals to search, interpret, or recall information such that it confirms or supports their prior beliefs (Nickerson, 1998). In practice this may occur, when people interpret ambiguous evidence, discard, or ignore contradictory information so their beliefs are further supported. Confirmation bias is stronger for emotionally linked beliefs, desired outcomes, and deeply ingrained beliefs (Levine et al., 2001; Safer et al., 2001). It is also thought to offer an explanation for why some false beliefs form or persist. That is, if it is not a type of false belief about association but a reason why such false beliefs sometimes arise.

Confirmation bias is seen in many situations. For example, many individuals with strong political beliefs may choose news sources which also align with their beliefs. The more individuals use the news source, the stronger the confirmation bias, as they are only searching for information that supports their beliefs. With so many news sources available, it is possible to find a new source that will support any belief. With the increase in social media, this effect can be amplified, as due to algorithms showing individuals stories or information that is similar to what the individuals already follow, whilst excluding contrary information.

Explanations or the mechanisms through which confirmation bias occurs, fall predominantly under three categories: biased search for information, biased interpretation of information, and biased memory recall (Lord et al., 1979). In addition to this, confirmation bias is thought to explain four specific effects (Althubaiti, 2016; Hart et al., 2009): illusory correlations, irrational primacy effect, belief perseverance and attitude polarisation.

1.2.7 Halo effect

The halo effect is a cognitive bias that occurs when people have a positive impression of a person, or company in regard to one attribute, which positively influences their feelings or opinions of the given person or company on another attribute (Thorndike, 1920). Within decision making research, this phenomenon is often found to occur when evaluators give more positive judgements based on their previous judgements, thus using unrelated information to make their evaluations (Gibson & Gore, 2016; Nisbett & Wilson, 1977). This effect can occur in the negative direction as well, whereby if an individual dislikes one aspect of something, they will have negative predisposition toward everything about that object or person (Nisbett & Wilson, 1977).

The halo effect is a common bias that occurs in everyday life. For example, in the workplace, one may see a co-worker or client who is smartly dressed, and thus assume they are more competent at their work than their colleague who wears a t-shirt to work. Another example is the ease with which often we base our overall impression of celebrities based on one or two of their positive attributes, such as seeing celebrities being successful or attractive, and then also thinking of them as being intelligent and kind.

In a typical halo effect study, participants often given a positive characteristic of an individual, group or company. Afterwards they provide evaluations of said individual on other characteristics. The halo effect has been studied in many contexts such as intelligence (Landy & Sigall, 1974; F. R. Moore et al., 2011; Talamas et al., 2016), politics (Berggren et al., 2017; Poutvaara et al., 2009; Surawski & Ossoff, 2006; Verhulst et al., 2010), marketing (Burke et al., 2018; Vance et al., 2016), and judiciary systems. (Dahl-Monroe, 2017; Efran, 1974; Sigall & Ostrove, 1975). For example, research within the jury system, had found that the halo effect to influence judgements on sentencing a great deal. Individuals are often more lenient in sentencing attractive people over unattractive ones even if they have committed the same crime (Efran, 1974). The authors suggested this may be due to the societal perception that attractiveness is related to successful future. Due to the extensive study of the halo effect, the paradigms and methodology used vary a great deal, however as the effect is found with such consistency it is considered to be a robust effect.

The halo effect is similar to contingency learning, as a bias in perception occurs which then affects the future judgements individuals make. Importantly it is a bias in perception of co-occurrence of variables, attractiveness is perceived to be negatively associated with culpability, when it is likely not to be associated. However, an immediate or first impression, is what biases individuals with the halo effect, rather than the experience from a series of absent or present cue and outcomes. In addition to this, in a contingency a biased perception of association is made, as opposed to just a bias in perception as made due to the halo effect.

To summarise, the halo effect occurs due to a positive first impression influencing future judgements, on unrelated matters. Whilst the halo effect is a form of biased perception, it differs from contingency judgements, as they are based on cue-outcome associations.

1.2.8 Discussion

To summarise, a number of different types of causal beliefs have been discussed, in which people may form false beliefs, which do not accurately reflect the underlying relationship between events or objects. To note, judging contingencies is different to judging causality. Whilst contingency judgements are made based on observations of the probability of an outcome following on from particular cue, such as the probability of B occurring is highest after A occurs. Causality refers to the direct influence of an event causing another event, such that C has caused D. Inferring causality is a major challenge as the evidence on which these inferences are made can be faulty or incomplete thus leading to incorrect reasoning hence errors arising in this process are to be expected. In this thesis, I am interested in situations in which people infer a relationship between a cue and outcome, when there is no relationship. A contingency approach has been taken up, and the contingency between cue and outcome is explored across this thesis. The study of contingency may shed further light on the understanding of causality. To be able to examine the contingency between cue and outcome, all the outline causal beliefs cannot be examined, particularly as in some cases there may be some underlying factor which links the cue and outcome. As such in this thesis, the illusion of causality is examined, as the illusion arises when there is no relationship between the cue and outcome, however people tend to believe there is some relationship between the cue and outcome.

1.3 Illusions of causality

Successfully identifying associations between events is fundamental to cognition (Shanks, 2010). This form of pattern detection is a basic building block of the learning required for adaptive functioning, ranging from learning which words represent which objects or ideas, to learning which outcome(s) reliably follow a specific type of action (Batterink & Neville, 2011; Shanks, 1993; Shelton & Martin, 1992; Werker et al., 1998). Learning theorists have therefore rightly been concerned to understand whether and how associations and causes can be accurately identified—and to identify any conditions that hinder people’s ability to effectively gauge the strength of an association or causal path between a cue and a related outcome. One such hindrance relates to the illusion of causality.

Although studies of action–outcome and cue–outcome, where action is when someone or something acts and cue judgments have generally found that people are sensitive to the strength of positive and negative contingencies, judgments of causation consistently deviate from the Δp rule when there is no contingency between the potential cause and the outcome (i.e., $\Delta p = 0$) (Allan & Jenkins, 1983; Shanks & Dickinson, 1988). Effects such as this, where the attribution of causal relationships between two non-contingent events, when ΔP is zero, are commonly referred to as an illusion of causality (Matute et al., 2015). Illusion of causality is an important phenomenon because it represents a consistent error in human learning that is thought to contribute to the development and maintenance of superstitious beliefs and pseudoscientific thinking (Matute et al., 2015). A framework through which the illusion of causality may be examined is the contingency account.

1.4 Theoretical frameworks for contingency learning

Below the contingency learning account is outlined along with the associative perspective and probabilistic contrast model. This thesis has taken a contingency approach for the experimental paradigms. However, the associative and probabilistic models may offer alternative or better explanations for the effects examined in this thesis, compared to the contingency account. As such in the discussion chapter the results will be compared to these two models as well.

1.4.1 Contingency learning

The contingency learning account is a statistical theory of learning, originally developed to explain how people learn the underlying contingencies between two or more events (Pineño & Miller, 2007). A contingency refers to the covariation between two or more variables, when each variable only has two possible values (i.e., present vs absent). In typical contingency experiments participants are shown or given instances of two or more events. The individual events may be present or absent on each trial. After a set of trials participants are asked to rate the contingency between the event, or some form of causality between the events using a numerical scale. The higher the judgement, the higher the perceived contingency by participants.

Whilst there are many models of contingency, the ΔP contingency account by Allan, (1980) is one of the simpler models (Shanks et al., 1996). In the ΔP account information regarding the cue and outcome is acquired according to the 2 by 2 contingency table, Table 1. Each cell represents whether the cue and outcome are each present or absent, giving 4 possible combinations. The contingency learning account states that the belief about events should correspond with the difference in the probability that two events co-occur, *P*(Event 1|Event 2), and the probability that one of these events occurs in isolation, *P*(Event 1|No Event 2). This difference is quantified as Δ*P*—the degree of contingency (Allan, 1980). As such Δ*P* can range from -1 to +1. For both positive and negative Δ*P* values there is some relationship between cause and outcome, whereas if the potential cause has no total relationship on the outcome, then Δ*P* is zero. For a positive contingency the probability of the outcome occurring is higher with the presence of the cue. For a negative contingency the outcome is more likely to occur in the absence of the cue.

Research has found that participants' ratings of contingencies often match the Δ*P* resulting from different probabilities of cue and outcome (Wasserman et al., 1990, 1996). However contradicting results have been found, and the contingency account has been found to fail to take into account effects such as cue competition.

Whilst the contingency account may not be a perfect account of human contingency learning, the model is simple in its explanation of contingency learning. In particular it focusses on belief formation based on binary variables. Also the account is not made complex by considering causality, but simply refers to the covariation between events. As this thesis is not focusing on causality, but how people form beliefs based on events they experience or witness, the contingency account is appropriate here. In addition to this, the paradigms based on the contingency account allow for the examination of belief formation in ways that are intuitive for participants to grasp. Also, the paradigms can be easily changed to reflect different contexts as well.

For all the experiments that use zero contingencies (i.e., Δ*P* =0), according to contingency learning theory, there would be an expectation that contingency judgements should be zero. Moreover, even if participants had a response bias when providing contingency judgments (e.g., an aversion to using the scale endpoint representing zero contingency) the contingency learning account cannot explain *differences* in judgment between conditions with identical Δ*P* between the high- and low-density conditions examined in this thesis. That is, across experimental manipulations, where the relationship between cue and outcome remained unchanged (at Δ*P* =0), no systematic differences in contingency judgements should arise.

However Wasserman (1990) found that people tend to weigh the cells of a contingency table (Table 1) according to, A > B > C > D for judgements of data presented in contingency tables. In the study participants weighed the 2 by 2 contingency tables to give contingency judgements. And found that people tend to weigh cell A trials the most followed, cell B and C trials, and cell D trials are given the least weighting. For example, this implies in the medication treatment scenarios used in this thesis, participants will give most weight to occurrences of patients given medication and recovering, followed by patients given medication not recovering, then patients no given medication and recovery, and the least weight to patients who do not receive medication and do not recover. In addition to this, it was also suggested that the trials of cells A and D lead to increased positive judgements whilst cell C and B counts decrease contingency judgements. Taking the medical treatment scenario again, when occurrences of patients given medication and recovering, along with not given medication and not recovery, results in more positive judgements. As such, patients given medication not recovering, along with patients no given medication and recovering decrease judgements. For instance, looking at Table 5 this is demonstrated, as increases in cells A or D, increases the positivity of the contingency, whilst increasing cells B and C decreases the positivity of the contingency. With this extension added to the contingency account, the results of this thesis may be better explainable by the account.

1.4.2 Associative perspective

The associative perspective considers the associative strength of two events, the conditioned stimulus (CS) and unconditioned stimulus (US). The Rescorla and Wagner (1972) model is one of the more prominent associative models, which focus on the formation and strengthening of cue and outcome associations. For the purposes of comparison here the CS can be considered similar to a cue, whilst the US can be considered as an outcome. When these two events are paired it leads to an internal connection between the events. Strength of connection is determined by how often the events are paired together, how often events occur in isolation and how often both are absent. It is this mechanism that allows the model to be sensitive to CS–US contingencies in animal and human learning.

Associative strength is determined by $\Delta V^{n}_{cue}=\alpha \cdot \beta \cdot (\lambda - V^{n-1}_{Total})$. In the equation ΔV^{n}_{cue} refers to the change in the associative strength of the cue on any given trial, *n*. There are two learning rate parameters, α and β, which represent the associability (salience) of the cue and outcome respectively. The learning paraments are bound between 0 to 1, and function according to their values of salience. The final term $(\lambda - V^{n-1}_{Total})$ represents the discrepancy between the amount of associative strength of the outcome λ, and the current

Table 5 Wasserman et al's. (1990) explanation of how cells A, B, C and D change contingency

Table A Comparison contingency calculation

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C) = .5
Absent	10	10	P(O/~C) = .5

ΔP = .5-.5=0

Table B Increased cell A trials

	Outcome		
Cue	Present	Absent	
Present	15	10	P(O/C) = .6
Absent	10	10	P(O/~C) = .5

ΔP = .6-.5= 0.1

Table C Increased cell B trials

	Outcome		
Cue	Present	Absent	
Present	10	15	P(O/C) = .4
Absent	10	10	P(O/~C) = .5

ΔP = .4-.5= -0.1

Table D Increased cell C trials

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C) = .5
Absent	15	10	P(O/~C) = .6

ΔP = .5-.6= -0.1

Table E Increased cell D trials

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C) = .5
Absent	10	15	P(O/~C) = .4

ΔP = .5-.4= 0.1

total associative strength that has been acquired until trial n-1, by all the cues present on trial n, (V^{n-1}_{Total}). The value of λ, is dependent on whether the outcome is present (1) or absent (0).

Experimental context also raises associative strength and can compete with cues to predict outcomes (De Houwer & Beckers, 2002). These models assume learning occurs on a conditioning trial only to the extent that the US is surprising. On the first conditioning trial when a CS is first paired with a US, as nothing signals the US, the US is surprising and therefore learning occurs. As conditioning trials proceed, CS will come to predict the US and the US will become less surprising. At some point the CS will predict the US perfectly and at this point no further learning will occur (Miller et al., 1995).

1.4.3 Probabilistic contrast models

The probabilistic contrast model is another rule-based account similar to the contingency accounts. According to the model, individuals reason about cause-effect relationships by recalling all the events of the contingency and calculating a similar statistic to Δp (Cheng & Novick, 1990). Individuals compare situations that differ due to the presence of a target cue, whether the cue is present or absent. By comparing these situations inferences about the relationship between cue and outcome can be learnt. Causal power is driven by the idea that for any cause of an effect, there are always potential alternatives causes. So, knowledge of the base rate of the effect (frequency of outcome) will let one make a contrast between the candidate cause and alternative causes such as experimental context (Buehner et al., 2003; De Houwer & Beckers, 2002). Base rate is 1 when the outcome always occurs when the cue is absent. Base rate is 0 when the outcome never occurs when the cue is absent. The following equation is used by the model, $Pi = \frac{\Delta p}{1-P(e/\bar{\imath})'}$.

In this equation, e, is the occurrence of the outcome and, i, is the occurrence of the event. When the cause is preventative, the denominator is P(e/ i). So, the power for an event to cause an outcome is a joint function of the overall contingency (P) and the probability of the outcome in the absence of the event.

1.4.4 Summary

In each empirical chapter the extent to which the results are predicted by the above accounts will be discussed.

1.5 Manipulations in Contingency learning research

1.5.1 Outcome density

Unfortunately, people do not always behave according to the normative contingency account; contingency beliefs are often biased. People tend to believe that events that occur often must somehow be related, even if they are not (Allan & Jenkins, 1983). One of these effects is known as the outcome density effect—the tendency to make an association between cue and outcome due to high frequency of outcome when there is no actual association between cue and outcome (Allan et al., 2005; Matute, 1995; Musca et al., 2010). When the contingency is held at zero then the probability with which the outcome appears may bias judgments such that a higher probability of the outcome occurring may lead to higher contingency judgements (Alloy & Abramson, 1979; Shanks & Dickinson, 1991; Vallée-Tourangeau et al., 2005). Hence, higher outcome densities may bias beliefs. Taking Table 6 for examples of high and low outcome densities, the ΔP for both is zero, as such there is no relationship between the cue and outcome. The only difference between both calculations is the probability of outcome occurring. It is often found that in when outcome density is higher, participants are more likely to judge a contingency between the cue and outcome.

Alloy and Abramson (1979) conducted one of the first outcome density studies using a modified paradigm previously introduced by Jenkins and Ward (1965). In a series of contingency problem trials, participants were tasked with turning on a green light as often as possible by either pressing a button or not pressing a button, which would or would not result in a green light, in any combination. These researchers manipulated both the contingency between button pressing and the light turning on, as well as the overall frequency of the light turning on – referred to as the outcome density. Regarding the contingency level, in some conditions, the probability that the light would turn on after pressing the button was higher than chance and the probability that it would turn on after not pressing the button was lower than chance. This corresponds to a positive contingency between button pressing and light turning on. In other conditions, this probability was equal to chance, equivalent to no contingency. Regarding the outcome density, in some conditions

Table 6 outcome density at Zero Contingency

Low outcome density calculation

	Outcome		
Cue	Present	Absent	
Present	25	75	P(O/C)=.25
Absent	25	75	P(O/~C)=.25

ΔP=.25-.25=0

High outcome density calculation

	Outcome		
Cue	Present	Absent	
Present	75	25	P(O/C)=.75
Absent	75	25	P(O/~C)=.75

ΔP=.75-.75=0

the probability on any given trial of the light turning on was higher than 50%, whereas in others it was not, representing high and comparatively low density, respectively.

Participants took part in a total of 40 trials, after which they provided a judgement of control. An outcome density effect was found, whereby outcomes were judged to be more controllable when outcome density was high rather than low, even though there was no contingency between the cue (button press) and the outcome (green light)[1].Subsequently, the outcome density effect was further examined and found to be robust across variations of the original paradigm (Blanco et al., 2013; Crump et al., 2007; Gillan et al., 2014; Matute et al., 2007; Msetfi et al., 2005; Vallée-Tourangeau et al., 2005). These variations include an allergy paradigm in which the participant assesses which foods cause allergy (Blanco et al., 2013) and a medical treatment paradigm in which the participant judges the effectiveness of a treatment for a disease (Blanco et al., 2011).

Chapter 3 of this thesis reports a study which tested the outcome density effect with zero contingency in a passive contingency task.

1.5.2 Cue density

When considering density effects and their influence on associations, the focus has been predominantly on outcome density effects (Musca et al., 2010; J. C. Perales et al., 2005). Contrastingly, cue density effects (e.g., treatments for diseases) have not been extensively examined. Cue density effects are another type of causal illusions that can also be considered focussing on examining the formation of beliefs. Cue density effects occur in

[1] This original study shows that the tendency to ascribe cause when none existed was diminished in a sub-group of participants with depression—a finding which some describe as a demonstration of ‘depressive realism’

zero contingencies when the probability of cue is manipulated to be high or low and participants judge a greater contingency between the cue and outcome in the high condition compared to the low condition (Table 7 for comparison of contingency calculations between high and low cue and outcome densities). As such the density of cues may also influence perceived associations, and in a similar way to outcome densities when making judgments (Blanco et al., 2013).

Taking the example of allergy paradigm, much is known about how judgements changes when a constant number of medications is taken, but the disease outcome varies. However, little is known about what occurs when the medication taken varies, but the outcome remains constant. In the real world such variations in cue may occur for example in differential treatment strategies between hospitals or doctors, where the treatment (cue) frequency may vary, however recovery (outcome) may stay relatively consistent across settings. The cue density, also known as cause density, has been studied comparatively less than the outcome density. The methodological similarities and differences of cue density studied compared to outcome density studies are summarised below, along with the findings from cue density studies.

1.5.2.1 Methodological similarities or differences between outcome density only studies and cue density only studies

Often when cues are manipulated, it is to see whether the manipulation leads to differences in predicting the outcome. Hence many cue studies often include an additional question per trial, that appears after the cue, asking the participant whether they think the

Table 7 Low- and High-density calculations, for cue and outcome density at Zero contingency. When cue is manipulated outcome probability is constant, when outcome is manipulated cue probability is held constant

Table A Low cue density at zero contingency

	Outcome		
Cue	Present	Absent	Outcome density
Present	10	10	P(O/C)=.5
Absent	40	40	P(O/~C)=.5
Cue density	P(C/O)=.2	P(C/~O)=.2	
	Cue Δp=.2 - .2=0		
	Outcome Δp=.5 - .5=0		

Table B High cue density at zero contingency

	Outcome		
Cue	Present	Absent	Outcome density
Present	40	40	P(O/C)=.5
Absent	10	10	P(O/~C)=.5
Cue density	P(C/O)=.8	P(C/~O)=.8	
	Cue Δp=.8 -.8=0		
	Outcome Δp=.5 - .5=0		

Table C *Low outcome density at zero contingency*

	Outcome		
Cue	Present	Absent	Outcome density
Present	10	40	P(O/C)=.2
Absent	10	40	P(O/~C)=.2
Cue density	P(C/O)=.5	P(C/~O)=.5	
	Cue Δp=.5 - .5=0		
	Outcome Δp=.2 - .2=0		

Table D *High outcome density at zero contingency*

	Outcome		
Cue	Present	Absent	Outcome density
Present	40	10	P(O/C)=.8
Absent	40	10	P(O/~C)=.8
Cue density	P(C/O)=.5	P(C/~O)=.5	
	Cue Δp=.5 - .5=0		
	Outcome Δp=.8 -.8=0		

outcome will follow or not in a yes or no format. Once participants have answered this, the outcome may or may not follow (Matute et al., 2011a). As such participants are making predictive judgements after each trial, rather than a causal judgement. After the set number of trials, participants typically make their causal judgement on a numerical scale, whereby they judge the strength of the causal link between the cue and the outcome.

The few cue density studies that exist focused on the difference in outcome predictions and causal judgments. Outcome density and cue density effects have often been found in studies assessing causal judgements, however these effects have not always been found in outcome predictions, whereby the predictions are made every trial (Allan et al., 2005; J. C. Perales et al., 2005). As causal judgements are made at the end of a learning phase after multiple trials a more consolidated perception of causality can be made. However, with prediction judgements as they are made trial by trial the judgement can change from trial to trial and occurs with more limited evidence. As such, as predictive judgements are not essential in determining cue and outcome density effects and does not guarantee evidence for cue or outcome effects, this approach is not being carried forward within the experiments of this thesis.

1.5.2.2 Cue manipulation only

In Vadillo et al’s (2011) study, cue density was manipulated to be either low or high. The standard allergy contingency paradigm was changed to use space aliens and carrots in a between subjects’ paradigm. The cue density bias was found to be stronger for outcome predictions compared to causal judgements. Specifically, they found that people believed more strongly that a substance was an allergen when the substance in question featured in more than 50% of learning trials, despite the contingency between substance and allergic reaction to be the same irrespective of the overall frequency of the substance.

Similar results were obtained in a study using the same original standard allergy paradigm (Matute et al., 2011a), with high outcome density, and manipulated high and low cue density. Participants were asked at the end of 100 trials, ‘To what extent do you think that Batatrim is the cause of the healings from the crises in the patients you have seen?’ (Causal question), and ‘To what extent do you think that Batatrim has been effective in healing the crises of the patients you have seen?’ (Effectiveness question). The effect was present in both types of questions and both conditions, however it was strongest in the effectiveness question for both high and low cue density.

Yarritu and Matute (2015) used a passive contingency task to examine the effect of cue density and prior expectation on judgement. One hundred and fourteen participants took part in a fully between-subjects experiment, using the same cover story as before. The probability of cue was manipulated to be high (0.8) or medium (0.5), and half the participants were presented with trials in which the cue occurred with a probability of 0.8, whilst the other half experienced trials with the cue occurring with a probability of 0.5. The probability of the outcome was 0.8 across all conditions. During the training phase, for each trial participants were asked a predictive question regarding whether the participants would recover from the illness, which was followed by the outcome. After all trials participants gave contingency judgements. There was a main effect of cue such that participants in the high cue probability condition gave higher judgements than in the medium density condition. There was no main effect or interaction effect for prior expectation.

1.5.2.3 Cue and outcome manipulations

Perales et al. (2005) extended this earlier work by manipulating not only cue density, but simultaneously varying whether the contingency was null or positive. Whilst outcome density varied between contingencies, it was held constant within each contingency. An effect of cue density was found, with judgements being higher in the high-density condition compared to the low-density condition. Similarly, an effect of contingency was found with the higher judgements in positive contingencies compared to null contingencies. Overall, the strongest causal judgements were made in high cue density positive contingency condition.

Few studies have simultaneously examined cue density and outcome density effects. In one such rare investigation, Blanco et al. (2013) used an allergy paradigm with zero contingency between cue and outcome, and employed a 2 × 2 design with cue density (high vs low) and outcome density (high vs low) to examine the interactive effects of cue and outcome. They found an effect of outcome density in all conditions, however, there was only a small effect of cue density, and only when outcome density was high.

1.5.2.4 Other findings

Within contingency paradigms, the densities of cues and outcomes are typically not manipulated within the same studies. However, a small number of studies have shown that when cues are manipulated to be on the same screen as the outcome or to disappear before the outcome appears, there is an effect of cue density (Musca et al., 2010). Participants are more likely to make higher causal judgements when the cue and outcome are on screen together, than when they are not even if the outcome does follow the cue.

Overall, the cue density effect occurs both when cue and outcome are seen together and when cue and outcome are in sequence, but the effect is stronger when they are seen together. When the cue and outcome occur together on screen it would be clear to see an association between the two events. However, it may also cause confusion in terms of the clarity between which event is the cue and which event is the outcome. In addition to this, in real life situations cues and outcomes do not often occur together at the same time, in fact there tends to be a time delay between the cue occurring which then results in an outcome.

1.5.2.5 Summary of cue density

Overall, the cue density effect has not been studied as extensively as the outcome density effect within contingency learning or using contingency paradigms. However, these studies do often follow a similar format to outcome density studies. An effect of cue density has been found across many studies, however the effect seems to be more prominent in the presence of a higher outcome density, when cue density is also high.

Chapter 4 of this thesis reports a study which tested the difference in effect size of cue and outcome density effects with a zero contingency in a passive contingency task.

1.6 Presentation Order effects

The next section will focus on the order in which the cue and outcome are presented within any individual trial, and whether the size of the cue and outcome density effect changes depending on whether the cue or outcome is presented first. Unfortunately, there is not a great of research on this particular area, as such the literature on trial order information (primacy and recency effects) are discussed, which are then linked to cue and outcome presentation order.

1.6.1 Trial order effects - Primacy and Recency effects

The order in which information is presented and its influence on decision making has been studied in several ways across a number of different disciplines. Two key effects that have emerged from this research are the primacy and recency effects.

The primacy effect occurs when presented with a long series of information, individuals focus more on the initial information presented to them, at the expense of the subsequent information (Marshall & Werder, 1972; Rundus, 1980). Individuals also tend to weigh the initial information more in a series of information use this as an anchor from which to make their judgement. For example, within the evaluation and judgements of juries, it has been found that when opposing arguments of equal persuasiveness are similar in all factors of importance, the first presented argument has the most sway with the juries (Lawson, 1968).

Yates and Curley (1986) examined these effects in contingency paradigms, to determine whether contingency judgements were affected by the order of information presentation across trials. The order of the trials was manipulated in such a way that trials presented early indicated a contingency in the opposite direction of the trials presented later. Half the participants were informed they would be required to recall the frequencies of the trials they saw, whilst half the participants were not. After a series of trials, participants recalled the frequencies of the trial types they saw. The primacy effect was present for the participants who had not been informed of the recall test compared to those who had been informed. Similar findings have been found linking the primacy effect to contingency judgements (Dennis & Ahn, 2001; Marsh & Ahn, 2006).

Contrastingly, the recency effect reflects a tendency for individuals to be more likely to remember the more recently presented information in comparison to other information presented at the start or the middle of a series of information. Across four experiments the recency effect was tested in human contingency judgments (López et al., 1998). Similar to Yates and Curley (1986), the trials of these four experiments were manipulated such that early trials indicated a different contingency to the later trials. The results showed a significant effect of trial order, in particular a recency effect was seen.

However, combined these two effects and studies provide contradictory accounts of contingency judgements. A series of four experiments on contingency judgements conducted by Collins and Shanks (2002) found effects of both primacy and recency on contingency judgements. Participants were assigned to either make a contingency judgement at the end of 80 trials or make incremental judgements after every 10 trials. The final contingency judgements participants made were influenced by primacy when participants integrated information from across trials. When participants made incremental judgements recency effects were found.

Overall, the evidence shows that the order in which information is presented across trials may influence contingency judgments. However, little is known about how judgements may be influenced when the presentation order of cue-outcome pairings is manipulated within trials.

1.6.2 Cue outcome order in contingency learning research

Contingency learning paradigms tend to manipulate the densities of cues and outcome independently, such as Blanco et al. (2013). Hence, insights into interactions between cue and outcome are minimal. However, studies such as that by Musca et al. (2010) has shown that when cues are manipulated to be on the same screen as the outcome or to disappear before the outcome appears, there is an effect of make higher causal judgements when the cue and outcome are on screen together, than when they are not even if the outcome does follow the cue. However, studies within this area have not explored the order in which cues and outcomes are presented, for example where participants are presented with the outcome before the cue, and then asked to make causal judgements.

However, Musca et al. (2010) did not manipulate the order of cue and outcome presentation, so while their findings may be interpreted in line with outcome and cue density effects, it is conceivable that these effects actually relate to the order of event occurrence, that is, manipulating the density of the 'second event' may have a stronger effect than manipulating the density of the 'first event.'

Shanks and Lopez (1996) examined the effect of manipulating order of cue and outcome presentation in regard to cue selection. Cues were manipulated to be abstract or concrete. Participants were able to form causal relationships between the cue and outcome, regardless of which way around the cue-outcome pairs were shown. However, this experiment used more than one cue on trials, and this was an active task in which participants gave predictive responses after each trial. At present, the literature does not seem to consist of any studies in which the order of cues and outcomes, in a series of cues and outcome pairings, are manipulated to be shown in a different order (such that the outcome is presented before the cue) in a passive contingency paradigm. It is therefore unknown how this may affect judgements in a passive contingency paradigm.

1.6.3 Summary

To summarise, when considering the effects of order the research has focused on the effect of trial order on judgements, particularly the primacy and recency effects. Within the contingency literature one study has examined the effect of the outcome appearing after the cue, with the cue staying on screen, in comparison to the cue appearing, then disappearing which is followed by the outcome. Whilst another study has examined effects on cue-outcome presentation order within an active task.

Chapter 5 of this thesis reports two studies which tested the effect of cue and outcome interchangeability with zero contingency in a passive contingency task.

1.7 Contingencies

With zero contingencies, within any blocks of trials, the relative frequency of the outcome does not vary between cue-present and cue-absent trials. The literature has shown that with higher probability of outcome, the more positive judgements become, even though within the zero contingency condition there is no relationship between the cue and outcome (Crump et al., 2007; López et al., 1998; Msetfi et al., 2005; Vallée-Tourangeau et al., 2005). The zero contingency condition will be one focus for this project, as this is the condition in which people often form (categorically) false beliefs, such that cues, and outcome are associated. However, I am also interested to see how judgements of cue-outcome pairings vary in positive and negative contingencies compared to zero contingencies. Within a positive contingency, a cue does predict the appearance of the outcome, meanwhile in a negative contingency the cue predicts the absence of the outcome. See Table 8 for contingency calculation, for examples of contingency calculation for zero, negative and positive contingencies for outcome density. The tables show a total of 40 trials per contingency set – yet depending on the number of trials per cell type, the overall contingency varies across the sets. Increasing cue-no outcome and no cue-outcome pairings results in the negative contingency, whilst increasing cue-outcome and no cue-no outcome pairings results in a positive contingency. Overall, in a positive contingency the probability of outcome occurring with cue is higher than that the probability of outcome occurring without the cue, however the opposite is true for negative contingencies. Within both positive and negative contingencies, there is an association between the cue and outcome, unlike a zero contingency. However, I am interested to see how type of contingency can influence judgements.

Table 8 Contingency calculations for outcome density across zero, positive, and negative contingencies

Outcome density at Positive Contingency

	Outcome		
Cue	Present	Absent	
Present	12	8	P(O/C)=.6
Absent	8	12	P(O/~C)=.4
	ΔP=.6-.4=.2		

Outcome density at Zero Contingency

	Outcome		
Cue	Present	Absent	
Present	8	12	P(O/C)=.4
Absent	8	12	P(O/~C)=.4
	ΔP=.4-.4=0		

Outcome density at Negative Contingency

	Outcome		
Cue	Present	Absent	
Present	8	12	P(O/C)=.4
Absent	12	8	P(O/~C)=.6
	ΔP=.4-.6=-.2		

Within the contingency paradigm literature, the influence of positive and negative contingencies has also been studied, positive contingencies more extensively than negative contingencies. Crump et al. (2007) used a 2 (zero vs positive contingency) by 2 (low vs high outcome density) by 2 (contingency ratings vs frequency rating) within subjects design to examine an adapted version of the contingency paradigm with 37 participants. In the paradigm participants viewed a stream of 60 cue-outcome pairs per trial, with 20 trials per block. After each trial participants were either asked to provide a frequency rating or a contingency rating, hence participants complete 10 frequency ratings and 10 contingency judgements per block. Main effects of both contingency and outcome were found, participants gave higher judgements in the positive contingency compared to the zero contingency, and participants gave higher ratings for the high outcome density compared to the low outcome density. An interaction effect of contingency by outcome density was also found, such that the difference between high and low outcome density conditions was larger for the zero contingency than positive contingency condition. The experiment suggests that participants do provide a more positive judgement in positive contingencies compared to zero contingencies.

Perales et al. (2005) replicated a similar effect in an experiment which looked at cue density effects, in a 2 (between subjects: zero vs positive) by 2 (within subjects: high vs low density) design and collected predictive responses after each trial, and contingency judgements at the end the task from 44 participants. The study found main effects of contingency; participants gave higher judgements in the positive contingency compared to the zero contingency.

Cavus and Msetfi (2016) examined the effects of cognitive load and intertrial interval (ITIs) on judgements of control across zero and positive contingencies. ITIs occur between the gap between trials, and are thought to potentially affect they way in which participants view contingencies. In a 2 (between subjects: no vs high cognitive load) by 2 (within subjects: high vs low density) by 2 (within subjects: action vs context) by 2 (within subjects: positive vs zero contingency) design 34 participants gave judgements of control ratings for themselves and the context. Participants were randomly assigned to either no load or high load condition, in which participants could control a remote control which may turn on the music, however the music could also turn on by itself. Main effects of both outcome density and contingency were found, participants gave higher ratings in the positive contingency compared to the zero contingency, and higher ratings in the high outcome density condition than the low outcome density condition. This replicates previous findings.

Maldonado et al. (2006) investigated the effect of inattentional blindness in positive and negative contingencies. In a 2 (between subjects: positive vs negative contingency) by 2 (between: high vs medium density) by 2 (within subjects: attended cue vs incidental cue) design, 64 participants viewed a series of ill patients who presented with two different symptoms. Participants focused on one symptom, which was the attended cue as participants were actively attending to it, and after every 8 trials gave a judgement on the strength of the relationship between the attended symptom and the disease, for a total of 32 trials. At the end of the task participants were asked to make a judgement on the relationship between the incidental cue which sometimes co-occurred with the attended to cue or by itself. Participants detected both the negative and positive contingencies for the attended cue, however for the incidental cue they only detected the positive contingency.

This experiment shows that people can detect both positive and negative contingency when they are paying attention to the cue. However, this experiment used paired cues, and did not assess the outcome density effect itself.

Maldonado et al. (1999) also investigated the effect of training with a zero-contingency condition first, and then on subsequent judgements in positive and negative contingency conditions. In a 2 (between: positive vs negative contingency) by 2 (between: trained vs untrained) design, 64 participants viewed 68 fictitious patients who did or did not have tropical herpiasis disease, and who did or did not have frontopigmentosis. After every 8 trials participants provided a judgement for the strength of the relationship between the symptom and the disease based on the patients they have seen thus far. Both the trained and untrained groups detected the positive and negative contingencies, however those trained in the zero contingency group took longer to detect the contingency condition they were in. This shows that people can detect both positive and negative contingencies, however prior experience can delay the time taken to learn a new contingency.

Msetfi et al (2007) conducted an experiment which investigated the influence of Inter Trial Intervals on judgements across zero, positive and negative contingencies, between depressed and non-depressed participants. In a 2 (contingency: zero, positive) by 2 (length of intertrial interval: short, long) by 2 (mood: nondepressed, depressed) by 2 (sex: female, male) fully between-subjects design 96 participants provided judgements of control. They found significant main effect of contingency, participants gave higher judgements in the positive contingency compared to the zero contingency, similarly an interaction effect of ITI by mood indicated that when ITI was long non-depressed participants gave higher judgements in both zero and positive contingencies compared to depressed participants. In a follow up experiment with 48 participants, the same effect was examined in a negative contingency condition (Msetfi et al., 2007). An interaction effect of mood by ITI was found, suggesting that non-depressed participants gave higher judgements of control in the long ITI condition compared to the short ITI. Depressed participants gave more negative judgements in the long ITI condition. Whilst this experiment has explored the influences of judgements across zero, positive and negative contingencies, these were not conducted within the same experiment. In addition to this outcome density was not considered here in this experiment.

Later Msetfi et al. (2013) explored the effect of depression on contingency judgements across positive, negative and zero contingencies with high and low outcome density conditions. Using a 3 (within subjects: negative, zero, positive contingency) × 2 (within subjects: low, high outcome density) × 2 (within subjects: action, context cue) × 2 (between subjects: low vs high BDI) design, 50 participants provided contingency judgements. Participants imagined they were in a house with a hidden stereo system, they could control the music turning on in different rooms using a remote. However, the remote only works intermittently, and sometimes the music turns on by itself. Participants had to test the remote in each room. For each experimental trial participants had to press a space button (remote) after which an auditory signal may be presented. There were 40 trials per block, with 6 blocks (3 contingencies by 2 outcome density conditions). At the end of each block participants gave a rating of their own control (action) and the control of the context over the auditory signal, on a -100 to 100 scale. Main effects of contingency were found, suggesting that participants gave highest judgements in the positive contingency and lowest judgements in the negative contingencies. Main effects of outcome density were also found suggesting that participants gave higher judgements in the high outcome density condition which were more toward the positive end of the judgement scale, compared to the low outcome density condition. For negative contingencies, the high outcome density judgements were closer to zero compared to the low outcome density judgements.

However, this experiment uses a partially between factor design for low and high Beck's Depression Inventory (BDI), as such there are only 24 (low) and 26 (high) participants for the BDI groups, giving relatively small samples for the other conditions. The study used a symmetrical calculation (Table 9) across all the calculations, which allows for more direct comparison across contingencies. Although the study uses conditions for the positive condition under which the P(O/C)=1, which means there are no trials where the cue occurs without the outcome. Similarly for the negative contingency the P(O/~C)=1, which means there is a complete absence of no cue-no outcome trials. These probabilities may possibly lead to exaggerated positive or negative biases, depending on which cell type has the most influenced. Wasserman et al. (1990) found that participants gave weighting to trial types

Table 9 Msetfi et al. (2013) contingency calculations

Low outcome density at Positive Contingency

P(O/C)=.5
P(O/~C)=0
ΔP=.5-0=.5

High outcome density at Positive Contingency

P(O/C)=1
P(O/~C)=.5
ΔP=1-.5=.5

Low outcome density at Zero Contingency

P(O/C)=.25
P(O/~C)=.25
ΔP=.25-.25=0

High outcome density at Zero Contingency

P(O/C)=.75
P(O/~C)=.75
ΔP=.75-.75=0

Low outcome density at Negative Contingency

P(O/C)=0
P(O/~C)=0.5
ΔP=0-.5=-.5

High outcome density at Negative Contingency

P(O/C)=.5
P(O/~C)=1
ΔP=.5-1=-.5

according to A > B > C >D when making judgement. As such this may have an affect the results of the study, due to participants giving unequal weighting to trail types.

1.7.1 Limitations

Within the negative contingency literature, it is difficult to make comparison to much of the literature, due to the focus often being on inhibition and control (Karazinov, 2018; Lee & Lovibond, 2021; Avellaneda et al., 2016; Williams, 1995). Within this literature, there are often multiple types of cues presented simultaneously, whereby some cues may be paired with another cue, which results in no outcome. Hence these types of trials and cues are paired to assess inhibition. Whilst this literature does provide information regarding negative contingencies to some extent, the pairing of two or more cues together means many of these studies cannot be compared to the studies by Msetfi et al. (2013) or those reported in this thesis, which present singular cues. A similar issue occurs within the positive contingency literature as well where by cues are paired and unpaired throughout trials (López et al., 1998)

Another issue which arises for comparison across the literature, is that many of the contingency studies, do not include both high and low outcome density conditions. Most often high outcome density conditions are only included, as it is well known that these produce biased judgements (Matute, 1996). As such, comparisons of the outcome density effect across positive and negative contingencies are difficult. Finally, much of the literature in this area focuses on active paradigms whereby participants can respond or provide the cue which results in an outcome or not (Blanco et al., 2011; Shanks & Dickinson, 1991). Whilst in some studies the outcome is programmed to occur on a certain number of times, due to participants being able to provide the cue however many times they may want to, this results in the experienced contingency of participants varying greatly.

1.7.2 Summary

Overall, the contingency effects have been studied well. However, the research has primarily focused on zero and positive contingencies and/or active contingency paradigms.

Chapter 6 of this thesis reports a study which tested contingency learning with zero, positive and negative contingencies in a passive contingency task.

1.8 Overview of research questions

The previous sections outlined and reviewed the literature on contingency learning. This thesis will further examine the outcome density effect, in relation to cue density and contingency to explore how this results in biased judgements. Across all the experiments I will be examining whether cue and outcome density effects are robust effects? The hypothesis that belief in an association between two events will increase, with higher frequency of alleged cause and target outcome, even if there is no actual relationship between the two is tested. In addition to this, I will be examining some of the boundary conditions under which the outcome density effect occurs, as well as comparing the size of outcome density effects and cue density effects for equivalent manipulations of density. Do methodological variations of the contingency paradigm effect the size of the outcome density effect? Finally, I will be examining the influence of affect and prior beliefs on judgement. Do factors beyond methodological variations influence the outcome density effect? These three research questions work together to understand the conditions under which false beliefs may be formed, and potentially provide conditions under which the false beliefs may be reduced or overcome.

1.9 Overview of chapters

Chapter 2 outlines a scoping review of the methodological variations seen across contingency paradigms, along with the methodological framework of this thesis. There are 7 empirical well-powered studies included in this thesis, outlined across chapters 3 to 9. The empirical chapters examine the influence of different factors on the outcome density effect, these are outlined in further detail below. The concluding chapter (chapter 10) discusses the findings, implications, and limitations of the project.

1.9.1 Chapter 3 to 5: Outcome and cue density effects

Four experiments are reported to replicate the outcome density effect and test the effect in lab and online settings. In addition to this, the studies address the gap in research between outcome and cue density effects by testing these effects against each other, and the effect presentation order of cue and outcome on these effects.

Experiment 1 investigated the hypothesis that false beliefs or judgement of perceived association increased with higher frequency of alleged cause and target outcome, even if there is no actual statistical relationship between the two. In both online and lab settings, I examined the influence of outcome density on judgements of effectiveness of medication, in a contingency paradigm, to test if outcome density effects differ between online and lab settings. The experiment (*N*=58) confirmed that the outcome density effect, is a very large effect, participants give higher judgement of causality when the probability of outcome is higher compared to when the probability of the outcome is lower. Whilst the effect was found in both locations, overall, the effect was larger in the lab compared to online, however there was no main effect of location.

In Experiment 2, I compare cue density and outcome density effects directly, by manipulating these densities in a contingency paradigm. By investigating cue and outcome density effects at the same time, I investigate if the same cognitive principles underlie both, or if there is something qualitatively different between them. Experiment 2 (N=85) found an effect of both cue and outcome density; however, the cue effect seems to be smaller than the outcome density effect.

In Experiments 3 and 4, I manipulate the order of cue and outcome presentation to compare the effects of the first seen event vs. the second seen event. This novel extension reflects that while, logically, a cause must precede an outcome, sometimes people observe or learn of an outcome *before* they see or hear of its possible causes. These event presentation order manipulations permit exploration of interaction effects involving *both* density level and event presentation order. I can therefore consider whether the effect of cue or outcome density on contingency judgment is altered by reversing the order in which event information is seen. Experiment 3 (N=96) found that the outcome density effect is present in an equal size, even when the order of event presentation is manipulated. Experiment 4 (N=97) found that the cue density effect is present in a similar size, even when the order of event presentation is manipulated. Although the effect was not as large as the outcome density effect. These results together suggest that the outcome density effect is a robust effect. Whilst the cue density effect is also found across these conditions it is a smaller effect. In addition to this order of event presentation makes little or no difference to the outcome and cue density effects.

1.9.2 Chapter 6: Scales and Contingencies

Studies have not only found the outcome density effect with zero contingencies but also for both positive and negative contingencies. I further examined the outcome density effect across positive, negative and zero contingency conditions within the same study to compare the effect of outcome density. In addition to this the effect of scale length is explored be comparing a frequently used 0-100 judgment scale with a -100 to +100 scale. Experiment 5 (*N*=131) found an effect of outcome density, contingency and scale. In addition to this, people naturally seem to be making associations that are more positive than the contingency would indicate.

1.9.3 Chapter 7 and 8: Beliefs and cover stories

Prior literature has examined the influence of prior beliefs on judgements, and shows that prior beliefs can affect judgements, such that strong prior belief is linked to better learning when the to be learned association matches the prior belief. However if the learning is in the opposite direction of the prior belief then learning takes longer to occur and judgements can be more biased (Biele et al., 2009; Yarritu et al., 2015). In addition to this, research also suggests that the effects of prior belief can be modified based on factors such evidence of the predictive value of cues (Catena et al., 2008). However these have often focused on prior beliefs which have been induced within the experimental setting (Biele et al., 2009; Catena et al., 2008; Czupryna et al., 2021; Evans et al., 2005; J. C. Perales et al., 2007; Yarritu et al., 2015). The influence of pre-existing prior beliefs on judgements in passive contingency tasks has not been explored thus far, as far as I am aware. As such, this study explores this factor.

In addition to this, I examine whether the outcome density effect occurs in different contexts (e.g., beyond a standard medical treatment paradigm). First, a pilot was conducted to explore the most appropriate contexts to be used. As the prior beliefs one holds may affect judgments in the new contexts, these contexts were chosen to be based around real life scenarios, and prior belief was also assessed for each context. Experiment 6 (*N*=204) found the outcome density effect to remain robust. There was some variability in judgement and prior beliefs across contexts. Overall, there was evidence to suggest that context-specific prior beliefs influence contingency judgements.

1.9.4 Chapter 9: Affect

Individual differences in mood have also been studied in relation to contingency judgements. For example, it has been shown that positive mood may lead to increased belief that causes and effects are linked (Clore & Huntsinger, 2007). Alloy and Abramson (1979) showed individuals with depression had more accurate judgements of the relationship between their actions and outcomes compared to non-depressed individuals. To date research has focused on whether low mood influences contingency judgements, however the influence of other negative emotions such as anxiety, fear, or anger, in comparison to positive emotions such as joy, gratitude or amusement on contingency judgements, have been studied less. In addition to this much of the previous work has largely examined the influence of affect in active contingency paradigms, however I extend this work by using a passive contingency paradigm. This experiment measured mood to identify whether naturally occurring differences in emotional state relate to differences in judgment. Experiment 7 (*N*=169) found affect to influence contingency judgements. Specifically, there in the low outcome density conditions there was an influence of positive affect, negative affect, and anxiety on contingency judgements. In the high outcome density conditions, there was an influence of anger and stress on contingency judgements.

Chapter 2: Methodological Framework

This chapter outlines the methodological framework of the experiments in this thesis. Firstly, an overview of methodology used to study the outcome density effect is presented, which is then used to provide the methodological justifications for the experiments within this thesis. Following this, the general methodology is presented. Finally, the ethical approval details are given.

2.1 Methodological Review of the Outcome Density Effect

2.1.1 The typical outcome density paradigm

Studies exploring the outcome density effect typically use the same paradigm: participants are exposed to a sequence of trials in which two stimuli—usually a cue and an outcome—may or may not co-occur. This results in four possible combinations of cue and outcome, as presented in Table 10. The contingency between these stimuli is manipulated by changing the frequency of stimuli co-occurrence. In some studies, participants are given training with similar stimuli to those in the main trials, to ensure participants understand the task. Next, participants take part in the main task. Participants also provide their estimation of the degree to which the cue and outcome are related to each other.

2.1.2 Methodological variation in outcome density studies

Table 10 Visual presentation of the possible cue-outcome pairings in a contingency paradigm

	Outcome	
Cue	Present	Absent
Present	A	B
Absent	C	D

Alloy and Abramson (1979) conducted one of the first studies on outcome density in contingency learning, using a modified paradigm previously used by Jenkins and Ward, 1965. The studies presented depressed and non-depressed participants with a series of contingency problems. Participants were tasked with turning on a green light as often as possible by either pressing a button or not pressing a button, which would or would not result in a green light. Pressing the button would cause the light to turn on with a probability of X; not pressing the button would do so with a probability of Y, with X and Y dependent on the condition. The contingencies and outcome densities were varied between-subjects. Participants took part in a total of 40 trials, at the end of which they were asked to provide their judgements of control on a 5-point scale (1=complete control, 5=no control). Since these studies, the outcome density effect has been further explored using the same methodology as well as variations, briefly discussed below.

Methodological variations across studies

In this section the methodological variations seen within experiments which use the contingency learning task to examine the outcome density effect will be discussed. Twenty-seven studies were found in this[2], and some will be discussed below to demonstrate the different variations of the contingency paradigm.

Scale type

The way in which participants' estimation of the degree to which the cue and outcome are related vary across studies. Some of these variations are linked to the scale wording used. Some use a numerical scale of 0 (no control) to 100 (complete control) (Gillan

---

[2] The search terms "outcome densit*" AND "contingenc*" were used for this literature search in July 2020.

et al., 2014a) whereas others use a visual analogue scale ranging from 0% (no improvement) to 100% (full recovery) (Chow et al., 2019b). The scale type used is typically dependent on the context of the experiment. Chapter 6 in this thesis reports a formal comparison of two types of scale and demonstrates and discusses how the choice of scale can affect the results of a contingency learning experiment.

Measurement type

Similar to scale type, participants' estimation of contingency is often not the only measurement taken. First, the wording of the measures differ, such as judging your own ability to illuminate a light bulb (Gillan et al., 2014a) or that of another person (e.g., when children assessed the ability of two characters to grow plants; Moreno-Fernández et al., 2017). In other cases, participants assess the effectiveness of an object. For example, the effectiveness of powder versus liquid in the ability to grow plants (Moreno-Fernández et al., 2017). In many studies, where participants are actively controlling the occurrence of the cue, the level of control is measured (Chow et al., 2019b). Finally, some studies require participants to judge the perceived relationship at the end of each trial (Blanco et al., 2013), at the end of each block (Msetfi et al., 2015), or only at the end of the entire experiment (Vallée-Tourangeau et al., 2005). These variations depend on the context of the experiments, for example experiments using active tasks often assess control, whilst passive tasks often use judgements of contingency.

Cover stories

A methodological variation is the context or cover story used in contingency learning tasks differ between studies. Some cover stories are fairly simple, such as participants' aiming to illuminate a light bulb (Alloy & Abramson, 1979; Gillan et al., 2014a; Jenkins & Ward, 1965; Msetfi et al., 2005), controlling a blue flash of light on a computer screen (Matute et al., 2007b), or assessing the extent to which pressing the spacebar on a computer keyboard caused the appearance of a geometric shape (Crump et al., 2007; Vallée-Tourangeau et al., 2005). Other cover stories have a more real-world reflection. For example, in the commonly used allergy task, participants are given information about patients with the fictitious Lindsay syndrome who may be treated with a fictitious medication which may lead to patient recovery (Barberia et al., 2019; Blanco et al., 2013; Chow et al., 2019b; Yarritu et al., 2015). Some cover stories are more imaginative, such as participants pressing a button and shooting lasers at Martian spaceships to prevent them from landing and invading earth, in a video game format (Matute et al., 2014). In sum, different cover stories have been used to study the outcome density effect in contingency learning tasks. While different cover stories are used in different studies, there is little (if any) research that formally examines variation in contingency judgment as a function of cover story. Experiment 6 in this study reports such an investigation, using four different real-world scenarios to examine the outcome density effect.

Passive vs active

Often dependent on the context of the study, participants' active or passive involvement varies between studies. In active tasks, participants may press a button whilst in passive tasks, participants view a series of cue and outcome combinations. In Alloy and Abramson (1979), participants were actively taking part in the study by choosing to press the button within three seconds of the beginning of the trial, regardless of whether they have control over a green light appearing. Hence when participants are actively involved in a task, participants assess the relationship between their own action (cue) and the outcome. The active paradigm is often used to assess the illusion of control, whereby participants tend to indicate they have greater control over an outcome than the programmed contingency indicates.

In the standard passive allergy paradigm, participants are told to imagine they are medical doctors. They must find out if a fictitious medication is effective in treating a fictitious illness and therefore leads to patient recovery. Participants view a series of information on patient illness and recovery in the presence of absence of medication (Blanco et al., 2013). In this paradigm, participants are learning the possible contingency between a cue and outcome by passively observing. Passive contingency tasks are often used to avoid differences in the relative frequency of cues derived from active intervention by participants, and to hinder any contributions from self-performance that may complicate the evaluation of the outcome density as a source of bias itself. Other techniques have been used here and these are sometimes problematic and may add noise. Due to this, the choice between active or passive participant involvement in outcome density experiments depends on the context and stimuli used. The reasons for researchers using an active or passive task are not always explicit. Presumably, some researchers choose the context which ends up being an active/passive task; while other researchers may choose an active/passive task and identify a suitable context, accordingly.

There are other variations in these experiments, such as participants viewing the context and stimuli from a first-person view, most often seen in active participant involvement studies, compared to third person view, often seen in passive involvement studies.

Between vs within-subjects

Between-subjects and within-subjects designs have both been used in contingency learning tasks. However, between-subjects designs typically have more participants than within-subjects designs. For example, using a within-subjects design, Crump et al. (2007), Moreno-Fernández (2017), Vallée-Tourangeau (2005), used samples of *N*=28, *N*=37 and *N*=34, respectively. In contrast, Barberia et al. (2019), Matute et al. (2011b), Vadillo et al. (2011a) using a between-subjects design used samples of *N*=150, *N*=108 and *N*=144, respectively. Due to this, within-subjects designs are an alternative used in contingency learning tasks, as a smaller number of participants are usually required. Moreno-Fernández et al. (2017), used an experiment with a within-subjects design, featuring two outcome density conditions (high vs low) to test if high probability of outcome encouraged causal illusions in children, which had been shown in adults previously (Matute et al., 2015). Participants played the role of a judge in a planting contest and assessed the ability of two different characters to make plants grow. Participants saw the two characters perform sequentially (one after the other), after which they were required to award a trophy to the best farmer. There were four blocks of trials, each consisting of 12 trials, two blocks for each condition, which were counterbalanced across participants. Whilst within-subjects designs are associated with order effects, this can be addressed by counterbalancing the order of conditions for participants. Some studies may also use a mixture of between and within-subjects design features depending on the measures being studied (Cavus & Msetfi, 2016). When designing an outcome density study, choosing between-subjects versus within-subjects design may be dependent on recruitment constraints, interest in specific individual difference, and design complexity.

Lab vs online

Most contingency learning studies of the outcome density effect have been conducted in lab settings similar to Alloy and Abramson (1979), perhaps due to the ease of conducting experiments in lab settings. However, in recent years, experiments have been carried out in online settings. One study by Matute et al. (2007) compared outcome density effects between lab and online settings and concluded that online research was suitable for this line of research. The first experiment in this thesis extends this comparison of lab and online settings.

Contingency

The contingency literature has also examined positive and negative contingencies along with zero contingencies (Msetfi et al., 2007, 2013b; Vallée-Tourangeau et al., 2005), under various conditions. Chapter 6 in this thesis extends this work by directly comparing contingency judgements and the outcome density effect across contingencies.

Manipulations within studies

A factor that varies across and within studies is the outcome density itself. Several studies use only high outcome densities (Barberia et al., 2019), whilst other studies use both high and low outcome densities within the same experiment for comparison across conditions (Moreno-Fernández et al., 2017). Another way outcome density is varied, is by using continuous high and low outcome densities rather than binary high and low outcome densities. Chow et al. (2019) used a between-subjects design, 2 (high vs low outcome density) × 2 (binary vs continuous outcome density), in which participants imagined they were medical researchers and judged the effectiveness of a new experimental drug in treating an illness. In the binary outcome density condition, outcome density was the exact value of either 80% (high outcome density) or 20% (low outcome density). In the continuous outcome density condition participants sampled from a bimodal distribution of outcomes from either the low distribution (*M*=20, *SD*=5) and high distribution *(M*=80, *SD*=5). The study found no interaction effects between outcome density and outcome variability, suggesting that fixed and variable outcomes did not produce significantly different effects. Depending on the experiment, more than one outcome density may be used, however most often high density is only used, as the outcome density effect is seen with high outcome density.

2.1.3 Summary

Contingency theory provides an account of causal belief, setting out the idea that when the contingency is zero, there should be no perceived contingency between a cue and outcome. However contingency beliefs can be biased by the outcome density effect, particularly when outcome density is high. Contingency learning tasks exploring the outcome density effect vary in methodology in terms of cover stories, designs, and measurement of dependent variables, online and lab settings, and binary and continuous outcome density used. Hence, when designing an outcome density study is important to consider these factors.

2.2 Methodological justification

In the previous section I outlined key variations in methodology across outcome density studies. Whilst these are all valid variations that are applicable to the contexts of each study, for the experiments in this thesis, all of these variations cannot be applied. As such, below I outlined the methodological choices and the justifications for these choices that will be used throughout this project.

2.2.1 Zero contingency

Throughout this thesis, all experiments will feature zero contingencies, (experiment 5 also features positive and negative contingencies). The main interest of this thesis is to examine the formation of false beliefs in situations where such beliefs should not occur as there is no association between the cue and outcome. As such zero contingencies are the most suitable contingency to use, even though the frequency of cues and outcomes can be varied, within a zero contingency paradigm the presence of the cue does not predict the occurrence of the outcome.

2.2.2 Passive vs active

In this area of research active tasks are more common and, as such, outcome density effects are often described as indicating an illusion of control. For the purposes of this project, however, I am interested in the passive tasks, because false beliefs formed through fake news are often caused by hearing or seeing information without personal intervention. People often do not actively take part in the formation of fake news. Also the key research objective for this thesis is to explore the underlying mechanism of false belief formation, and not judgments of control. As such the passive version of the contingency tasks will be used, where participants will observe a series of events and provide their judgements of contingencies. Accordingly I will use an adapted version of the standard allergy paradigm (Blanco et al., 2013) which used a medical context will be used. This paradigm is often used within this field and the outcome density effect has been found with this paradigm. In addition to this the passive judgement paradigm is easy for participants to grasp.

2.2.3 Scale type

As the passive contingency paradigm in the medical context is being used for this project, the scale used will assess the effectiveness of the given medication for the initial experiments. In addition to this, the 0-100 scale will also be used, with the appropriate labels, such a scale is common in contingency learning research, and therefore supports comparison between my findings and those in the literature. For the later experiments, these design features are manipulated to develop a better understanding of the properties of the scale most commonly used in this line of research.

2.2.4 Between vs within subjects

For all but the sixth experiment in this thesis, outcome density is manipulated within subjects such that all participants complete one low and one high outcome density condition and provide a contingency judgement for each density. This is to allow for the precise comparison between densities, eliminating individual variability in judgment, to determine the presence and size of the outcome density effect, which is the basis of this project. Other than outcome density, the other factors are manipulated between and within subjects however this is dependent on experiment. Details of these are provided in each empirical chapter.

Due to the within subjects' manipulation, consideration has been given towards the stimuli used across the experiment. All experiments, except experiment 6 used the adapted version of the allergy paradigm (Blanco et al., 2013). However as the paradigm is repeated multiple times, the use of the same stimuli can lead to learning effects, which may bias the judgements in the subsequent conditions. Maldonado et al. (1999) found a training phase with a zero contingency, followed by trials of either positive or negative contingency, led to participants taking longer to accurately detect the programmed contingency, in comparison to those who did not experience a zero contingency training phase. Whilst only one experiment in this thesis uses positive and zero contingencies, to mitigate effects such as these which may be applicable to zero contingency conditions as well, three fictitious disease and illness names were used, such that participants never encountered the same stimuli within an experiment.

2.2.5 Number of Trials

To avoid any fatigue effects due to too many trials, the number of trials will be kept low. However, it is also important to ensure that participants sample enough trial types out of the four possible trials. As such each condition will contain 40 trials. This allows for a minimum of 4 trials per trial type (A, B, C or D, see Table 10), while maintaining a zero contingency for outcome densities as high (low) as 0.80 (0.20). As participants may not fully be paying attention for every trial, reducing the trial number to 10 or 20, so that only 1 or 2 trials (see Table 11 for calculations for different trial numbers) of certain trial types may affect the overall validity of the experiment. Within each block of 40 trials, all the trials will be fully randomised, to mitigate trial order effects.

2.2.6 Other considerations

When designing an outcome density paradigm, it is vital to consider the combinations of cue and outcome that participants may sample, that may not have been purposefully built into the experiment. Taking the example of intertrial intervals (ITIs), the duration of time between the onset of one trial and the onset of the next trial, may also be considered as the absence of both cue and outcome, as seen in Table 10. Due to this, increased or decreased intertrial intervals could change the contingencies and outcome densities present in a study. For example, if there are long intertrial intervals, participants may experience this as the absence of both cue and outcome beyond what the experiments manipulated, as seen in Table 12. This could then result in the contingency no longer being 0

Table 11 Calculations for different numbers of trials

10 trials

	Recovery		
Medication	Present	Absent	
Present	1	4	P(O/C)=.2
Absent	1	4	P(O/~C)=.2

Δp=.2-.2=0

20 trials

	Recovery		
Medication	Present	Absent	
Present	2	8	P(O/C)=.2
Absent	2	8	P(O/~C)=.2

Δp=.2-.2=0

30 trials

	Recovery		
Medication	Present	Absent	
Present	3	12	P(O/C)=.2
Absent	3	12	P(O/~C)=.2

Δp=.2-.2=0

40 trials

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C)=.2
Absent	4	16	P(O/~C)=.2

Δp=.2-.2=0

even though this may have been the programmed contingency. Rather, participants may experience a positive contingency, and which would then bias their estimates of the contingency and bias the results. Due to this, when designing outcome density experiments, potentially unintended combinations of cue and outcome should be taken into consideration, so the designed contingency and outcome density is maintained. For the experiments within this thesis, the presence of cue and outcome, and the absence of cue and outcome are defined as individual images, Figure 1, as such even though there may be intertrial intervals, these should not be defined as no cue – no outcome pairings.

2.3 General methods

I employed a standard contingency learning paradigm throughout these studies with small, but important, incremental variations from one experiment to the next. I discuss the common design and experiment idiosyncrasies below. Details of specific differences are outlined within the relevant methods section for each study.

Table 12 Zero contingency calculations without and with 5 intertrial intervals

Zero Contingency with no intertrial interval

	Outcome		
Action	Present	Absent	
Present	15	5	P(O/C)=.75
Absent	15	5	P(O/~C)=.75

ΔP=.75-.75=0

Zero Contingency with intertrial interval

	Outcome		
Action	Present	Absent	
Present	15	5	P(O/C)=.75
Absent	15	5+5	P(O/~C)=.60

ΔP=.75-.60=0.15

2.3.1 Design

Figure 1 Stimuli images. Adapted from Blanco et al. (2014) with their permission

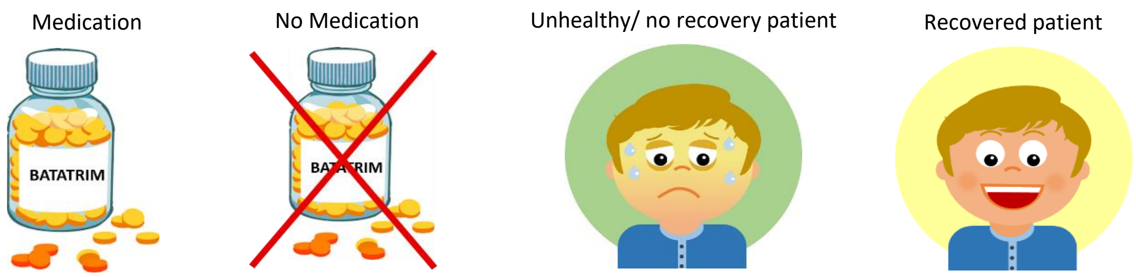

Across all experiments density level was manipulated within subjects. The dependent variable was always participants' judgements of the effectiveness of a medication, except experiment 6 which used different cover stories.

Density level

Outcome density was manipulated to be either low (0.2) or high (0.8) (Table 13) while maintaining a zero contingency. The density levels refer to the probability of an outcome occurring in the presence of a cue or the absence of a cue. The higher the density, the higher the probability the outcome occurs with and without a cue. Cue density was held at a medium density of 0.5, for both high and low outcome density. Except experiment 5 which also included non-zero contingencies, the same contingency calculation was used for all studies. The experiments were run across two sessions. For each session there was a block of high and a block of low outcome density trials; each block consisted of 40 trials. The block shown first was chosen randomly. To reduce any carry-over or learning effects across

Table 13 Low and high outcome density calculations in zero contingency

Low outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C)=.2
Absent	4	16	P(O/~C)=.2
	Δp=.2-.2=0		

High outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	16	4	P(O/C)=.8
Absent	16	4	P(O/~C)=.8
	Δp=.8-.8=0		

sessions, different medication and illness names were used for each session, randomly assigned to block and experiment session for each participant.

Judgements

Participants indicated at the end of each block to what extent they thought the medication presented to them had been effective in treating the illness the patients were suffering from (except experiment 6). They did so on a slider scale ranging from 0 (it was not effective at all) to 100 (it was perfectly effective). For experiments 5 and 6, the slider scale ranged from -100 (It caused all symptoms in all cases) to 100 (It prevented all symptoms in all cases). There was no slider handle displayed until a response was given to avoid judgment anchoring (Mussweiler et al., 2012). Participants could not move to the next screen until a response was given.

2.3.2 Task and stimuli

The experiment is a conceptual replication of Blanco et al. (2014), part of the material used the stimuli from Blanco et al. (2014) with their permission. The task was built on Gorilla (Gorilla Experiment Builder, 2020), as an online task so participants could complete the task from their location of choice for the online portion of the experiment. The illness, Lindsay syndrome, and medication name, Batatrim, for the task were names used in previous studies (Blanco et al., 2014; Yarritu et al., 2015). Three additional sets of illness and medication names were created to permit counterbalancing and to ensure participants were provided new stimuli covers for each session, thereby counterbalancing of illness/medication across conditions to reduce carryover effects across conditions. The following new names were also used: Amplexia syndrome and Fexolise; Marasia syndrome and Lysopran; and Sawyer Syndrome and Quitrile. All the illnesses and medications were fictional. See Figure 1 stimuli images. Participant saw two of the four orders in Table 14, one for each session. The same stimuli were used for all studies except experiment 6.

2.3.3 Procedure

The study was designed on Gorilla (Gorilla Experiment Builder, 2020) and took place online (except experiment 1). Participants completed a contingency estimation task four times (except experiment 7). The sessions were a minimum of three days apart to minimize carry-over or learning effects. Participants were recruited through Prolific, an online research participant platform with approximately 100,000 registered research volunteers (Prolific, 2020) and were prompted to complete the first session of the experiment. After 3 days, participants were invited to complete the second session of the experiment, which followed the same procedure as the first session. Participants could complete the task from their location of choice.

Table 14 Order of densities, illness and medication shown to participants for each session

			High Outcome Density Block	Low Outcome Density Block
Session 1				
	Outcome density	Order 1	Lindsay Syndrome and Batatrim	Marasia Syndrome and Lysopran
		Order 2	Marasia Syndrome and Lysopran	Lindsay Syndrome and Batatrim
Session 2				
	Outcome density	Order 1	Amplexia syndrome and Fexolise	Sawyer Syndrome and Quitrile
		Order 2	Sawyer Syndrome and Quitrile	Amplexia syndrome and Fexolise

Upon choosing to take part in the study, participants read the information sheet, provided consent to participate and answered demographic questions (gender, age, ethnicity, an ability to communicate in English), and then read the instructions. Participants imagined that they were a doctor specialising in new illnesses and their potential treatments. Once participants started the task, they were told the following patients are suffering from syndrome and may be treated with a specific medication. They then proceeded to the first trial.

Within each trial, the following images were shown (see Figure 1), in a standard presentation order: fixation cross for 750ms, unhealthy patient image for 2000ms, medication image or no medication image for 3000ms, recovered patient image or not recovered patient for 3000ms. These images automatically followed each other, without input from the participant. Between trials, participants clicked 'next' to continue. Each block contained 40 trials; each session had 2 blocks (low vs. high density). Trial orders were randomised for each participant. Each block takes approximately 8 minutes to complete.

Participants indicated at the end of each block to what extent they though the medication presented to them had been effective in treating the illness the patients were suffering from. They did so on a slider scale ranging from 0 (it was not effective at all) to 100 (it was perfectly effective). Each session featured two points of measurement, one after the high outcome density block, and one after the low outcome density block. Over the 2 sessions, in total there were 4 points of measurement. At the end of the task a short debrief was given thanking participants for completing the experiment. Upon completion participants were paid.

2.3.4 Participants Recruitment

Participants were over the age of 18, fluent in English and registered on Prolific platform as a UK resident, see Table 15 for participant sample size and demographics across experiments. For a 2x2 within-subjects interaction of assuming (a) an attenuated interaction pattern, (b) any mean differences being 0.5 (all *SD*=1), (c) alpha=.05, (d) 40 people per group, (e) 0.5 common correlation with no sphericity, a power of 87% is achieved for 80 participants. As such all the studies except the first one* provides ample sample size to test the hypothesis (*originally enough participant were recruited, however due to Covid-19 recruitment had to be stopped resulting in only 58 participants fully completing the experiment). This power analysis is applicable to all experiments in this thesis.

2.4 Ethics

All experiments in this project involve human participants, and all the experiments were conducted in accordance with the ethical standard of the research ethics committee of King's College London. See Table 16, for the ethics identifier for each experiment. For all the experiment low risk ethics were gained, except experiment 7 which required high risk ethics

Table 15 Demographics of participants included in analyses

	N	Age (years), modal group, range	Female gender, *N (% of sample*)
Experiment 1	58	18-20 *(18-50)*	47 *(81%)*
Experiment 2	86	21-25, (*18-64)*	19 (*46.3%)*
Experiment 3	96	26-30 (*18-71+)*	62 (64.6*%)*
Experiment 4	97	31-35 (*18-70)*	60 (61.9*%)*
Experiment 5	131	21-25 (18-65)	50 (38.2%)
Pilot 6	100	26-30 (18-70)	25 (25%)
Experiment 6	204	31-35, 36-40 (*18-70+)*	93 (45.6*%)*
Experiment 7	169	18-20 (18 -71+)	118 (69.8%)

as participants affect information was also collected. Prolific (Prolific, 2020) and the Department of Psychology research participation pool at King's College London (SONA) were used to recruit participants. All data was collected using GDPR compliant participant recruitment services.

Table 16 Ethics identifier for each experiment

Ethics Identifier	Experiment
MRSP-19/20-17678	Experiment 1
MRSP-19/20-19281	Experiment 2, 3 ,4
MRSP-20/21-25363	Experiment 5
MRSP-21/22-29497	Pilot 6
MRSP-22/23-34307	Experiment 6
HR-19/20-19492	Experiment 7

Chapter 3: Outcome density in Online vs Lab

Most contingency learning tasks in the past on the outcome density effect have been conducted in lab settings similar to Alloy and Abramson (1979), perhaps due to the ease of conducting experiments in lab settings. However, it may perhaps be that this effect is due to the artificial lab settings used in previous experiments. In recent years lab experiments were replicated, in online settings, to both ease recruitment of a more representative sample and test if effects seen in controlled lab settings manifest in comparatively less controlled online settings.

Blanco et al. (2013) compared the effects of cue and outcome density effects in an online contingency paradigm. They used an adapted version of the allergy paradigm typically used in human contingency learning. Participants imagined they were doctors who had to judge whether a given medication was effective in treating patients in recovering from an illness. A significant effect of outcome density was found. Later Msetfi et al. (2015) conducted three experiments online assessing the strength of the outcome density effects based on collectivistic or individualistic values, whether participants were rating the control exerted by the context or their own control over the outcome via the action, and mood. Whilst an outcome density effect was found across the three studies, the effect was not present under all conditions.

Matute et al. (2007) have directly compared the outcome density effect in a lab and online setting experiment. In a 2 (lab vs online) × 2 (control vs no control) × 2 (obtain vs avoid flashes) between-subjects contingency learning task participants viewed a series of blue flashes on the computer screen for which they could choose to press a button (Matute et al., 2007b). The experiment was conducted online via a virtual lab and replicated in laboratory settings. Half the participants in both conditions were told they either had to control the flashes or they had no control. Of these groups, half were also told they had to avoid or obtain blue flashes. The study found a comparatively large effect size for the illusion of control effect in both the laboratory and the internet condition, even though all conditions used zero contingencies. Hence recent research is showing that outcome density experiments may be conducted in both lab and online settings.

Whilst Matute et al. (2007) has conducted an experiment which examined the outcome density effect directly in online and lab settings, the experiment focused on the control participants had over the outcome. For this thesis, the adapted version of the allergy paradigm using medication and illness stimuli used in Blanco (2013) is of more interest as it is more reflective of how people may generally take in information. However thus far there has been no study which has been conducted examining the outcome density effect using the adapted standard allergy paradigm, across online and lab settings. As such, the first experiment of this thesis examines if there is a difference in the size of the outcome density effect in lab and online settings, using a within subjects' design. This experiment aims to see if the outcome density effect is the same in lab and online settings, and if so, the subsequent experiments will be conducted online. This is both for ease of recruitment, and to reflect a more "real-life" setting in which people typically receive information.

The main goal of this chapter is to examine whether the outcome density effect is robust. This chapter also examined how methodological variations in the contingency learning paradigm affect the size of the outcome density effect. The experiment investigates the hypothesis that false beliefs or judgement of perceived association increases with higher frequency of alleged cause and alleged outcomes, even if there is no actual statistical relationship between the two. In both online and lab settings, I examined the influence of outcome density on judgements of effectiveness of medication, in a contingency paradigm, to test if outcome density effects differ between online and lab settings. The task and stimuli are as outlined in the general methods.

3.1 Methods

3.1.1 Design

A 2 × 2 × 2 mixed design was used, with judgement of effectiveness of medication in treating illness as the dependent variable. Three factors were varied in the experimental design: location (within subjects: online vs. lab) by outcome density (within subjects: high vs. low density) by session order (between subjects: online first vs. lab first; Table 17 and Figure 2).

Table 17 Order of outcome density for each session

		High Outcome Density Block	Low Outcome Density Block
Session 1			
	Order 1	Lindsay Syndrome and Batatrim	Marasia Syndrome and Lysopran
	Order 2	Marasia Syndrome and Lysopran	Lindsay Syndrome and Batatrim
Session 2			
	Order 1	Amplexia syndrome and Fexolise	Sawyer Syndrome and Quitrile
	Order 2	Sawyer Syndrome and Quitrile	Amplexia syndrome and Fexolise

Experiment Location

Participants completed a contingency estimation task four times, twice online in a location of their choice and twice in the lab under the supervision of the researcher. Order of location was counterbalanced: half the participants were set up to complete the task online first and the other half completed the task in the lab first. Participants completed both sessions.

Outcome density.

Outcome density was manipulated to be either low (0.2) or high (0.8) (Table 18), while maintaining a zero contingency.

3.1.2 Procedure and measures

Upon sign-up, participants selected two experiment dates, three to ten days apart. Next, participants were then randomly allocated to lab or online as their first session. The

Figure 2 Participant allocation to conditions across session

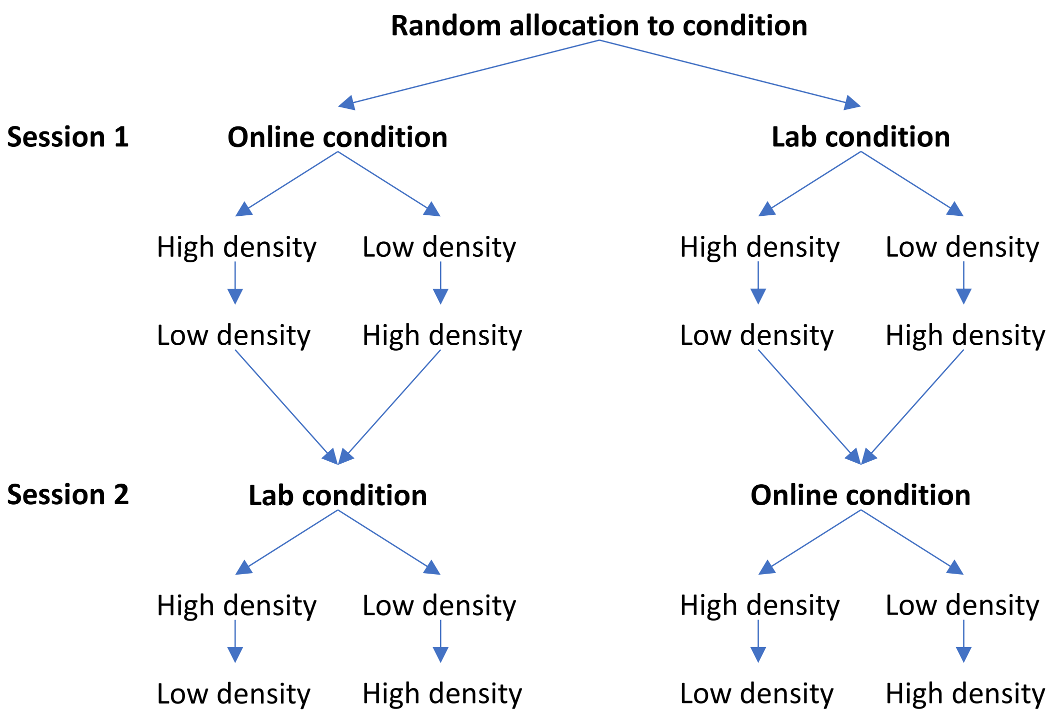

study information sheet and further instructions were emailed to participants in advance of taking part in either session. For the online session, login details were emailed at the start of the experiment. For the lab session, login details were completed by the researcher. The lab and online sessions follow the same procedure.

In the lab session participants logged into the experiment and completed it. A session took approximately 18 minutes to complete, and at the end of the second session a short debrief was given. If first session participants did not complete the online session within 48 hours, an email was sent every 24 hours reminding participants to complete the session, to ensure the session was completed at least 2 days before the lab session. Participants would then complete their second session either in the lab or online. Once participants completed both sessions, they were given either participation credits or £6.

3.1.3 Participants

A power analysis was performed to estimate required sample size, based on effect sizes from Chow et al. (2019), a recent study on outcome density. Specifically, I aimed for a sample size that would allow us to detect a possible interaction effect where a moderately

Table 18 Outcome Density manipulation calculations

Low outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C)=.2
Absent	4	16	P(O/~C)=.2
	ΔP=.2-.2=0		

High outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	16	4	P(O/C)=.8
Absent	16	4	P(O/~C)=.8
	ΔP=.8-.8=0		

sized (Cohen, 2013) outcome density effect existed in one location (δ=.33 e.g. in lab) but not the other (δ=0 e.g. online). To detect the corresponding 2 × 2 within-subject interaction effect with a power of 1 – β=0.80 and Type-I error of α=.05 (two sided), the sample size required was *N*=80. I aimed for 100 participants to allow for study cancellations and attrition between sessions.

Participants were 105 students recruited from the Department of Psychology research participation pool at King's College London. Participants confirmed they were over 18 and fluent in English. Eighty-four participants completed the first session of the study, and 58 participants completed both sessions before testing was stopped due to the imposition of restrictions on movement due to COVID-19 (March 2020). Table 19 contains participants' demographics.

3.2 Results

I conducted two paired-samples t-test to verify that the outcome density effect was present in both locations separately. Participants judged the effectiveness of the medication to be significantly lower in the low outcome density than the high outcome density condition both online, $t(57)=7.82$, $p < .001$, and in the research lab, $t(57)=6.00$, $p < .001$ (Table 20).

Table 19 Demographics

	First Session (N=84)	Second Session (N=58)
Age, modal group, *(range)*	18-20 *(18-50)*	18-20 *(18-50)*
Female gender, *N (% of sample*)	70 (*83.3%)*	47 *(81%)*
Online session, *N (% of sample*)	45 *(53.6%)*	36 *(62.1%)*

I ran a 2 × 2 × 2 mixed ANOVA with judgements of the effectiveness of medication as the dependent variable. Within-subjects independent variables were location (online or research lab) and outcome density (low or high). Session order (online first or lab first) served as between-subjects factor. I found the predicted significant main effect of outcome density on effectiveness judgements, $F(1, 56)=63.66$, $p < .001$, $\eta_p^2=.53$. Participants judged the medication to be significantly less effective in the low outcome density condition (*M*=22.56, *SD*=14.17) than the high outcome density condition (*M*=46.02, *SD*=21.02), with an effect size of *d*=1.31. There was no significant main effect of location $F(1,56)=0.60$, $p=.443$, $\eta_p^2=.011$, of first session location order $F(1,56)=0.001$, $p=.971$, $\eta_p^2=.00$. There was no significant location × first session location interaction effect, $F(1, 56)=0.25$, $p=.620$, $\eta_p^2=.004$.

The large outcome density effect was not significantly qualified by the location of testing, or by the location of the first session: there was no significant location × outcome density interaction effect, $F(1, 56)=1.43$, $p=.236$, $\eta_p^2=.025$; and there was no significant outcome density × first session location interaction effect, $F(1, 56)=0.41$, $p=.840$, $\eta_p^2=.001$.

There was, however, a significant three-way interaction for location × outcome density × first session location interaction effect, $F(1, 56)=5.45$, $p=.023$, $\eta_p^2=.089$. This triple interaction suggests that the size of the outcome density effect may vary according to the

Table 20 Mean (SD) judgement of medication effectiveness response by manipulation of density and location

	Location	
	Lab	Online
High density response	47.4 *(23.49)*	44.83 *(24.53)*
Low density response	19.57 *(16.95)*	25.45 *(20.88)*
95% CI for mean difference	20.70, 34.95	11.72, 27.14
Density effect size *d*	1.36	0.85

combination of testing location and the session order. I therefore compared the size of the outcome density effect across each combination of testing location and first session location (see Table 21 and Figure 3). This follow-up analysis demonstrated that the outcome density effect was large in each session (as defined by location and order) but seemed particularly large in the lab session when this was the first session that the participant had done, and seemed smaller for the online session if this was the participant's second session.

3.2.1 Session effect

To further examine the effect of sessions order, I ran two further ANOVAs for each session. I ran a 2 × 2 × 2 mixed ANOVA with location and outcome density as within-subjects factors and outcome density order during first session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant effect of location × density order during the first session, $F(1, 57)=0.97$, $p=.330$, $\eta_p^2=.017$), outcome density × density order during the first session, $F(1, 57)=3.00$, $p=.089$, $\eta_p^2=.051$), location × outcome density × density order during the first session, $F(1, 57)=0.17$, $p=.680$, $\eta_p^2=.003$).

Table 21 Mean (SD) response by order of first session location and location

	Lab session first		Online session first	
	Lab	Online	Lab	Online
High OD response	50.28 *(22.69)*	42.50 *(23.49)*	42.68 *(24.54)*	48.64 *(26.24)*
Low OD response	17.75 *(12.90)*	26.92 *(21.52)*	22.55 *(22.08)*	23.05 *(20.02)*
95% CI for mean difference	23.58, 41.48	6.45, 24.72	8.18, 32.10	10.98, 40.20
Outcome density effect size *d*	1.76	0.69	0.86	1.10

I ran a 2 × 2 × 2 mixed ANOVA with location and outcome density as within-subjects factors and density order during the second session (low density first or high density first) as independent variables, and effectiveness judgements as dependent variable. There was no significant effect of location × density order during the second session, F(1, 57)=0.58, p=.450, η_p^2=.010), outcome density ×

density order during the second session F(1, 57)=0.27, p=.870, η_p^2=.000), location × outcome density × density order during the second session, F(1, 57)=0.47, p=.500, η_p^2=.008).

3.2.2 Inclusion of all participants

The 2020 coronavirus pandemic prevented 26 participants from completing the second session of the study. Within-subjects ANOVA relies on full cases, thus excludes those

Figure 3 Mean effectiveness judgements by first session location, and location with 95% confidence

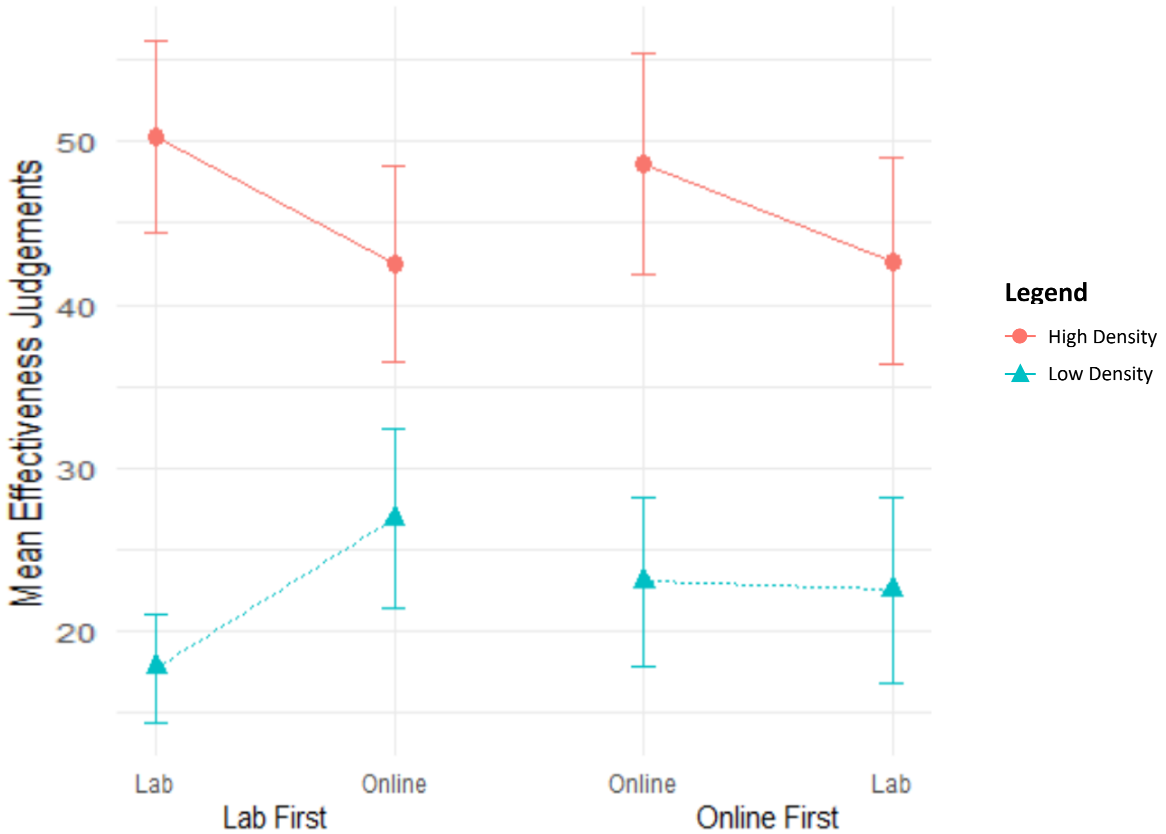

participants who completed only the first session. To verify that results upheld with and without the inclusion of these partial study completions I ran the same analysis using a random-intercept multilevel regression analysis (which retains partial responses) in SPSS. This type of analysis partitions residuals into a participant-level term (the random intercept) and observation-level term (residual error), thereby circumventing the (participant-level) independent error-terms assumption that regular regression analysis makes. I regressed participants' judgement on effect-coded location (-1 online, 1=research lab), outcome density (-1=low, 1=high), and the location × outcome interaction; a random intercept was assigned to participants. Similar to the ANOVA results, I found a significant main effect of outcome density $F(1, 280)=101.24$, $p < .001$, $B=11.60$, $SE=1.15$, 95%CI=[9.33, 13.87], $t(280)=10.06$, $p < .001$. Neither the main effect of location, $F(1, 280)=0.42$, $p=.518$, nor the location × outcome density interaction, $F(1, 280)=2.64$, $p=.105$, was significant.

3.3 Discussion

Experiment 1 examined the influence of outcome density on judgments of effectiveness of medication in treating illness, in a contingency paradigm in both online and lab settings. The findings suggest that the outcome density effect is present in both online and lab settings, supporting the hypothesis, suggesting that this paradigm can be effectively used online.

The findings of this study confirm that the outcome density effect is considerably large. Furthermore, participants gave higher judgement of causality when the probability of outcome was higher compared to when the probability of the outcome was lower. Wasserman (1990) found that people tend to weigh the cells of a contingency table (Table

22) according to, A > B > C > D, with more trials of cells A and D leading increased positive judgements whilst cell C and B counts decreased positive contingency judgements. Looking at Table 18, for the contingency calculations across the low and high conditions, there are more cell A trials for the high outcome density condition, which increases the positive judgements that participants make. There are more cell B trials for the low outcome density condition, which decreases the positive judgements. There are more cell C trials for the high-density condition, which decreases positive judgements, and more cell D trials for the low-density conditions which increases the positive judgements. Cell A trials are weighted more highly, as are cell B trials, consequently this explains the difference judgements seen for high and low outcome density conditions.

The findings also reflect the findings of Matute et al. (2007) who found evidence of the similar effects between online and lab settings. Blanco et al. (2013) found an outcome density effect, η_p^2=.40, which is similar in size to the findings of this study η_p^2=.53, although partial eta squared does adjust for according to other factors examined. Of course, the actual comparisons of effect sizes across papers are difficult, as effect sizes along with means and standard deviations are often not reported in these papers; not to mention methodological and putative population differences. Indeed, here I am interested in the effect sizes in passive contingency tasks, whilst the vast majority of research in contingency learning has focused on judgements of controls in active contingency learning tasks. For

Table 22 Replication of Table 1 - Visual presentation of the possible cue-outcome pairings in a contingency paradigm

	Outcome	
Cue	Present	Absent
Present	A	B
Absent	C	D

example, Jenkins and Ward's (1965) study examined the influence of judgements of control on different levels of outcome density and contingencies. Looking at high outcome density compared to the low outcome density condition in zero contingency, an effect size of 0.55 can be calculated for the judgement of control of the spectator who observed the participant actively take part in the experiment. The Jenkins and Ward (1965) experiment is not an exact like to like of the current experiment, however it is interesting to compare the effect sizes to one of the early experiments in this field. The different in the effect sizes may be expected the experimental design are very different.

Similarly, Msetfi et al. (2005) explored the influence of intertrial intervals (ITI), gender and depressive affect on judgements of control using low and high outcome densities. Calculating the effect sizes of the outcome density provides the following effect sizes: women short ITI d=0.11, men short ITI d=0.20, women long ITI d=0.23, men in long ITI d=1.26 (the effect sizes for participants with depressive affect is not calculated here). Here the effect sizes vary quite bit between men and women, and short and long ITI. I have not examined gender difference in outcome bias, as it was not part of my research question, and participants are not balanced across gender to run such tests. However, when considering the effect of ITI, the effect sizes do vary quite a bit between short and long ITI. The contingency paradigm within this experiment, and the following experiments, have been programmed such that ITI should not be an issue. An ITI is considered the absence of both the cue and outcome, which does occur in between trials. However, in the experiments of this project the cues and outcomes are clearly defined at the start of each block of trials, such that no cue and no outcome are still images that participants see. Hence any blank time between trials is not defined as an ITI in the same way as in Msetfi et al.

(2005) experiment. Hence when considering the effect sizes for the short ITI these are much lower than the effect sizes found in this experiment.

There was a three-way interaction, suggesting that outcome density effect may vary depending on the order of session by location. This may suggest that for this paradigm, there is a larger effect in the first session, compared to the second session, dependent on the location, with the largest effect in the online location when presented first. Despite the three-way interaction the effect was present throughout, and therefore conducting studies online is acceptable, and practically more feasible.

3.3.1 Theoretical Evaluation

Contingency account

The contingency account would predict participants to judge a zero contingency, when the programmed contingency is zero, regardless of probability of outcome. Taking Table 23 (which were the calculations used in experiment 1), as an example, for both the high and low outcome probabilities the overall contingencies are zero, as such the contingency account would predict participants to also make similar judgements. Even

Table 23 Outcome Density manipulation calculations used in experiment 1

Low outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	4	16	P(O/~C) = .2

High outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	16	4	P(O/C) = .8
Absent	16	4	P(O/~C) = .8

though the probability of the outcome occurring is far higher in one condition than the other. As such the contingency account does not account for the outcome density effects found in this thesis.

When considering Wasserman's extension to the contingency account, the extension has slightly different predictions. Looking at Table 24 , for the contingency calculations across the low and high outcome density conditions, for the high outcome density conditions, the majority of cells are present in cells A and C, outcome present trials, whilst for the low outcome density conditions the majority of trials are present in cells B and D, outcome absent trials. For the high-density condition even though there are 16 trials of A and C, cell A trials are weighted more and increase the positive judgements. For the low outcome density condition, whilst there are an equal number of B and D trials, cell B trials are weighted more highly, and cell B trials decrease the positive contingency. As such using

Table 24 Low- and High-density calculations, for cue and outcome density at Zero contingency

Table a *Low cue density at Zero Contingency*

Medication	Recovery	
	Present	Absent
Present	4	4
Absent	16	16
	P(C/O) = .2	P(C/~O) = .2

Δp = .2-.2=0

Table b *High cue density at Zero Contingency*

Medication	Recovery	
	Present	Absent
Present	16	16
Absent	4	4
	P(C/O) = .8	P(C/~O) = .8

Δp = .8-.8=0

Table c *Low outcome density at Zero Contingency*

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	4	16	P(O/~C) = .2

Δp = .2-.2=0

Table d *High outcome density at Zero Contingency*

	Recovery		
Medication	Present	Absent	
Present	16	4	P(O/C) = .8
Absent	16	4	P(O/~C) = .8

Δp = .8-.8=0

this extension it would be predicted that for these calculations participants would give more positive judgements in the high outcome density conditions, and less positive judgements in the low outcome density conditions. These predictions reflect the findings of the experiment.

As such the contingency account does not predict the findings of experiment 1, however with the Wasserman extension the findings of experiment 1 are accounted for.

Associative account

Overall, the associative perspective can predict cue and outcome density effects, however these effects would be expected to decrease in size the more trials that occur. This is because the overall strength of an association between a cue and an outcome is determined by the predictive value of the cue, rather than the overall number of coincidences between both events. In the short term, however, before associative strength reaches its asymptote, the 'coincidence' can have a temporary biasing effect. The outcome density effect is understood as a transient pre-asymptotic bias arising from accidental cue-outcome pairings before the context becomes an effective competitor (Shanks, 2007). If there is unequal salience for the events on the contingency, the associative models can predict outcome density effects.

The difference between the high and low outcome density effect could also be accounted for. With the high outcome density condition there are more opportunities for learning to take place as the outcome occurs in more trials in which an association could be formed. This would also suggest that learning would take place faster with the high outcome density condition. As such if there were more trials the differences between the conditions may not be seen.

Probabilistic account

Probabilistic models can predict more extreme levels of outcome density results in higher judgments of control in positive and negative contingency conditions. However, they cannot predict outcome density effects in zero contingency, as power is 0. When power zero is assumed, there is no relationship between the cue and outcome (Cheng & Novick, 1990). However, it could be argued that a zero contingency is misperceived by people as being slightly positive or negative rather than completely zero from the standpoint of the probabilistic account. Note that this account predicts that a 0-100 scale will artificially bias responses in a zero-contingency condition even in the absence of demand characteristics. Some participants will perceive a positive contingency (and will give a non-zero response), others will perceive some negative contingency and are not permitted to give a negative response. Even if the participants feel that they should respond with zero (even though they perceive some contingency) the mean will be above zero. Using this idea, the model could predict outcome density bias. This may for example occur in within-subjects designs where participants judge multiple contingencies. And due to comparisons between contingencies one could be perceived to be more or less positive than the other, leading to the outcome density effect. This could possibly explain the outcome density effect found in this experiment, although the same medical cover story was used with different stimuli for each condition.

3.3.2 Limitations

Due to the 2020 pandemic, the estimated sample size from the power analyses, could not be met, as data collection had to be stopped early. Hence only 58 participants completed both sessions of the study. However even with the sample size constraints a large effect size was found across the conditions. In addition to this, future studies using the same paradigm and effects, could use smaller samples sizes as well.

Due to the nature of the design, whilst participants were encouraged to leave a gap of 3 days between sessions, this was not always possible, and some participants only had a 2-day gap, whilst others had a 10-day gap. For future studies, to ensure more consistency a 3-day gap was used for all participants when there were two sessions. As with any online studies there will be some issues, one participant for example, lost their internet connection during the testing phase. However, the online sessions were conducted to simulate a more real-life scenario where people may receive information about different types of events and form beliefs and judgements.

3.4 Conclusion

Overall, these findings suggest that the outcome density effect is present in both online and lab settings, supporting the hypothesis, suggesting that this paradigm can be effectively used online. The three theoretical accounts are able to predict these findings to an extent.

Chapter 4: Outcome vs Cue density effects

The previous experiment replicated the outcome density effect in lab and online settings. Specifically, in zero contingency tasks, where there is no relationship between the cue and outcome, increasing the probability of the outcome, leads people to perceive a positive relationship between the cue and outcome. Experiment 2 complements these findings by taking a closer look at *cue* density effects in comparison to *outcome* density effects.

When considering density effects and their influence on associations, the focus has been predominantly on outcome density effects (Musca et al., 2010; J. C. Perales et al., 2005). Contrastingly, cue density effects (e.g., treatments for diseases) have not been extensively examined. Studies have found stronger cue density effects when cue density is high compared to when cue density is low in contingency paradigms (Matute et al., 2011b; Vadillo et al., 2011b). Few studies have simultaneously examined cue density and outcome density effects. In one such rare investigation, Blanco et al. (2013) used an allergy paradigm with zero contingency between cue and outcome and employed a 2 × 2 design with cue density (high vs low) and outcome density (high vs low) to examine the interactive effects of cue and outcome. They found an effect of outcome density in all conditions, however, there was only a small effect of cue density, and only when outcome density was high. This study manipulated cue and outcome density at the same time to be high and low, as such it is not fully clear what the difference in the sizes of effects are between cue and outcome density effects, when the other is held constant.

Overall, the cue density effect has not been studied as extensively as the outcome density effect within contingency learning or using contingency paradigms. However, these studies do often follow a similar format to outcome density studies. An effect of cue density has been found across many studies, however the effect seems to be more prominent in the presence of a higher outcome density, when cue density is also high.

The main goal of this chapter is to examine whether the outcome density effect and cue density effect are robust. This chapter also examined how methodological variations in the contingency learning paradigm affect the size of the outcome and cue density effects. This experiment investigates the density effect on contingency judgment of equivalent manipulations of cue density and outcome density, to see if, and how much, outcome effects and cue effects differ in size. As such this study adds to the literature by conducting a highly powered, within-subjects comparison of cue and outcome density effects. To provide further data on the relative contribution of cue density and outcome density to the over-estimation of contingency, in Experiment 2, I compare cue density and outcome density effects directly, by manipulating these density effects in a contingency paradigm, extending the work of Blanco et al. (2013). By investigating cue and outcome density effects at the same time, I investigate if there is something qualitatively different between them. If the effect sizes differ and/or one effect is conditional on the level of the other variable (with the other variable always having an effect) then this suggests that there are different things at play. To ensure equivalent comparison when cue is manipulated to be high or low, outcome density will be held at a medium density and vice-versa. Previous findings suggest cue effects are likely to be smaller than outcome effects, as such I predict outcome density effects to be larger than cue density effects.

4.1 Methods

This experiment investigated the influence of cue and outcome density on judgements of effectiveness of medication in treating illnesses, in a contingency paradigm, and sought to compare the sizes of the corresponding effects to each other. This study extends Blanco et al. (2013) by manipulating cue and outcome to be either high or low whilst holding the other constant. The task and stimuli are the same as outlined in the general methods.

4.1.1 Design

A 2 × 2 × 2 mixed design was used, with judgement of effectiveness of medication in treating illness as the dependent variable. Three factors were varied in the experimental design: manipulation of event (within subjects: cue density vs. outcome density) by density level (within subjects: high density vs. low density) by session order (between subjects: cue density first vs. outcome density first).

Cue and Outcome density manipulation.

Cue density was manipulated to be either low (0.2, Table 26) or high (0.8,Table 26) while maintaining a zero contingency. Outcome density was held at a medium density of 0.5, for both high and low cue density. Outcome density was manipulated to be either low (0.2,Table 26) or high (0.8, Table 26) while maintaining a zero contingency. Cue density was held at a medium density of 0.5, for both high and low cue density. Each session featured a low and a high-density block, with 40 trials. The block shown first was chosen randomly without replacement. Participants saw two of the four orders in Table 25, one for each session. Within each density type, cue or outcome, there were two levels to the densities, either high (0.8) or low (0.2). Participants completed two sessions. Half the participants completed the cue density condition during the first session, and outcome density condition

Table 26 Low- and High-density calculations, for cue and outcome density at Zero contingency

Table a *Low cue density at Zero Contingency*

Medication	Recovery	
	Present	Absent
Present	4	4
Absent	16	16
	P(C/O)=.2	P(C/~O)=.2

Δp=.2-.2=0

Table b *High cue density at Zero Contingency*

Medication	Recovery	
	Present	Absent
Present	16	16
Absent	4	4
	P(C/O)=.8	P(C/~O)=.8

Δp=.8-.8=0

Table c *Low outcome density at Zero Contingency*

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C)=.2
Absent	4	16	P(O/~C)=.2

Δp=.2-.2=0

Table d *High outcome density at Zero Contingency*

	Recovery		
Medication	Present	Absent	
Present	16	4	P(O/C)=.8
Absent	16	4	P(O/~C)=.8

Δp=.8-.8=0

Table 25 Order of Densities, Illness and Medication shown to participants for each session

			High Outcome Density Block	Low Outcome Density Block
Session 1				
	Cue or Outcome density	Order 1	Lindsay Syndrome and Batatrim	Marasia Syndrome and Lysopran
		Order 2	Marasia Syndrome and Lysopran	Lindsay Syndrome and Batatrim
Session 2				
	Cue or Outcome density	Order 1	Amplexia syndrome and Fexolise	Sawyer Syndrome and Quitrile
		Order 2	Sawyer Syndrome and Quitrile	Amplexia syndrome and Fexolise

in the second session, with the second half of participants completing outcome density condition first and cue density second (Figure 4). The manipulation of cue and outcome is referred to as event manipulation in this section. This manipulation is conceptualised as an event manipulated, as participants complete one session of cue density manipulation and one session of outcome density manipulation which is counterbalanced. This could potentially be conceptualised as a session order manipulation; however, participants do not complete both cue and outcome density manipulations in both sessions. As such a within subjects comparison of cue or outcome manipulation in session one and two cannot be done. Instead, it is thought that it would be better to conceptualise this as an event manipulation.

4.1.2 Procedure and measures

Figure 4 Participant Allocation to conditions across sessions

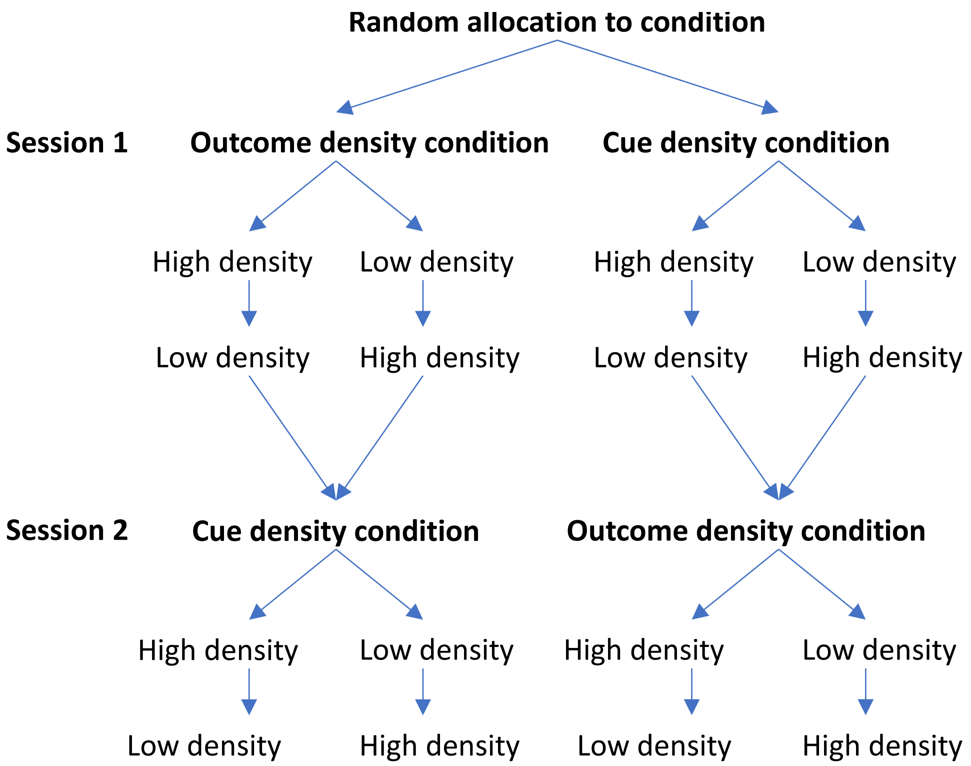

The procedure was similar to the previous study, with the exception that it took place online only. Participants signed up on the Prolific research participant platform (Prolific, 2020) and were prompted to complete the first session of the study. The task for each session took about 18 minutes to complete, and at the end of the task a short debrief was given. After 3 days, participants were invited to complete the second session of the study, which followed the same procedure as the first session. Participants were paid £2.50 for each completed session and had to complete both sessions to be paid.

4.1.3 Participants

Ninety-one participants were recruited to allow for participants to drop out between sessions; 91 completed the first session of the experiment, and 85 participants completed both sessions. Participants were over the age of 18 and fluent in English; see Table 27 for participant demographics.

4.2 Results

I ran a 2 × 2 × 2 mixed analyses of variance (ANOVA) to test the hypothesis that outcome density effects are larger than cue density effects. Judgement of effectiveness of the medication was the dependent variable. The three independent variables were: event manipulation order (within subjects: cue density vs. outcome density), density level (within subjects: high density vs. low density), and session order (between subjects: cue first vs.

Table 27 Study demographics across density types

	Second Session (N=85)	
	Cue Density (N=41)	Outcome Density (N=45)
Age, modal group, *(range)*	21-25,26-30 (*18-70)*	21-25, (*18-64)*
Female gender, *N (% of sample*)	31 (70.5*%)*	19 (*46.3%)*

outcome first). In this factorial analysis (and the equivalent analyses in our subsequent studies) it is the effects involving density level that are of prime interest for our hypotheses because these relate to the impact that cue/outcome density has on contingency judgments.

As predicted, there was a significant main effect of density, $F(1, 83)=101.91$, $p < .001$, $\eta_p^2=.55$. Participants judged the medication to be more effective in the high-density conditions (*M*=43.58, *SD*=19.06) than the low-density conditions (*M*=24.88, *SD*=17.22). There was a main effect of event manipulation order on effectiveness of judgements, $F(1, 83)=11.16$, $p=.001$, $\eta_p^2=.12$. Participants judged the medication to be less effective in the outcome density condition (*M*=31.09, *SD*=18.06) than the cue density condition (*M*=37.37, *SD*=18.38). I found a significant interaction effect for event manipulation order by density level, $F(1, 83)=8.93$, $p=.004$, $\eta_p^2=.097$. This suggests that the size of the density effect varies according to the order in which the event was presented. I therefore compared the size of the density effect, across both cue and outcome density (Table 28 and Figure 5). This analysis demonstrated that the density effect was nearly three times as large in the outcome density condition compared to the cue density condition.

To further understand the two-way interaction effect, both pairs of simple main effects were examined. Two paired-samples *t*-tests verified that the density effect was

Table 28 Mean (SD) judgement of medication effectiveness response by manipulation of event and density.

	Density Type	
	Cue	Outcome
High density response	43.45 *(20.80)*	43.74 *(16.89)*
Low density response	31.29 *(23.95)*	18.61 *(16.85)*
95% CI for mean difference	6.54, 17.77	19.41, 30.85
Density effect size *d*	0.54	1.49

present in both cue and outcome density conditions. Participants judged the effectiveness of the medication to be significantly higher in the high-density conditions than the low density conditions, both in the cue density condition, $t(84)=4.54$, $p< .001$, and in the outcome density condition $t(84)=8.74$, $p<.001$. Two paired-samples *t*-tests were conducted to verify whether there was an effect of event manipulation order in both the high density and low-density condition. Participants' judgements of the effectiveness of the medication did not differ significantly between cue and outcome density in the high-density condition, $t(84)=-.092$, $p=.927$. However, it did differ significantly in the low-density condition $t(84)=5.10$, $p<.001$, with participants giving lower judgements for the outcome density condition than for the cue density condition (Figure 5).

All other main effects and two-way interactions were small and non-significant, there being no main effect of session order, $F(1, 83)=.46$, $p=.467$, $\eta_p^2=.006$; no interaction effect of density by session order, $F(1, 83)=.92$, $p=.342$, $\eta_p^2=.011$; and no effect of event manipulation order by session order, $F(1, 83)=1.65$, $p=.202$, $\eta_p^2=.02$.

Figure 5 Mean effectiveness judgments by manipulation of event for high and low densities, with 95% confidence intervals

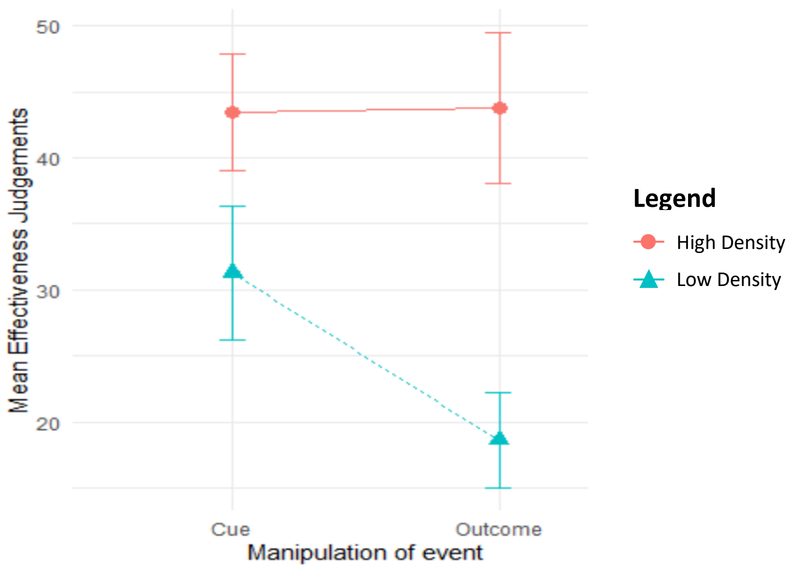

There was a significant three-way interaction for event manipulation order by density level by session order, $F(1, 83)=5.51$, $p=.021$, $\eta_p^2=.062$. This suggests that the size of the density effect may vary according to the combination of the event type (cue density or outcome density) that is manipulated and session order. Therefore, I compared the size of the density effect across each combination of session order, and manipulation of event (cue or outcome) (Table 29 and Figure 6). This analysis demonstrated that the density effect was present for both cue and outcome in the first session, with both effects being large. In the second session the outcome density effect was present again. However, there was no significant effect of cue density.

In sum, from the two-way and three-way interactions involving density level and event manipulation order I conclude that the effect of density level is smaller for cue density manipulations than it is for outcome density manipulations of an equivalent scale, such that a cue density effect on contingency judgment is only detected in some of the conditions of this experiment. The cue effects appear to be small enough that context or contrast effects can counteract the effect of cue density as indicated by the moderating effect of session

Table 29 Mean (SD) response by order of first session cue and outcome, and manipulation of cue or outcome

	First session		Second session	
	Cue Density Manipulation	Outcome Density Manipulation	Cue Density Manipulation	Outcome Density Manipulation
High density condition	46.95 *(21.49)*	47.59 *(26.19)*	40.18 *(19.81)*	39.61 *(27.33)*
Low density condition	27.83 *(20.34)*	19.39 *(16.69)*	34.52 *(26.72)*	17.78 *(17.19)*
95% CI for mean difference	11.69, 26.56	19.76, 36.65	-2.50, 13.81	13.94, 29.72
Outcome density effect size *d*	0.91	1.28	0.24	0.96

order. This is in line with previous research that examined the interactive effects of cue densities and outcome densities on contingency judgments, which found that cue effects where only found with high outcome densities (Blanco et al., 2013). Thus, the findings suggest that there is an effect of both cue and outcome density but that the cue effect is smaller than the outcome effect.

I ran a 2 × 2 × 2 mixed ANOVA with manipulation of event and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant main effect of density order during first session F(1, 83)=2.01, p=.152, η_p^2=.025. There were no interaction effect of manipulation of event × density order during first session, F(1,

Figure 6 Mean effectiveness judgments by first session order of event, and manipulation of event for high and low densities, with 95% confidence intervals

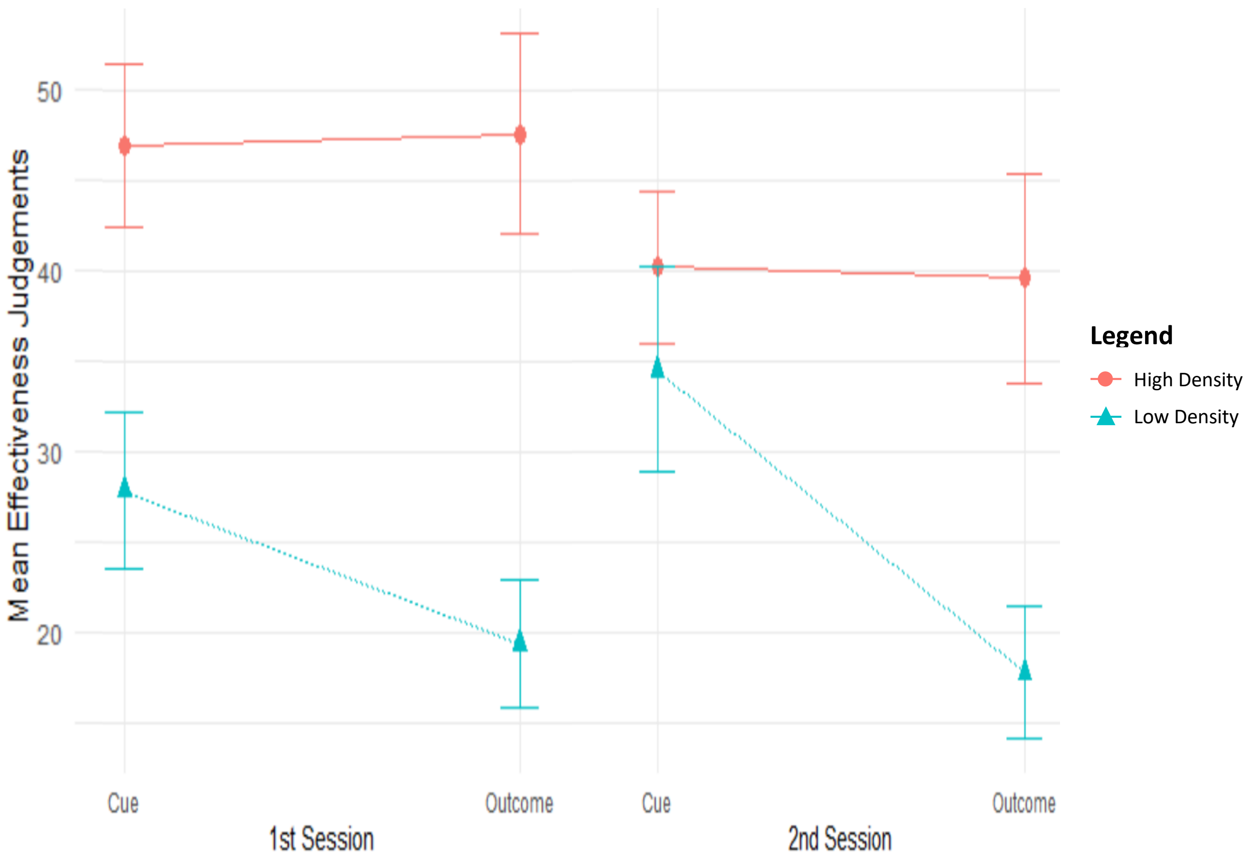

83)=.340, p=.562, η_p^2=.004, density × density order during first session, F(1, 83)=.547, p=.462, η_p^2=.007, manipulation of event × density × density order during first session, F(1, 83)=.020, p=.887, η_p^2=.000.

I ran a 2 × 2 × 2 mixed ANOVA with manipulation of event and density as within-subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant main effect of density order during second session F(1, 83)=1.92, p=.170, η_p^2=.023. There was a significant interaction effect of manipulation of event × density order during second session, F(1, 83)=5.63, p=.020, η_p^2=.064 (Table 30). There was no significant interaction effect of density × density order during second session, F(1, 83)=3.72, p=.057, η_p^2=.043, manipulation of event × density × density order during second session, F(1, 83)=1.38, p=.243, η_p^2=.016.

4.3 Discussion

This study examined the influence of cue density and outcome density on judgements of effectiveness of medication in treating illnesses, in a contingency paradigm online. In particular, the study tested the potential difference in impact of cue density and outcome density effects. I replicated the findings for outcome and cue density effects in contingency learning from Blanco et al. (2013), as well as the outcome density findings from Experiment 1 in online settings. My finding shows that the outcome density effects are larger than cue density effects on average. The relative difference in size appears to depend on other factors. Thus, the findings suggest that there is an effect of both cue and outcome density, however the cue effect seems to be smaller than the outcome density. In addition, the cue effect does not seem to be present under all conditions tested in this study.

The aim of this study was to examine if the outcome density effect was larger or the same size as the cue density effect. If both the cue and outcome have the same effect on judgements of effectiveness, I would have expected to see results that show effect sizes to be of similar sizes, however the results do not show this. This may suggest that the outcome density effect is more of a robust effect regardless of context, whilst cue density effect is not as a robust of an effect.

Table 30 Mean (SD) response by order of second session order density, and manipulation of event and density

	Low density first				High density first			
	Cue	Outcome	95% CI for mean difference	Event density effect size *d*	Cue	Outcome	95% CI for mean difference	Event density effect size *d*
Low density response	32.58 *(24.41)*	18.38 *(16.75)*	7.60, 20.80	0.68	29.46 *(23.51)*	18.94 *(17.23)*	2.72, 18.31	0.51
High density response	49.80 *(18.44)*	44.32 *(23.94)*	-2.33, 13.89	0.26	34.37 *(20.84)*	42.91 *(30.97)*	-19.18, 2.09	0.32
95% CI for mean difference	9.75, 24.69	18.83, 33.05			-3.43, 13.25	14.07, 33.88		
Outcome density effect size *d*	0.80	1.26			0.22	0.96		

The experiment compared the magnitudes of the outcome density effect and the cue density effect. Consistent with the findings from Blanco et al. (2013), who conducted a similar study but with a between-subjects design with only 27 participants per condition, the outcome density effect is more of a robust effect regardless of other manipulations, whilst the cue density effect is less robust and more dependent on other experimental manipulations. Thus, the outcome density effect was large and present in all conditions, whilst the cue density effect was smaller and dependent on session order. The cue density effect was present in the first session but not in the second session. This may result from participants in their second session comparing the perceived effectiveness of the medication in the (current) cue density condition against the perceived effectiveness of the medication in the (previous) outcome density condition.

When compared to the outcome density means in the second session, the high-density results are very similar, whilst the low-density condition, is where the cue density differs a great deal from the outcome density condition. This seems to be for both first and second session. This may suggest that the cue density relies on a higher outcome density, to show an effect, or to show a bigger effect. And with a lower outcome density, the effect diminishes. With these factors combined, second session, learning effects and lower density, may explain the lack of cue density effect in the second session.

This possibility is suggested by the fact that the mean effectiveness judgments are less extreme for the cue density condition in Session 2 (Figure 6). That is, mean effectiveness with high event density is lower for cue density manipulation than for outcome density manipulation, while the reverse is true for mean effectiveness with low event density. If such comparisons were made, I would expect an outcome density effect in the second session that is larger than the cue density effect seen in Session 1. This is not apparent in Figure 6, though it is true that for low event density, mean effectiveness is more extreme (i.e., lower) for the outcome density manipulation than for the cue density manipulation.

Also, to note, Blanco et al. (2013) looked at the interaction of low and high outcome and cue densities, as such they did not look at cue density effects when outcome when held at medium. With our results, I can additionally add to the literature, that cue density effects can occur with medium outcome density.

4.3.1 Theoretical Evaluation

Contingency account

According to the contingency account, as all four conditions of this experiment are zero contingencies, participants would be expected to make similar judgements across all conditions. See Table 26 for calculations of the experiment.

When considering the contingency calculations of Table 26, direct comparisons of the cell types can be made between the two low density conditions and the two high density conditions. According to Wasserman (1990) people tend to rank the cells weights as, A > B > C > D, with more trials of cells A and D leading increased positive judgements whilst cell C and B counts decreased positive contingency judgements. When examining the low contingency conditions for both cue (Table 26a) and outcome density (Table 26c), cells A and D are the same, whilst cells B and C are reversed. Comparing conditions, there are more counts of cell B for the low outcome density condition compared to the low cue density condition, which participants weighed more than the cell C trials. However, this higher weighting would decrease the positive contingency judgements, as cells B and C decrease positive contingency judgements. This matches the findings of this experiment, whereby for the low-density condition participants have given more positive judgements for the cue density condition than the outcome density condition.

When examining the high contingency conditions for both cue (Table 26b) and outcome density (Table 26d), again cells A and D have the same frequency counts, whilst cells B and C are reversed. For the high cue density condition there are more counts of cell B, than there are for the high outcome density condition. As such lower judgements would be expected for the high cue density condition, than the high outcome density condition, which also reflects the results of the current study. Here the Wasserman extension is able to account for the findings of experiment 2.

Similar cross-session effects have been observed in other studies of judgment and preference, including when—as I did here—sessions are separated by at least one day (Rakow et al., 2020). I argue that such effects deserve further attention in contingency learning to better understand the role relative or comparative judgments play when people assess contingency from observed or remembered samples of information (Stewart et al., 2006). Whatever the mechanism responsible, it does seem that compared to the outcome density effect, the cue density effect relies on larger density manipulation to show an effect, or to show a large effect.

Associative account

Whilst the associative account could account for the outcome density effects until the asymptote is reached, it cannot account for cue density effects, as the equations $\Delta V =$

$\alpha * \beta * (\lambda - V_{total})$, only accounts for the presence and absence of the outcome with λ. It does not account for the presence and absence of the cue. As such when cue was manipulated in this experiment, outcome density was held constant across high and low cue density, the Rescorla-Wagner model would predict the same level of associative strength for both high and low cue density conditions. An assumption of the model is that learning only occurs if the cue is present, otherwise learning cannot occur, and associative strength cannot be formed. If the model is changed such that it accounts for the presence and absence of the cue as well, $\Delta V = \alpha * \beta * ((\lambda_{\alpha} * \lambda_{\beta}) - V_{total})$, then it would be able to account for cue density effects in the same way as outcome density effects.

However, the size of the cue and outcome density effects were different in this experiment. This could perhaps be accounted for by assuming differential salience of the cue and outcome. If the cue is thought to be more salient than the outcome, then smaller cue density effects may be expected, as the low cue density conditions would reach the asymptote quicker than the low outcome density conditions. As such the associative perspective could explain the findings of this experiment.

Probabilistic account

As mentioned above, if it is argued that zero contingencies are in fact misperceived as being slightly negative or positive, then cue density effects may be predicted in a similar way to outcome density effects. However even with this assumption, the difference in the size of effects of cue and outcome density effects cannot be accounted for.

4.4 Conclusion

Our experiments therefore clarify two key results that any theory of contingency learning must be able to explain. First, it clarifies that density effects in contingency learning are larger for outcome density manipulations than for cue density manipulations. Also it furthers supports the weightings indicated by (Wasserman, 1990) for the different cell types for contingency tables.

Chapter 5: Cue and Outcome Interchangeability

As an extension to the research, in Experiments 3 and 4, I manipulate the order of cue and outcome presentation to compare the effects of the first seen event vs. the second seen event. This novel extension reflects that while, logically, a cause must precede an outcome, sometimes people observe or learn of an outcome *before* they see or hear of its possible causes. For example, a medic may observe signs of infection in patients before learning what potential causes the patients were exposed to, or a manager might see that their salesforce has made a number of successful sales before they find out what steps their workers took in an attempt to achieve those successes. These event presentation order manipulations permit exploration of interaction effects involving *both* density level and event presentation order. I can therefore consider whether the effect of cue or outcome density on contingency judgment is altered by reversing the order in which event information is seen. Aside from cues and outcomes to be functionally different, they typically also differ in temporal order. To contrast cues and outcomes on the basis on their roles, then the temporal confound would need to be manipulated. Manipulating the order of cues and outcomes allows one to do so, and the impact of this temporal ordering can be directly assessed. For example, it may be that the effects are the same size, provided that they occupy the same temporal position (e.g., both presented together), but not when occupying different temporal positions (e.g., one first, one second), as is the default in studies so far. I can therefore assess whether previous observations that outcome density effects are larger than cue density effects might be better understood as reflecting that density effects are larger for whichever event is seen second rather than first. As such I wish to explore the effect of this temporal ordering, and do not know how this will affect the cue and outcome density effects. If information presentation order has an effect on outcome density, then it would be expected that the size or magnitude of the outcome density effect would differ between the two conditions. It would be expected that the density effect would be larger in the outcome-presented first condition, as the cue information would be interpreted through the lens of the outcome information. The main goal of this chapter is to examine whether the outcome and cue density effects are robust. This chapter also examined how methodological variations in the contingency learning paradigm affect the size of the outcome and cue density effects.

5.1 Experiment 3: Cue and outcome interchangeability—Outcome Density

Experiment 3 investigated whether the impact of cue and outcome density on judgements of effectiveness of medication depend on the order in which cue and outcome are presented. Experiment 3 investigated whether this manipulation of cue-outcome presentation order moderated the effect of *outcome density* on contingency judgment.

5.1.1 Method

A similar design to the one outlined in the General Methods was employed. The order of event presentation was manipulated, along with outcome density. In the trials of one session, the cue was shown first followed by the outcome—thereby following the standard order of presentation: fixation cross, unhealthy patient image, treatment cue image then treatment outcome image. In the trials of the other session, the outcome was shown first followed by the cue – the image order being fixation cross, unhealthy patient image, treatment outcome image then treatment cue image. See Figure 7 for flowchart of the event presentation order manipulation.

A 2 × 2 × 2 mixed design was used, with judgement of effectiveness of medication in treating illness as the dependent variable. Three factors were varied in the experimental design: the order of event presentation (within subjects: cue then outcome vs. outcome then cue), the corresponding outcome density level (within subjects: high vs. low density), and the session order (between subjects: cue then outcome in first session vs. outcome then cue in first session).

Design

Figure 7 Event order presentation for cue and outcome order manipulation

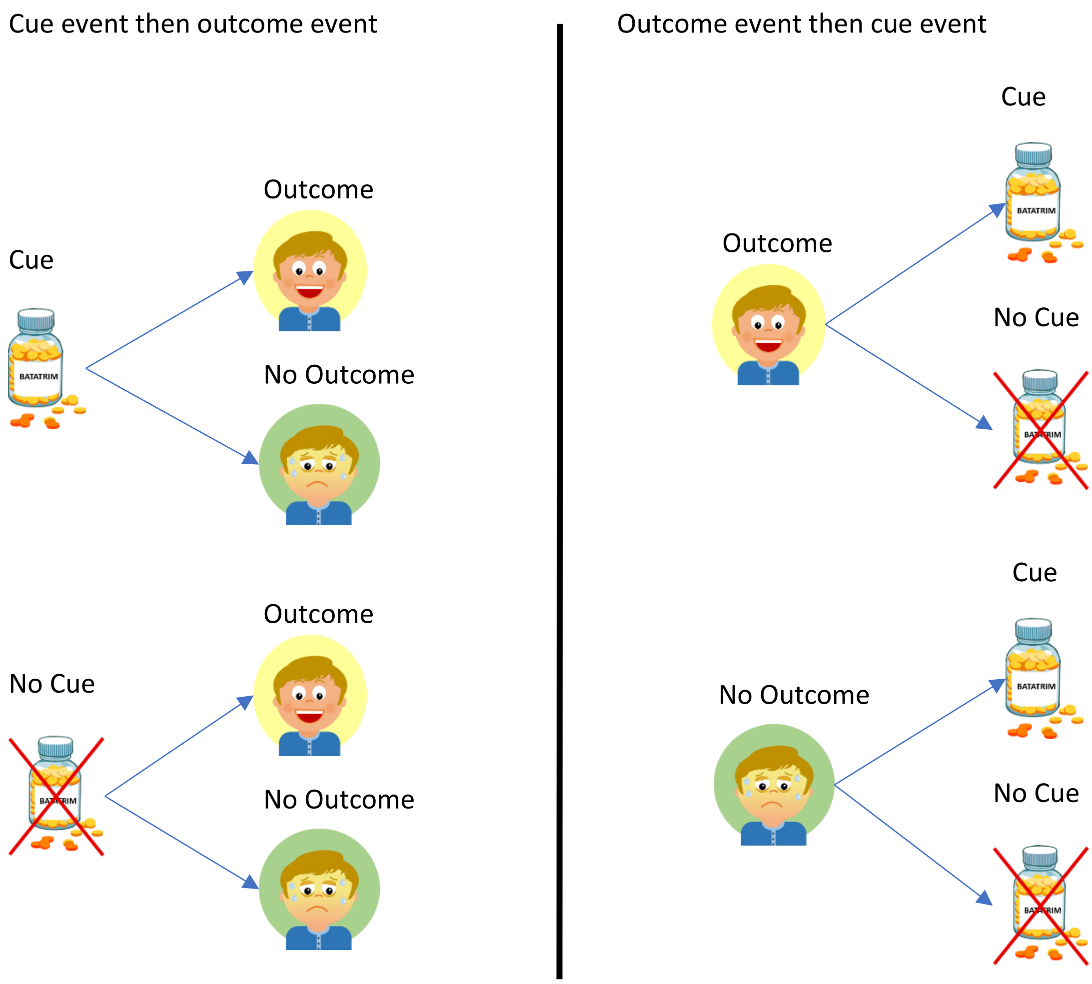

Order of cue and outcome presentation was manipulated, to either show the cue first followed by the outcome or the outcome first followed by the cue, in a session. Participants completed two sessions, 3 days apart (Figure 8). To reduce any carry-over or learning effects across sessions, different medication and illness names were used for each session, assigned at random. Outcome density was manipulated to be either low or high while maintaining a zero contingency. Cue density was held at a medium density of 0.5, for both high and low outcome density. For each session there was a block of 40 high density trials and 40 low density trials. The block shown first was chosen randomly without replacement. To reduce any carry over or learning effects across sessions, different medication and illness names were used for each session.

Participants

Figure 8 Participants Allocation to conditions across session

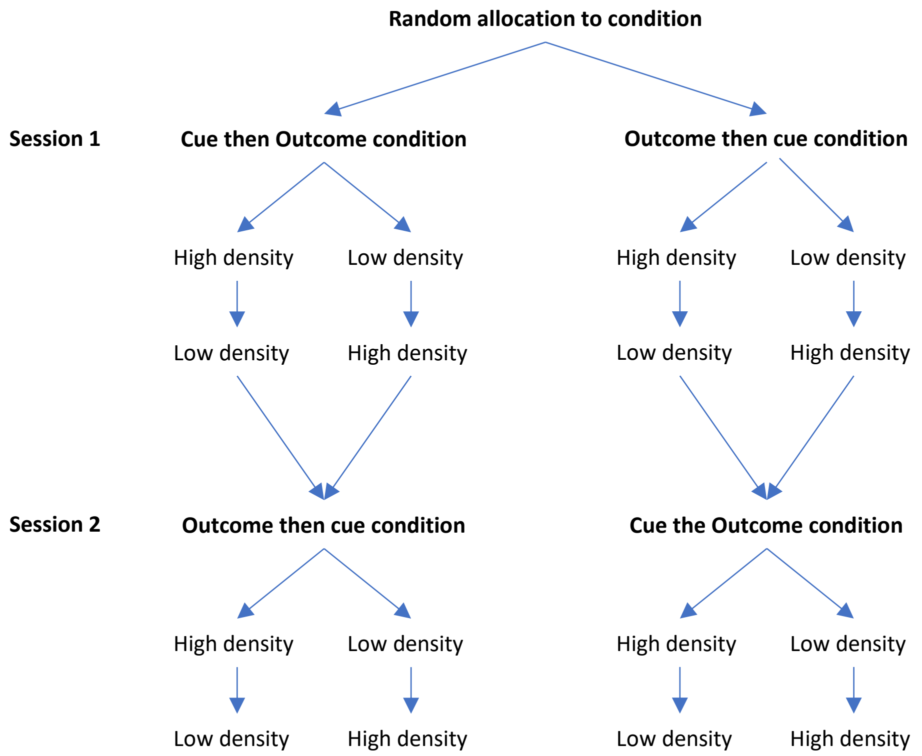

Participants were 110 Prolific members; 107 participants completed the first session, and 96 participants completed both sessions. Three participants withdrew their participation during the session, as such their data was not included in the analyses.

5.1.2 Results and Discussion

I ran a 2 × 2 × 2 mixed ANOVA. The dependent variable was judgement of medication effectiveness. The three independent variables were: event presentation order (within subjects: cue then outcome vs. outcome then cue); outcome density (within subjects: high vs. low density); and session order (between subjects: cue-then-outcome in first session vs. outcome-then-cue in first session). Outcome density effects were expected, in both event presentation order conditions.

As predicted, and replicating the findings of Experiment 1, there was a large and statistically significant main effect of outcome density, $F(1, 94)=141.34$, $p < .001$, $\eta_p^2=.60$. Participants judged the medication to be less effective in the low-density condition (*M*=19.16, *SD*=11.51) than the high-density condition (*M*=43.83, *SD*=22.16). There was no significant main effect of event presentation order on judgements of effectiveness, $F(1, 94)=2.64$, $p=.107$, $\eta_p^2=.027$. Judgments of medication effectiveness did not differ significantly between the cue first (*M*=30.13, 95% CI [19.77, 31.05], *SD*=16.82) and outcome first (*M*=32.86, 95% CI [19.03, 28.70], *SD*=16.42) conditions. As this is a well-powered experiment with reasonably narrow CIs for means and mean differences, these findings suggest that event presentation order makes little difference to the judgement of medication effectiveness.

There was no interaction effect of event presentation order by density, *F*(1, 94)=.21, *p*=.647, η_p^2=.002; the size of the density effect did not vary significantly by event presentation order (Table 31 and Figure 9). Thus, neither here, nor for the main effect of event presentation order, did I find statistically significant evidence that event presentation order made a difference. Additionally, there was no significant three-way interaction for event presentation order by density by session order, *F*(1, 94)=.001, *p*=.972, η_p^2=.000. See Figure 10.

Table 31 Experiment 3 and 4 mean (SD) judgement of medication effectiveness response by event presentation order and density

	Experiment 3: Outcome density		Experiment 4: Cue density	
	Cue First	Outcome First	Cue First	Outcome First
High density response	42.93 *(26.96)*	44.70 *(23.99)*	46.55 *(19.10)*	46.87 *(21.31)*
Low density response	17.52 *(14.94)*	20.83 *(15.65)*	35.70 *(20.33)*	31.58 *(23.23)*
95% CI for mean difference	19.77, 31.05	19.03, 28.70	6.07, 15.61	9.63, 20.95
Density effect size *d*	1.17	1.18	0.55	0.69

Figure 9 Mean effectiveness judgments by event order presentation for high and low densities, with 95% confidence intervals

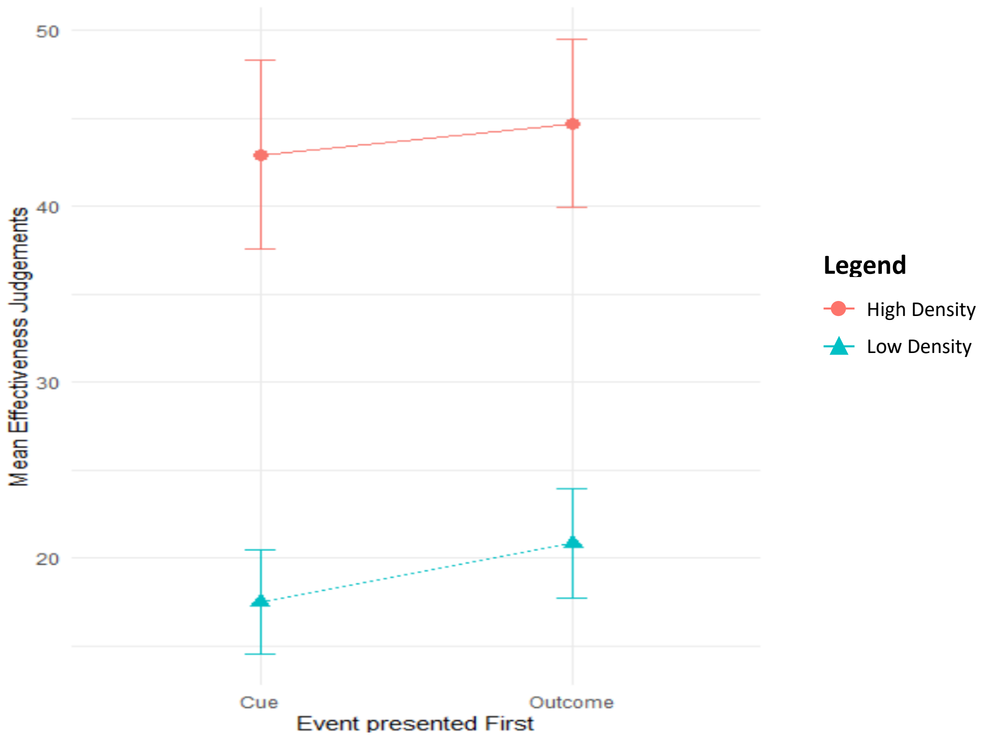

None of the other effects in this ANOVA have any bearing on the hypotheses of this experiment, but they are reported for completeness. There was no significant main effect of session order, $F(1, 94)=.000$, $p=.997$, $\eta_p^2=.000$; and no significant interaction effect for density by session order, $F(1, 94)=.071$, $p=.791$, $\eta_p^2=.001$, or for event presentation order by session order, $F(1, 94)=3.29$, $p=.073$, $\eta_p^2=.034$.

I ran a 2 × 2 × 2 mixed ANOVA with manipulation of event and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant main effect of density order during first session $F(1, 83)=2.01$, $p=.152$, $\eta_p^2=.025$. There were no interaction effect of manipulation of event × density order during first session, $F(1, 83)=.340$, $p=.562$, $\eta_p^2=.004$, density × density order during first session, $F(1, 83)=.547$,

Figure 10 Mean effectiveness judgments by session, for event order presentation by high and low densities, with 95% confidence intervals

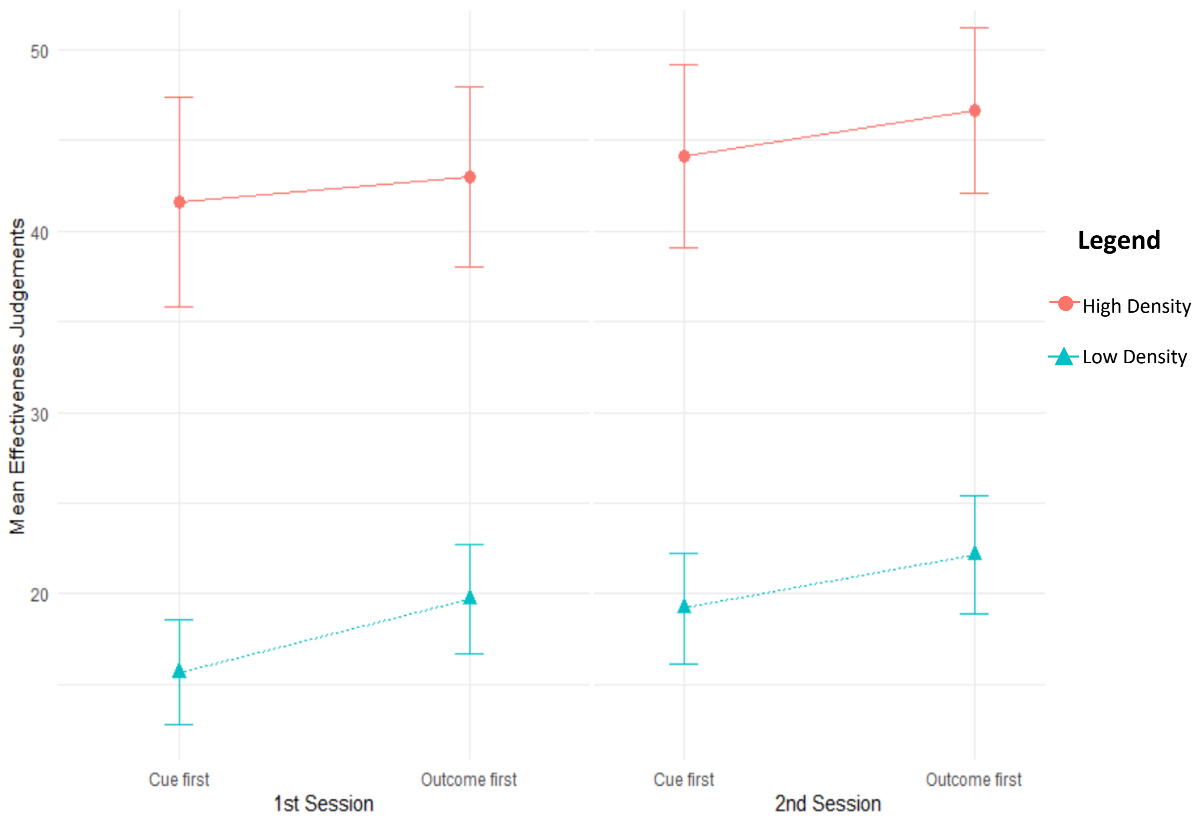

p=.462, η_p^2=.007, manipulation of event × density × density order during first session, *F*(1, 83)=.020, *p*=.887, η_p^2=.000.

I ran a 2 × 2 × 2 mixed ANOVA with manipulation of event and density as within-subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant main effect of density order during second session *F*(1, 83)=1.92, *p*=.170, η_p^2=.023. There was a significant interaction effect of manipulation of event × density order during second session, *F*(1, 83)=5.63, *p*=.020, η_p^2=.064. There was no significant interaction effect of density × density order during second session, *F*(1, 83)=3.72, *p*=.057, η_p^2=.043, manipulation of event × density × density order during second session, *F(*1, 83)=1.38, *p*=.243, η_p^2=.016 see Table 32.

Our primary interest in this experiment was whether event presentation order moderated the outcome density effect. I found no significant evidence for such moderation. However, null-hypothesis significance testing is not designed to evaluate evidence for the

Table 32 Experiment Mean (SD) response by session, and manipulation of cue or outcome

	First session		Second session	
	Cue First	Outcome First	Cue First	Outcome First
High density response	41.58 *(28.94)*	42.98 *(24.99)*	44.12 *(25.31)*	46.64 *(22.92)*
Low density response	15.64 *(14.44)*	19.69 *(15.13)*	19.18 *(15.32)*	22.13 *(16.30)*
95% CI for mean difference	17.25, 34.61	16.50, 30.09	17.30, 32.58	17.39, 31.64
Outcome density effect size *d*	1.13	1.13	1.19	1.23

absence of an effect. I undertook two further analyses to evaluate how well the data aligned with the null hypothesis that the outcome density effect is not affected by event presentation order.

First, I used confidence intervals to identify the plausible size of the effect that event presentation order has on the outcome density effect (Greenwald, 1975). To do so, I computed the outcome density effect for each participant (i.e., difference in judgment, for high outcome density minus low outcome density). The effect size for the difference in effect (cue first minus outcome first) was d=.047, 95% CI [-0.153, 0.247]. From this I conclude that if event presentation order moderates the size of the outcome density effect, that moderating effect is likely small, $|d| < 0.25$.

Second, I computed a Bayes Factor to assess the evidence for the point null hypothesis of zero effect (d=0) against an alternative prior (Dienes, 2014). Informed by our data from Experiment 2 and reasoned assumptions about the plausible effect of reversing the event presentation order (Dienes, 2019) I used a normally distributed alternative prior with a mean of d=0.75 and SD of 0.75 in units of d. These parameter values were chosen to be half of the outcome density effect size observed in Experiment 1. This prior places the greatest density at the point representing the possibility that half of the observed outcome density effect is contingent on the cue preceding the outcome, while also representing a range of other possibilities including an increased outcome density effect when the cue follows the outcome or a reverse density effect.[3] BF_{01}=10.24, meaning that the odds of the

[3] Our prior has 68% of its density on the interval [0, 1.5] where d=1.5 was the effect observed in Experiment 1. This reflects an assumption that the most plausible effect of presenting the outcome before the cue is to reduce the size of the outcome density effect. I made this assumption on the basis that most theories of contingency learning assume, explicitly or

data are around 10 times greater under H_0 than under the alternative prior. This is conventionally described as 'strong' evidence for H_0.

5.2 Experiment 4: Cue and outcome interchangeability—Cue Density

Experiment 4 investigated whether the impact of cue and outcome density on judgements of effectiveness of medication depends on the order in which cue and outcome are presented. Experiment 3 investigated whether this manipulation of cue-outcome presentation order moderated the effect of *outcome density* on contingency judgment, whilst—in a symmetrical fashion—Experiment 4 investigated if the same manipulation of cue-outcome presentation order moderated the effect of *cue density* on contingency judgment.

5.2.1 Methods

A similar design to the one outlined in the General Methods was employed. The order of event presentation was manipulated, along with cue density rather than outcome density. In the trials of one session, the cue was shown first followed by the outcome—thereby following the standard order of presentation: fixation cross, unhealthy patient image, treatment cue image then treatment outcome image. In the trials of the other session, the outcome was shown first followed by the cue—the image order being fixation cross, unhealthy patient image, treatment outcome image then treatment cue image.

---

implicitly, that the outcome follows the cue (see General Discussion). Therefore, reversing this event ordering is likely to disrupt the process of evaluating contingency (e.g., increasing judgmental noise thereby decreasing the size of the effect).

A 2 × 2 × 2 mixed design was used, with judgement of effectiveness of medication in treating illness as the dependent variable. Three factors were varied in the experimental design: the order of event presentation (within subjects: cue then outcome vs. outcome then cue), the corresponding cue density level (within subjects: high vs. low density), and the session order (between subjects: cue then outcome in first session vs. outcome then cue in first session).

Design

Order of cue and outcome presentation was manipulated, to either show the cue first followed by the outcome or the outcome first followed by the cue, in a session. Participants completed two sessions, 3 days apart, as seen in Figure 8. To reduce any carry-over or learning effects across sessions, different medication and illness names were used for each session, assigned at random. Cue density was manipulated to be either low or high while maintaining a zero contingency. Outcome density was held at a medium density of 0.5, for both high and low cue density. For each session there was a block of 40 high density trials and 40 low density trials. The block shown first was chosen randomly without replacement. To reduce any carry over or learning effects across sessions, different medication and illness names were used for each session.

Participants

I recruited 100 participants from Prolific for Experiment 4. All participants completed the first session of the experiment, and 97 participants completed both sessions.

5.2.2 Results and Discussion

I performed a 2 × 2 × 2 mixed ANOVA. The dependent variable was judgement of medication effectiveness. The three independent variables were: event presentation order (within subjects: cue then outcome vs. outcome then cue); cue density (within subjects: high vs. low density); and session order (between subjects: cue-then-outcome in first session vs. outcome-then-cue in first session). As per Experiment 3, it is the interaction effects involving *both* event presentation order and cue density that are important for the main contribution of this experiment to our investigation. In line with Experiment 2, cue density effects are expected in both event presentation order conditions, but these are expected to be smaller, on average, than the outcome density effects in Experiments 2 and 3.

Replicating the findings of Experiment 2, there was a statistically significant main effect of cue density, $F(1, 95)=42.09$, $p < .001$, $\eta_p^2=.307$. Participants judged the medication to be less effective in the low-density condition *(M=*33.63, *SD=*17.34) than the high-density condition *(M=*46.68, *SD=*15.34). There was no significant main effect of event presentation order on judgments of effectiveness $F(1, 95)=.89$, $p=.887$, $\eta_p^2=.009$. Participant judgment of medication effectiveness did not differ significantly between the cue first *(M=*41.12, *SD=*15.88, 95% CI [37.92, 44.32]) or outcome first *(M=*39.19, *SD=*17.02, 95% CI [35.75, 42.62]) conditions.

There was no statistically significant event presentation order by density interaction, $F(1, 95)=1.68$, $p=.198$, $\eta_p^2=.017$ (Figure 11). The size of the density effect did not vary significantly by event presentation order (Table 31). Together with the absence of a main effect of event presentation order, this suggests that there is little evidence that event presentation order made a difference. There was no significant three-way interaction of event presentation order by density by session order, $F(1, 95)=.12$, $p=.731$, $\eta_p^2=.001$, see Figure 12.

None of the other effects in this ANOVA analyses have any bearing on the hypotheses of this experiment, but they are reported for completeness. There was no significant main effect of session order, $F(1, 95)=2.09$ $p=.152$, $\eta_p^2=.021$; and no significant interactions for density by session order, $F(1, 95)=1.26$, $p=.265$, $\eta_p^2=.013$, or event presentation order by session order, $F(1, 95)=2.66$, $p=.106$, $\eta_p^2=.027$.

I ran a 2 × 2 × 2 mixed ANOVA with event presentation order and density as within-subjects factors and density order during first session (low first or high first) as independent variables, and effectiveness judgements as dependent variable, to investigate whether density order during first session has an order effect. There was a significant main effect of density $F(1, 94)=140.95$, $p < .001$, $\eta_p^2=.60$. There was a significant main effect of first session

Figure 11 Mean effectiveness judgments by event order presentation for high and low densities, with 95% confidence intervals.

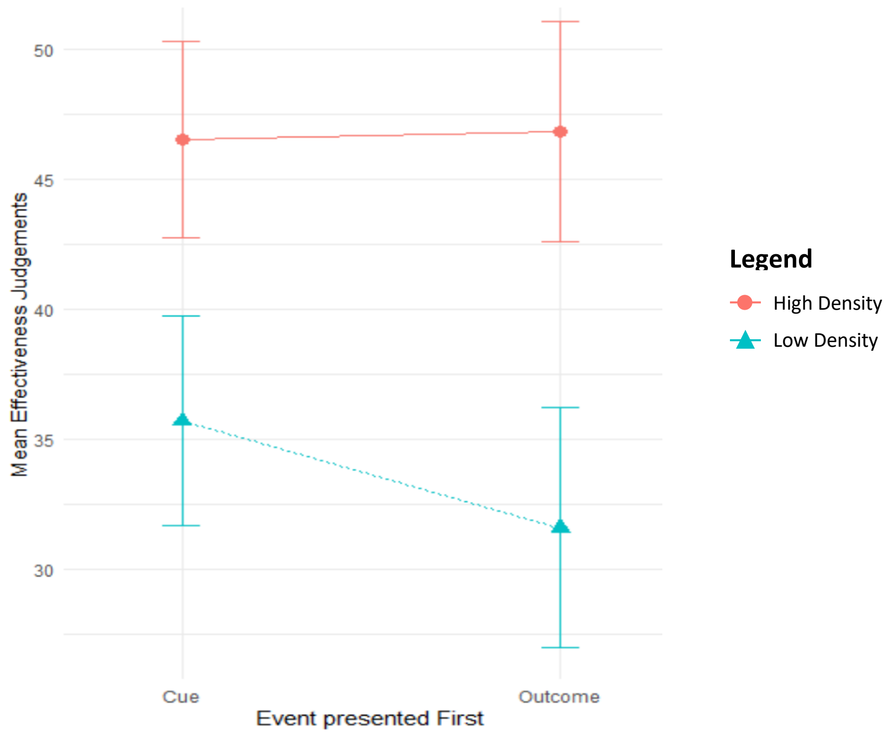

density order $F(1, 94)=4.22$, $p=.043$, $\eta_p^2=.043$, participants gave lower judgements if they were in the low density first condition *(M*=28.35, *SD*=20.60), compared to the high density first condition *(M*=34.27, *SD*=19.34), suggesting there may have been an anchoring effect present that is based on the first judgement they made. There was an interaction effect of event presentation order × density order during first session, $F(1, 94)=12.05$, $p=.001$, $\eta_p^2=.114$, suggesting that responses varied for event presentation order according to which density participants were shown first. There was a significant 3-way interaction of event presentation order × density × density order during first session, $F(1, 94)=.9.30$, $p=.003$, $\eta_p^2=.090$. There was no interaction effect of density × density order during first session, $F(1, 94)=1.14$, $p=.290$, $\eta_p^2=.012$.

Figure 12 Mean effectiveness judgments by session, for event order presentation by high and low densities, with 95% confidence intervals

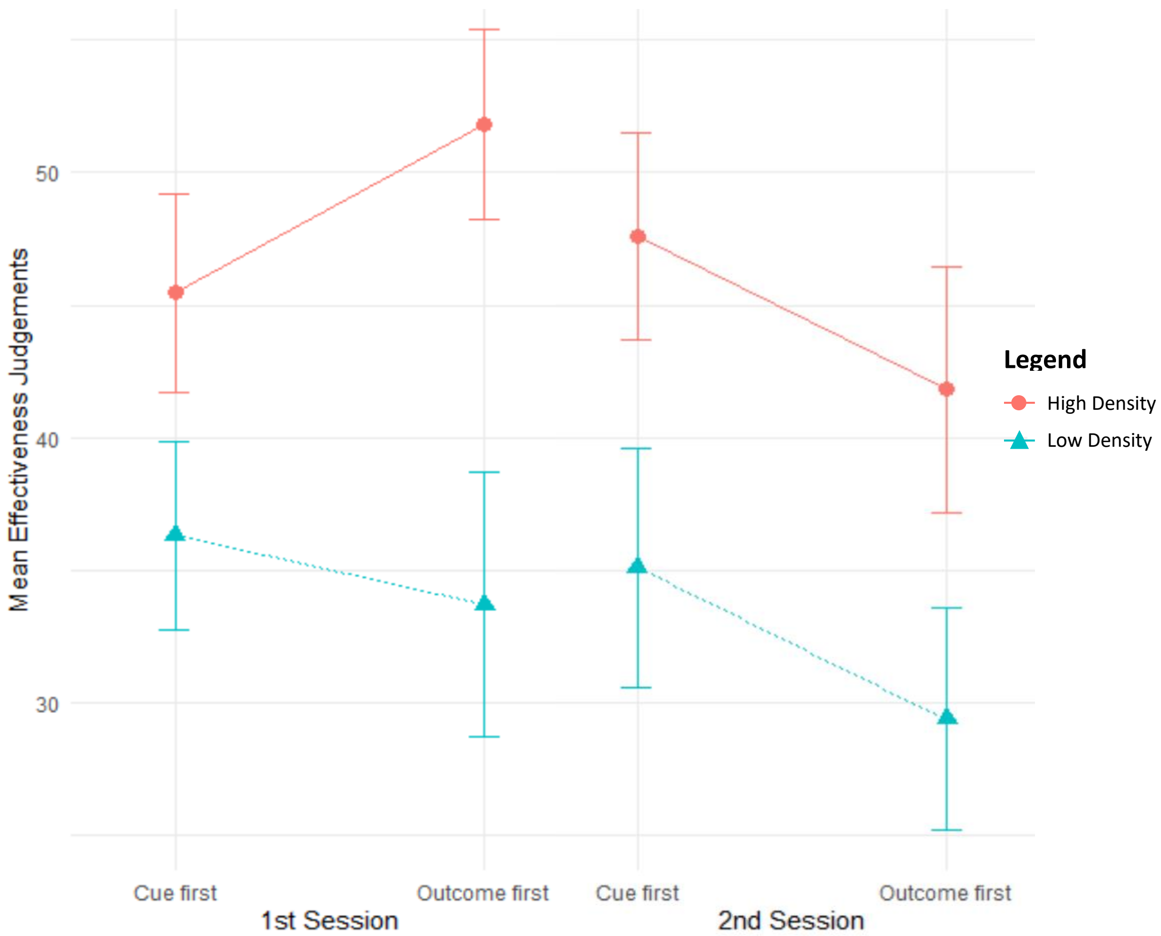

I ran a 2 × 2 × 2 mixed ANOVA with event presentation order and density as within-subjects factors and density order during second session (low first or high first) as independent variables, and effectiveness judgements as dependent variable. There was no significant main effect of density order during the second session $F(1, 94)=1.18$, $p=.280$, $\eta_p^2=.012$. There were no significant interaction effects of event presentation order × density order during the second session, $F(1, 94)=.80$, $p=.374$, $\eta_p^2=.008$, or density × density order during the second session, $F(1, 94)=.043$, $p=.836$, $\eta_p^2=.000$, manipulation of order of event × density × density order during the second session, $F(1, 94)=.072$, $p=.789$, $\eta_p^2=.001$, see Table 33.

I used approaches equivalent to those I used in Experiment 3 to re-examine the absence of a significant moderating effect of event presentation order on the cue density effect. This moderating effect of event presentation order was small, $d=-0.13$, with a 95% CI of [–0.33, 0.07] suggesting that the moderating effect is unlikely to be greater than small-to-medium in size. The Bayesian analysis again pitted a point null hypothesis of $d=0$ against a

Table 33 Experiment 4 Mean (SD) response by order of session, and manipulation of cue or outcome

	First session		Second session	
	Cue First	Outcome First	Cue First	Outcome First
High density response	45.48 *(18.80)*	51.82 *(18.10)*	47.59 *(19.53)*	41.81 *(23.27)*
Low density response	36.31 *(17.82)*	33.71 *(25.17)*	35.10 *(22.70)*	29.40 *(21.12)*
95% CI for mean difference	2.32, 16.02	9.29, 26.91	5.64, 19.34	5.10, 19.73
Outcome density effect size *d*	0.50	0.83	0.59	0.56

normal alternative prior with *M* and *SD* of half the effect size found in Experiment 2 (*d*=0.27). The data were again more probable under the null than under the alternative prior, with the Bayes Factor of BF_{01}=3.21 representing mere ‘anecdotal’ evidence for H_0.

5.3 General Discussion of experiment 3 and 4

Experiments 3 and 4 examined the influence of the order of cue and outcome presentation on judgements of effectiveness of medication, in a contingency learning paradigm. The findings suggest that event presentation order makes little or no difference to the judgement of medication effectiveness or to the size of the outcome density and cue density effect. Across the experiments, the size of effects for cue and outcome density were relatively consistent.

Experiments 3 and 4 examined whether the density effect differed when the presentation order of the cue and outcome was manipulated, to see if the outcome effect is due to frequency of outcome or being seen second. I find it intriguing that this manipulation made no difference to the presence of density effects, and little difference to their size. For example, the finding that presentation order did not matter allows us to cautiously discount the possibility that the cue density effect is smaller than the outcome density effect as a result of the cue being typically presented before the outcome. After all, little changed when I reversed this default event presentation order. Moreover, the finding suggests that contingency judgments do *not* rely on events being encountered in their ‘logical’ order with the putative cause preceding the outcome.

Our experiments therefore clarify two key results that any theory of contingency learning must be able to explain. First, it clarifies that density effects in contingency learning are larger for outcome density manipulations than for cue density manipulations. Second, it clarifies that contingency judgments and the size of density effects are mostly unaffected by the presentation order of the contingent events (cue and outcome).

5.3.1 Theoretical Evaluation

Contingency account

As discussed in chapter 4, the contingency account cannot account for cue and outcome density effects. In addition to this there is an assumption of the contingency account, and other human contingency learning and association accounts, that the cue occurs first, and then the outcome occurs (De Houwer & Beckers, 2002). However, the order in which cue and outcome must occur is not essential to contingency calculations.

The Wasserman extension was able to predict the cue and outcome density findings. If there is no assumptions about the order of cue and outcome then the Wasserman extension is able to account for these findings. For the contingency calculations, the order in which cue and outcome occur is not strictly necessary. So long as both the cue-values and outcome-values are encoded in memory, the cue and outcome could occur in any order. Therefore, if we assume the account does not need a cue to precede the outcome, it can account for the findings that presentation order of cue and outcome does not affect cue or outcome density effects.

Associative account

In regard to event presentation order, the associative account assumes that the cue is followed by the outcome, as otherwise there would be no surprise element and hence no learning (Shanks, 2007). Therefore, switching the order of the cue and the outcome should disrupt learning. As such, the setup of associative models would suggest that they cannot naturally account for the finding that contingency judgments were largely unaffected by reversing event presentation order to place the outcome before the cue. At minimum, this suggests that for an associative account to explain our data, the CS and US must be arbitrarily defined – the first presented event is being associated with the second presented event regardless of the assigned labels of cause and outcome. This is because, if surprise for one event following another is the mechanism for learning, our data suggest that surprise at outcome-value for given cue-value is approximately equivalent to surprise at cue-value for a given outcome value.

There has been a great deal of research within this area, which has not particularly worried about the order in which cues, and outcome are shown. And the findings of this experiment do show it is not a particularly important issue to consider. However as one of the implicit assumptions of the Rescorla-Wagner associative perspective is that the cue precedes the outcomes, it cannot account for the findings of this experiment.

Probabilistic account

Similar to other models, the base assumption of probabilistic contrast models is that cues are followed by outcomes. Taking the probabilistic contrast model calculation as an example, the account suggests that power is driven by the idea that for any cause of an effect, there are always potential alternatives causes. Accordingly, knowledge of the base rate of the effect (frequency of outcome) will let one make a contrast between the candidate cause and alternative causes such as experimental context, and causal power is defined by contingency of an outcome contrast. As such this model neither explicitly states anything regarding presentation order.

5.3.2 Limitations and Future Directions

Overall, the results of Experiment 3 and 4 suggest that the effects of cue density and outcome density, are not effects of first seen or second seen events, as the size of the density effects remain relatively consistent. However, the size of the cue density effects was larger when outcome was seen first, compared to when cue was seen first. In the paradigms used in these experiments, the cue and outcome occur within milliseconds of each other. In real world settings, cue-outcomes pairings may be separated by much longer time gaps. Future research would benefit from establishing if the interchangeability of cue and outcome persists across greater time durations.

5.4 Conclusion

In summary, this research has illustrated robustness of the outcome density effect, and its greater size in comparison to cue density effects. In addition to this, event presentation order has negligible effect on judgements, for both cue and outcome density effects. Thus, perhaps cue density effects are inherently smaller than outcome density effects. Furthermore, contingency judgments do not seem to rely on events being encountered in their ‘logical’ order with the cause preceding the outcome.

Chapter 6: Scales and Contingencies

Prior literature has examined positive and negative contingencies; however, these have most often been conducted on active paradigms whereby participants can respond by providing the cue which may result in the outcome (Blanco et al., 2011; Shanks & Dickinson, 1991). In some studies, the outcome is programmed to occur on a certain number of times, due to participants being able to provide the cue however many times they may want to, this results in the experienced contingency of participants to vary greatly from the programmed contingency. In addition to this, much of the prior literature on negative contingency has examined negative contingency in relation to inhibition and control (Karazinov, 2018; Lee & Lovibond, 2021; Avellaneda et al., 2016; Williams, 1995). Within this literature, cues are paired to assess inhibition. Whilst this literature does provide information regarding negative contingencies to some extent, the pairing of two or more cues together means many of these studies cannot be compared to the current study, which is only looking at singular cues. A similar issue occurs within the positive contingency literature as well whereby multiple cues are paired and unpaired with the outcome throughout trials (López et al., 1998). Finally, many of the contingency studies, do not include both high and low outcome density conditions. Most often high outcome density conditions are only included, as it is well known that these produce biased judgements (Matute, 1996). As such comparisons of the outcome density effect across positive and negative contingencies are difficult.

This experiment investigated the outcome density effect across positive and negative contingencies in comparison to zero contingency conditions. Positive and negative contingencies in direct comparison to zero contingency are of interest, to see if people make consistent errors across tasks, or whether biases are specific to the zero-contingency design. In addition to this, it will be interesting to see if people can accurately judge the extent to which an association is positive or negative or if people overestimate the programmed contingency. So far, the existing literature has shown that people judge zero contingencies to be more positive than they truly are. As such it may be expected that participants will judge positive contingencies to be more positive than the programmed contingency. However, with negative contingencies it may be argued that participants will make “more extreme” judgements and consequently judge a more negative contingency than the programmed contingency, or perhaps participants will give a less extreme judgment with a shift towards less-negative judgments. Based on previous findings it would be expected that high density conditions would give higher judgements than low density conditions in negative contingencies (Msetfi et al., 2013a).

A potential issue in studying negative contingencies is the 0-100 scale typically used in contingency paradigms. With the 0-100 scale participants cannot indicate a negative association. Whilst with a zero and positive contingency it would not be expected for participants to give a negative judgement. However, for the zero contingency, participants tend to make positive judgements, and it is unclear if this is due to the scale biasing them towards a positive association as there is no space to make a negative judgment. And out of the possible answer choices of 101 options, only one indicates for a zero contingency or no association whilst 100 indicate a positive judgement. As such the scale itself typically used for paradigms such as these could be leading towards the general positive bias seen in outcome density effects. In addition to this, when considering a negative contingency using the 0-100 scale participants cannot provide a judgement which would reflect the negative contingency they are experiencing.

A study was conducted to test the influence of interactive effects of the cue and outcome on contingency judgements (Blanco et al., 2013). The study consisted of two experiments which used a zero-contingency paradigm and varied the frequency of both the cue and outcome, with one experiment using a 0 to 100 scale and the other using a -100 to 100 scale. Outcome and cue density effects were found across the studies. However, the stimuli were also varied between studies and the study did not manipulate scale type within the same experiment, as such it is not clear how the scale type affects judgements of contingencies.

Neunaber and Wasserman (1986) examined the effects of unidirectional and bidirectional rating scales on contingency judgements. Fifty-two participants were allocated to either the unidirectional or bidirectional scale condition. They took part in 13 blocks of trials, there were four positive (+.80, +.60, +.40, +.20), four negative (-.80, -.60, -.40, -.20) and five zero contingency blocks. An active contingency task was used, whereby participants could respond by pressing a button which may or may not result in a bulb lighting up. After each block, participants gave a contingency judgement. Judgements in the unidirectional scale were less sensitive to variations in contingences. Judgements in the bidirectional scale for the positive and negative contingencies were highly symmetrical. Judgements in the bidirectional scale were also not biased by the probability of the outcome. However, this was an active contingency paradigm in which participants could freely respond which means the experienced contingency by participants may not match the programmed contingency.

Hence this experiment has two aims. Firstly, to investigate how judgement scores differ between positive, zero and negative conditions, by including all the contingencies. To ensure equal occurrence of outcomes, such that one contingency is not more “positive” or “negative” than the other, symmetrical contingency calculation will be used, to ensure even distribution of events across contingencies. Consequently, I predict a difference in judgement of effectiveness between the high- and low-density conditions, with higher judgement scores given in the high-density condition compared to the low-density conditions, across all contingencies and scale conditions. In the contingency variation experiment, I predict there will be a difference in judgement across the three conditions, positive, zero and negative, with highest judgement for the positive contingency, and lowest judgement for the negative contingency.

Secondly the experiment investigated the difference in judgement of perceived association when using two different types of judgement scales in a passive contingency task. I predict that people will give higher judgement scores in the 0-100 scale condition compared to the -100 to 100 scale condition, due to the larger range of judgements that can be chosen with the -100 to 100 scale. Finally, I predict that the outcome density effect will appear across all conditions. The main goal of this chapter is to examine whether the outcome density effect is robust. This chapter also examined how methodological variations in the contingency learning paradigm affect the size of the outcome density effect.

6.1 Methods

6.1.1 Design

A mixed design was used, with judgement of effectiveness of medication in treating illness as the dependent variable. The following independent variables were varied in the experiment design: scale (between subjects: 0-100 (Old scale) vs -100 to 100 (new scale) by contingency (within / between subjects: zero vs positive vs negative contingency) vs outcome density (within subjects: high vs low density).

In the old scale condition, only the positive and zero contingencies were tested, as measuring a negative contingency using a 0-100 scale may not be appropriate. In the new scale condition, all three contingencies were tested.

All participants completed two zero-contingency judgements, and two non-zero contingency judgements (positive or negative), giving a total of 4 contingency judgements per participant. See Table 34 for contingency calculation for high- and low-density conditions across positive, negative and zero contingencies. For previous experiments, the probability of outcome given cue was 0.2 for the low outcome density conditions and 0.8 for the high-density conditions. However due to the inclusion of the positive and negative contingencies in this experiment to ensure symmetry in the contingency calculations the probability of outcome given cue was changed for the zero contingency. Thus 0.35 for the low outcome density condition and 0.65 for the high outcome density conditions were used. This then allowed for the average probability of outcome occurring across low density blocks across all contingencies to also be 0.35, as there were 14 trials which the outcome occurred out of a total of 40, 14/40=0.35. Similarly, across all high-density blocks in all contingencies the average probability of outcome occurring is 0.65 as there were a total 26 trials on which the outcome occurred in all contingencies, 26/40=0.65. To ensure a symmetrical design, the contingency calculations for positive and negative contingencies were flipped. So, the probability of an outcome given cue in the positive contingency, was used for the probability of outcome given no cue, and vice versa. This allows for an equal comparison of the two contingencies.

The task is based on Blanco et al. (2014), with the same medical treatment paradigm used in previous chapters. See Figure 13 for the participants allocation to different arms of

Table 34 Contingency calculations for outcome density

Table a *Low outcome density at Positive Contingency*

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C)=.5
Absent	4	16	P(O/~C)=.2
	ΔP=.5-.2=.3		

Table b *High outcome density at Positive Contingency*

	Outcome		
Cue	Present	Absent	
Present	16	4	P(O/C)=.8
Absent	10	10	P(O/~C)=.5
	ΔP=.8-.5=.3		

Table c *Low outcome density at Zero Contingency*

	Recovery		
Cue	Present	Absent	
Present	7	13	P(O/C)=.35
Absent	7	13	P(O/~C)=.35
	ΔP=.35-.35=0		

Table d *High outcome density at Zero Contingency*

	Outcome		
Cue	Present	Absent	
Present	13	7	P(O/C)=.65
Absent	13	7	P(O/~C)=.65
	ΔP=.65-.65=0		

Table e *Low outcome density at Negative Contingency*

	Outcome		
Cue	Present	Absent	
Present	4	16	P(O/C)=.2
Absent	10	10	P(O/~C)=.5
	ΔP=.2-.5=-.3		

Table f *High outcome density at Negative Contingency*

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C)=.5
Absent	16	4	P(O/~C)=.8
	ΔP=.5-.8=-.3		

the study, including branching to between-subjects conditions and the order in which participants encountered conditions within-subjects. Thus, participants either made judgments using the new scale or the old scale but made judgments for two different contingency conditions. Within a given contingency, participants encountered both a high-outcome-density and low-outcome-density condition.

6.1.2 Procedures

The procedure is similar to the previous studies. Participants signed up on the Prolific research participants platform and were prompted to read the information sheet and then complete the consent form. After completing demographic questions all participants took part in 4 contingency judgement tasks in one session. Identical timings were used for all

Figure 13 Participant allocation across conditions. Outcome density=Density

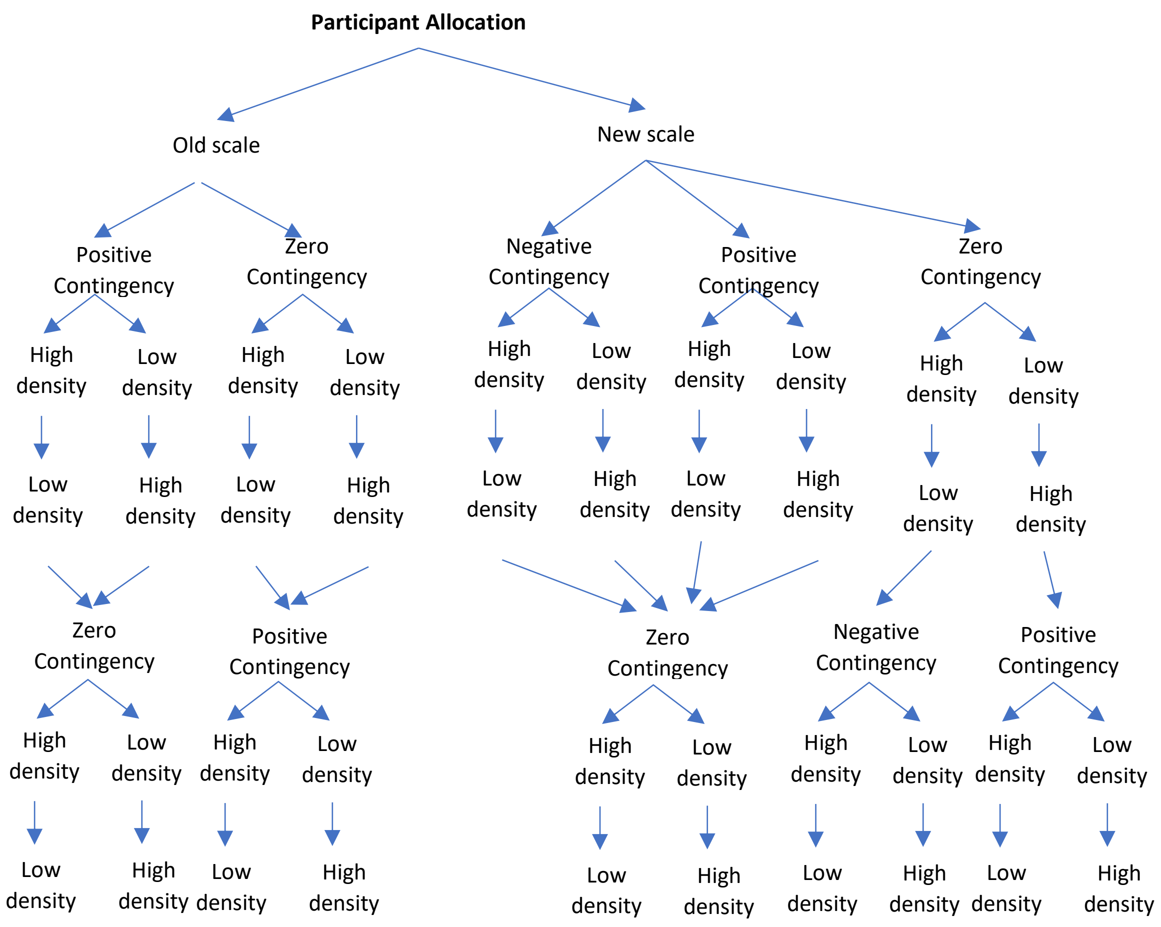

trials. Participants were randomly allocated to the conditions. The experiment took around 40 minutes to complete.

6.1.3 Participants

A total of 131 participants were recruited for the study, of whom 50 were female. The modal age group was 21 to 25, with ages ranging from groups of 18 to 65.

6.2 Results

6.2.1 Descriptive

See Table 35, Table 36 and Figure 14 for descriptive data. Table 35 shows descriptive statistics for each condition of the experiment. Table 36 shows the proportion of participants who perceived a positive, negative or zero contingency, for each condition. Figure 14 shows the spread of contingency judgements per condition.

6.2.2 Scale Type

Table 35 *Judgement scores descriptive statistics by Scale, Contingency (positive vs. zero vs. negative) and Density (high vs. low)*

	Old Scale				New Scale					
	Positive		Zero		Positive		Zero		Negative	
	High	Low	High	Low	High	Low	High	Low	High	Low
N	44	44	44	44	44	44	87	87	43	43
Mean	63.41	48.11	44.70	32.84	43.09	26.23	22.54	-5.14	-9.53	-20.84
Median	66	50	45	33	60	30	22	0	-8	-18
Standard Deviation	18.41	18.04	21.00	18.59	39.61	41.23	33.39	34.38	34.15	36.14
Range	85	75	91	92	168	153	155	175	140	145
Minimum	10	0	0	0	-83	-55	-75	-96	-90	-95
Maximum	95	75	91	92	85	98	80	79	50	50

As the analyses focusses on scale type, a 2 × 2 × 2 mixed ANOVA was used, with judgement of effectiveness of medication in treating illness as the dependent variable. Negative contingency was not included in this analysis as it was not measured on this scale type. Three factors were varied in the experimental design, which were the factors for the

Figure 14 Judgement scores including median and interquartile ranges

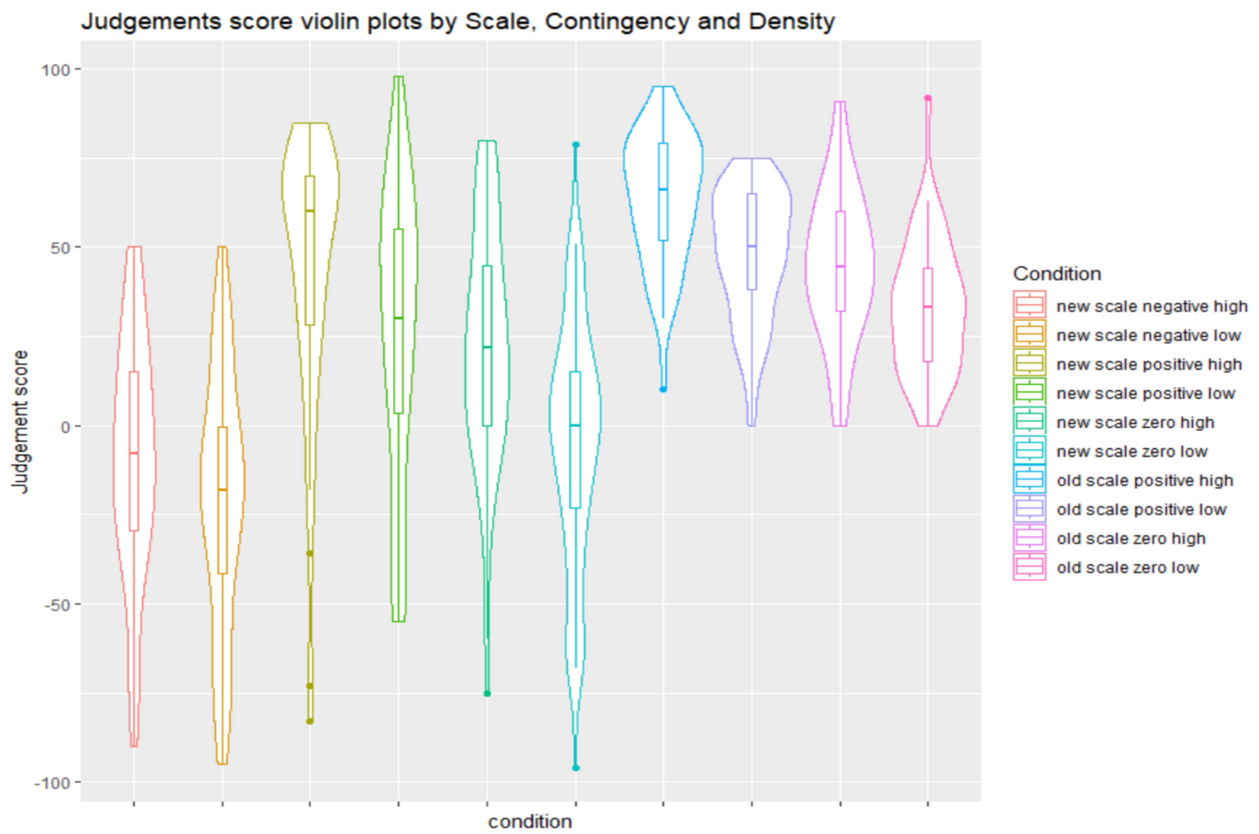

Table 36 Proportion of participants who perceived a positive, zero, or negative contingency by Scale, Contingency and Density

	Old Scale				New Scale					
Proportion of participants who perceived a…	Positive		Zero		Positive		Zero		Negative	
	High	Low	High	Low	High	Low	High	Low	High	Low
Positive contingency (%)	100.0	100.0	95.5	97.7	86.4	77.3	69.0	43.7	37.2	23.3
Zero contingency (%)	0.0	0.0	4.5	2.3	0.0	0.0	14.9	10.3	2.3	2.3
Negative contingency (%)	-	-	-	-	13.6	22.7	16.1	46.0	60.5	74.4

ANOVA: contingency (within subject: positive vs zero contingency) by density (within subjects: high vs low density) and scale type (between subjects: old scale vs new scale).

There was a large main effect of contingency F(1, 86)=44.60, p < .001, η_p^2=.341. Participants judged the medication to be less effective in the zero-contingency condition *(M*=23.51, *SD*=23.36) than the positive contingency condition *(M*=45.21, *SD*=23.26). There was also a large main effect of density F(1, 86)=31.17, p < .001, η_p^2=.266. Participants judged the medication to be less effective in the low-density condition *(M*=24.84, *SD*=22.53) than the high-density condition *(M*=43.89, *SD*=25.05). There was a large main effect of scale type F(1, 86)=47.11, p < .001, η_p^2=.354. Participants judged the medication to be less effective with the new scale *(M*=21.46, *SD*=24.94) than with the old scale *(M*=47.27, *SD*=24.94). See Table 37 (though ignore the rightmost column for this analysis), and Figure 15.

There was no interaction effect of contingency by density F(1, 86)=1.48, *p*=.227, η_p^2=.017, no interaction effect contingency by scale type F(1, 86)=2.10, *p*=.151, η_p^2=.024, no interaction effect of density by scale type F(1, 86)=2.57, *p*=.113, η_p^2=.0.26. There was no three-way interaction of contingency by density by scale type was F(1, 86)=3.69, *p*=.058, η_p^2=.041.

Table 37 Mean (SD) response by Scale type, and manipulation of contingency

	Old Scale		New Scale		
	Positive	Zero	Positive	Zero	Negative
High density response	63.41 *(18.41)*	44.70 *(21.00)*	43.09 *(39.61)*	22.54 *(33.39)*	-9.53 *(34.15)*
Low density response	48.11 *(18.04)*	32.84 *(18.59)*	26.23 *(41.23)*	-5.14 *(34.38)*	-20.84 *(36.14)*
95% CI for mean difference	8.56, 22.04	4.88, 18.85	0.18, 33.55	18.85, 36.51	0.36, 22.24
Outcome density effect size *d*	0.84	0.60	0.42	0.82	0.32

6.2.3 Contingency

A multi-level model instead of an ANOVA was used to analyse all conditions with the new scale, due to the uneven design where all participants completed two zero contingency conditions (high and low outcome density). However, half the participants completed two positive contingency conditions (high and low outcome density) whilst the other half completed two negative contingency conditions (high and low outcome density). As such a multi-level model, was used with judgement of effectiveness of medication in treating illness as the dependent variable. Two factors were varied in the experimental design: contingency (mixed: positive vs zero vs negative contingency) by density (within subjects: high vs low density).

Figure 15 Mean effectiveness judgments by Scale, for positive and zero contingencies by high and low densities, with 95% confidence intervals

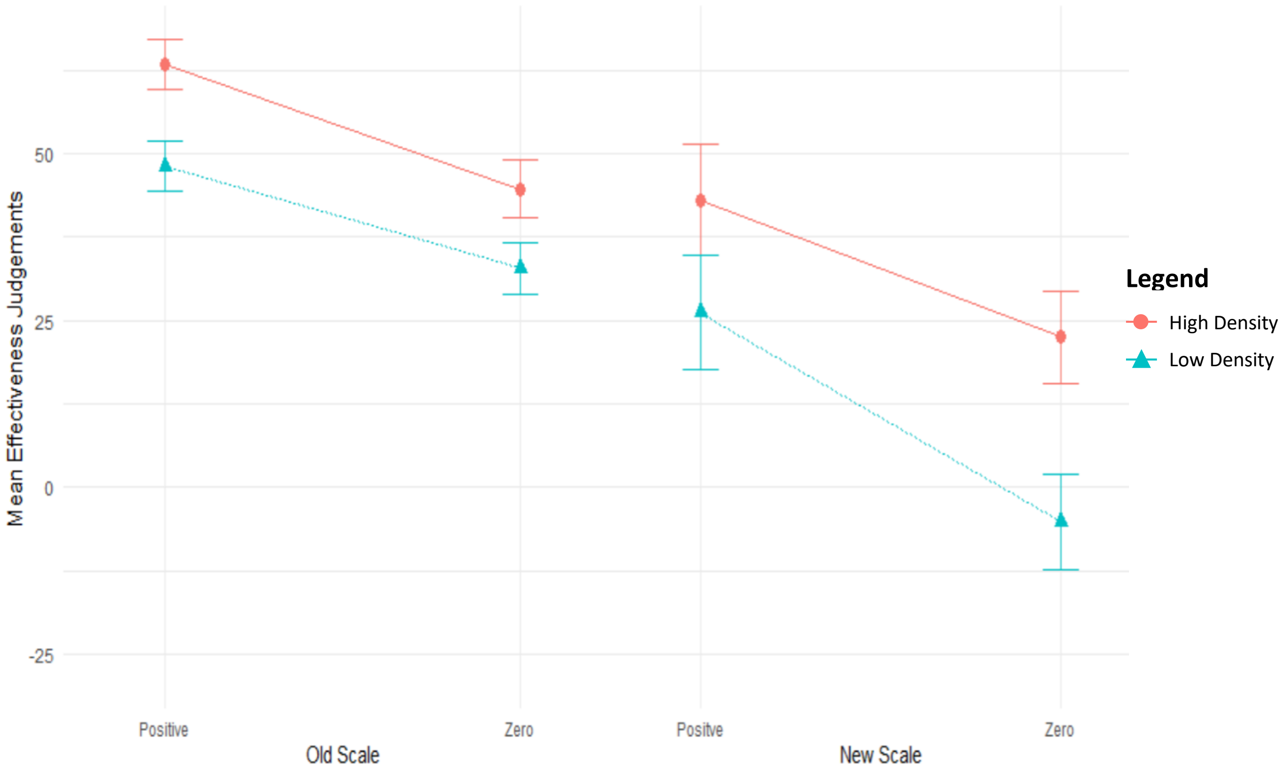

There was a large main effect of contingency F(2, 348)=42.61, p < .001. Participants gave higher judgement for medication effectiveness in the positive contingency condition (*M*=34.66, *SD*=35.64), than the zero-contingency condition (*M*=8.70, *SD*=25.38) or the negative contingency condition (*M*=-15.19, *SD*=36.05). See Table 38 for pairwise comparison, which shows significant differences between all pairs of contingencies. There was a main effect of density F(1, 348)=21.36, p < .001. Participants judged the medication to be less effective in the low-density condition (*M*=0.84, *SD*=26.72) than the high-density condition (*M*=18.70, *SD*=26.72), see Table 37 for density effect sizes.

No interaction effect of contingency by density was found F(2, 348)=1.71, *p*=.183. See Figure 16.

6.3 Discussion

This study used a learning paradigm to investigate whether scale type and contingency type influences judgement. I hypothesised that people would give higher judgement scores in the 0-100 scale condition compared to the -100 to 100 scale condition.

Table 38 Pairwise comparison of positive, zero and negative conditions, for new scale

	Positive - Zero	Positive - Negative	Zero - Negative
Mean difference (*SD*)	25.96* *(43.72)*	49.85* (*50.68)*	23.89* (*44.06)*
95% confidence interval for the mean difference	16.79, 35.13	39.22, 60.47	14.65, 33.17
Effect size *d*	0.84	1.39	0.77

*The mean difference is significant at the 0.05 level

In addition to this, I also predicted that there will be a difference in judgement across the three conditions, positive, zero and negative, with highest judgement for the positive contingency, and lowest judgement for the negative contingency. The study found support for all of these predictions, moreover there were no interaction effects across these. All significant effects were large. None of the non-significant effects were larger than small-medium in size.

6.3.1 Scale type

As expected, a scale type effect was also found. Overall participants gave higher contingency judgements on the 0-100 scale compared to the -100 to 100 scale, across both positive and zero contingency conditions. When considering the proportion of participants who conceived a positive or zero contingency, for the positive and zero contingency by

Figure 16 Mean effectiveness judgments by contingency for high and low densities, with 95% confidence intervals

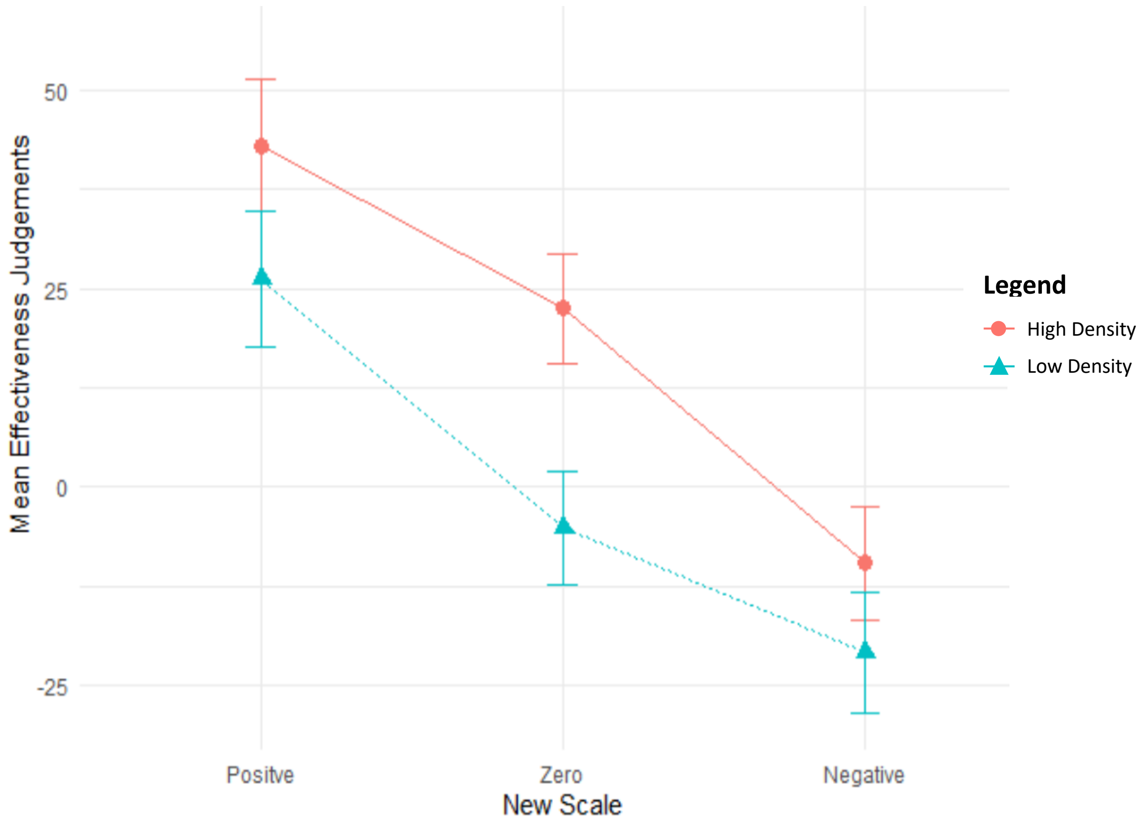

scales (Table 36), for the 0 to 100 scale in the zero contingency condition only a very small proportion of participants perceived a zero contingency in low (2.3%) and high (4.5%) outcome density conditions. For the -100 to 100 scale in zero contingency, the proportion of participants who perceived a zero contingency for the low (14.9%) and high (10.3%) outcome density conditions increased slightly, there was a larger proportion of participants who perceived a negative contingency for the low (16.1%) and high (46%) outcome density conditions. Overall, this suggest that for the zero-contingency condition, more people perceived the zero contingency condition accurately in the -100 to 100 scale. However, an even larger proportion perceived a negative contingency.

For the 0 to 100 scale all participants perceive a positive contingency for both the low- and high-density conditions, in the positive contingency condition. However, when considering the -100 to 100 scale, a small proportion of participations perceive a negative contingency in the low (22.7%) and high (13.6%) density conditions.

To my knowledge only Neunaber and Wasserman (1986) have looked at the effect of scale length on contingency judgements within the same experiment. This experiment does not replicate the findings of Neunaber and Wasserman (1986). They found judgements in the bidirectional scale to be symmetrical for the positive and negative contingencies which was not the case for this experiment. In addition to this, they found judgements in the bidirectional scale were not biased by the probability of the outcome. Again, an effect of outcome density or probability was found in the current experiment. The Neunaber and Wasserman (1986) was an active contingency paradigm in which participants could freely respond which means the experienced contingency by participants may not match the programmed contingency. This may potentially explain the differences in the results between the two studies.

The findings here suggest that that the 0 to 100 scale may force participants to provide more positive contingency judgements due to a lack of other options, as the scale provides 100 options for a positive judgement, and only one option for non-positive judgements. In contrast the -100 to 100 scale offers 101 options for non-positive judgements. By using the -100 to 100 scale, fewer participants are perceiving a positive contingency in the positive contingency, however more participants are perceiving a zero contingency in the zero-contingency condition. Looking at the spread of judgements made, across positive and zero contingencies for both scales (Table 35, and Figure 14), for the -100 to 100 the judgements are far more spread out across the entire scale, in particular for the zero contingency conditions whereby the judgements are more evenly spread across the scale length. As such, for previous contingency judgements experiments, whereby a large perceived positive contingency has been found for zero contingency (López et al., 1998; Moreno-Fernández et al., 2017; J. C. Perales et al., 2005), that is a large outcome density effect, may be due to the unidirectional scales being used. By offering a unidirectional scale which runs positively, participants may implicitly believe their judgement has to be positive, regardless of what they may perceive, and as such provide a positive judgement.

However, when considering the high conditions for the zero contingency with the new scale, a very large proportion of participants still detect a positive contingency, indicating a general tendency to make positive associations or contingency judgements when there is an increased probability of the outcome. As such the outcome density effect is still present in the -100 to 100 scale, however the effect may just be exaggerated with the use of the 0 to 100 scale.

6.3.2 Contingency type

As expected, a contingency effect was found. Across the three contingencies the outcome density effect was strongest in the zero contingency, and weakest in the negative contingency. With the new scale, a fair proportion of participants in the positive contingency for both high and low outcome density (Table 36), perceived a negative contingency as they gave a negative contingency judgment. An even larger proportion of participants perceived a positive continency, in the negative contingency condition. Looking at the zero contingency condition nearly 70% of the participants perceived a positive continency in the high-density outcome condition, however for the low-density outcome condition there was more of a high even split of participants who perceived a positive or negative contingency. This is interesting, in particular as the proportion of participants who perceived a zero contingency in comparison is quite smaller.

When considering the proportion of participants who detected the programmed direction of contingency correctly, when they were in the low outcome density conditions, it was a similar proportion across the positive (77.3%) and negative (74.4%) contingencies. However, the proportion of participants who detected the correct contingency in the high outcome density conditions for the positive (86.4%) and negative (60.5%) contingencies varied more. The contingency calculations were symmetrical for the positive and negative contingencies, however the effects seen are not symmetrical. Whilst the effect sizes for the mean difference of positive contingency – zero contingency (0.84) and negative contingency

– zero contingency (0.77) are both large and relatively close, the mean judgements are not equidistant from 0 as may be expected based on the symmetrical calculations. Rather all the mean judgements seem to be shifted more positively.

This suggests that cue-outcome cells are more important than no cue-no outcome cells for peoples' judgements of positive contingency. To understand this point, refer to Table 34. When comparing the contingency calculation for high outcome density at positive contingency (Table 34b) and low outcome density at negative contingency (Table 34e), the absence of the cue occurring with or without the outcome, is the same across both. The only difference across these conditions is presence of the cue with or without the outcome, which is reversed for these two conditions. As such judgements of positive contingency are affected by the increased presence of cue-outcome pairings, cell A trials, whilst judgements of negative contingency are affected by the increased presence of cue-no outcome pairings. This supports Wasserman et al. (1990) find that that cell A counts increase positive judgements whilst cell B counts decrease positive contingency judgements.

For negative contingency judgements, the presence of no cue-outcome pairings are more important. When comparing the calculations for low outcome density at positive contingency (Table 34a) and high outcome density at negative contingency (Table 34f), the cue occurs with the same frequency with the outcome and with no outcome across the two conditions. The only difference is the occurrence of the absence of cue with or without the outcome, these two figures are reversed across these conditions. As such the increase in no-cue-outcome in the high outcome density negative contingency seems to be leading to the increased negative contingency judgements. As such judgements of positive contingency are affected by the increased presence of no cue-no outcome pairings, cell D trials. Whilst judgements of negative contingency are affected by the increased presence of no cue-outcome pairings, cell C trials. This supports Wasserman et al. (1990) findings that cell D counts increase positive judgements whilst cell C counts decrease positive contingency judgements. The data show that this logic extends in a coherent manner to negative contingency judgments.

In sum, participants seem to be making associations that are, or average, more positive than the contingency would indicate but which are further shifted in a positive direction in the presence of high outcome density. Importantly, by including both positive and negative contingencies and using a new scale to accommodate both directions of contingency, I was able to demonstrate that increasing the outcome density makes judgements more positive rather than making them stronger. That is by increasing the probability of the outcome occurring, the judgements made are more positive, rather than participants making stronger judgments. As such it shows that researchers should not be saying ‘associations are over-estimated when outcome density is high’ or ‘associations are judged more strongly when outcome density is high’. But rather, ‘when outcome density is high, associations are judged to be more positive (less negative) than when outcomes density is low’.

As expected, the outcome density effect was robust and appeared in all conditions across the experiment. Participants consistently have higher contingency judgements in the high outcome density condition compared to the low outcome density condition, replicating the findings from the previous experiments in this thesis. The outcome density effect was the smallest for the negative contingency, and largest for the zero contingency conditions overall, the largest outcome density effect was for the positive continency in the old scale condition. However, even though this was a well-powered investigation, the study did not identify the significant interaction that would lead us to believe that density effects were different for positive and negative contingencies.

Within the zero contingency the outcome density effect was largest with the new scale. However, for the positive contingency the outcome density effect was largest with the old scale compared to the new scale, see Figure 17 for comparisons. This is interesting as the scale type may affect the outcome density effect between contingencies, however there was no interaction effect of scale type by contingencies found. On average the difference between mean judgement in the old scale and the new scale is 21, for high and low positive

Figure 17 Mean effectiveness judgments by contingency for high and low densities, with 95% confidence intervals, across all conditions for both old and new scale

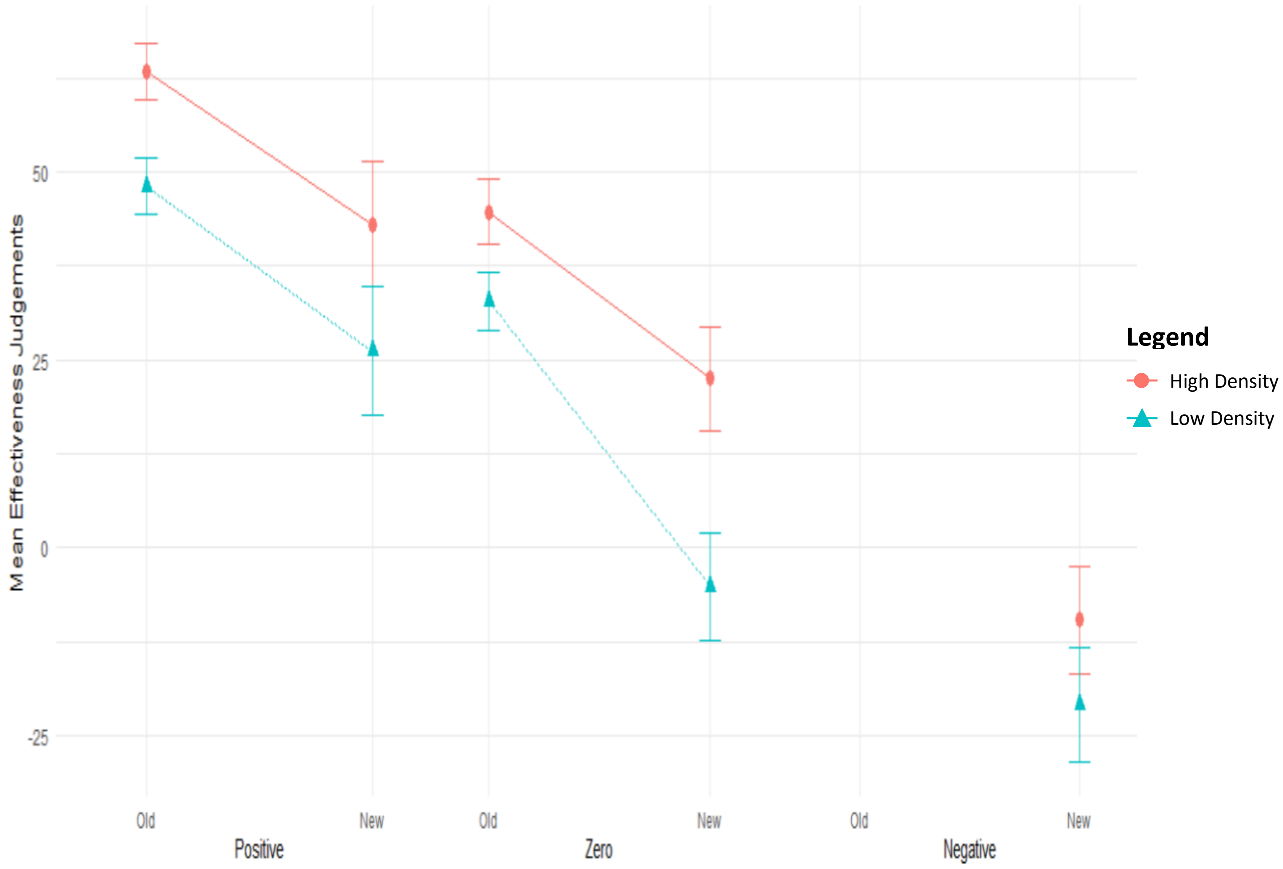

contingency, and high zero contingency, however for the low zero contingency the difference is 38, suggesting that this condition is more affected by the scale length.

6.3.3 Previous findings

This experiment replicates the findings of Maldonado et al. (1999, 2006) that people can detect positive and negative contingencies. Although the proportion of participants who correctly perceived the contingency condition, differs between the positive and negative contingency compared to the current experiment. The higher judgements for the positive contingency than the zero contingency replicates findings of previous studies (Shanks & Lopez, 1996; Cavus & Msetfi, 2016; Crump et al., 2007).

Crump et al. (2007) found both a main effect of contingency and outcome density as the current study. However, they also found an interaction of contingency by outcome density, such that the ratings between the low and high outcome density conditions were largest for the zero contingency than the positive contingency, which were not replicated here. Interestingly Crump et al. (2007) found much lower means for the low (-29.9) and high (5.6) zero contingency condition, and the low (4.7) and high (18.8) positive contingency conditions compared to the current study. The study used outcome probability of 0.2 for the low and 0.8 for high zero contingency condition, 0.33 for the low and 0.67 for the high positive contingency condition. In addition to this whilst the study used a contingency paradigm, participants were viewing a stream of information about abstract cues and outcome and making both contingency and frequency estimate judgements. As such, methodologically there are many differences between the current experiment, and the Crump et al. (2007) experiment, which may provide some explanation for the interaction effect and much lower mean judgements made for the different condition, however similar main effects have been found.

Perales et al. (2005) looked at cue density effects, in a 2 (between subjects: zero vs positive) by (within subjects: high vs low density) and collected predictive responses after each trial, and contingency judgements at the end the task from 44 participants. The study found main effects of contingency; participants gave higher judgements in the positive contingency compared to the zero contingency. The current study has replicated these findings in a well powered experiment, with almost double the number of participants per condition, compared to the Perales et al. (2005) study.

Vallée-Tourangeau et al. (2005) measured causal ratings in a 3 by (contingency: positive, zero and negative) by 3 (base rate of outcome: P(O)=.25, medium, P(O)=.5, or high, P(O)=.75) in a fully within subjects design. Thirty-four participants assessed the extent to which pressing the spacebar on a keyboard resulted in the appearance of a geometric shape and gave judgements on -100 to 100 scale. They found a main effect of contingency, such that participants gave more positive judgements in the positive contingency compared to negative and zero contingencies, and a more negative rating in the negative contingency compared to the zero contingency condition. Similarly, they also found a main effect of outcome density, such that increased probability of outcome led to higher judgements, but there was no interaction of contingency by outcome density found. However the mean judgements in the Vallée-Tourangeau et al. (2005) experiment for the zero and negative contingencies were far more negative then what was found in the current experiment. This may in part be due to the use of an active paradigm, whereby the programmed contingency, may not be the experienced contingency. In addition to this, there were no separation of the different conditions, and as such the stimuli did not vary between conditions, which may have affected the causal ratings.

6.3.4 Theoretical Evaluation

Contingency account

The contingency account can explain the effect of contingency in experiment 5, as participants are expected to make more positive judgements the more positive the overall ΔP is. In addition to this, when the contingency calculations of the experiment are examined, further support for Wasserman et al. (1990) extension can be found. To understand this point, refer to Table 39 (reproduced from Chapter 6). When comparing the contingency calculation for high outcome density at positive contingency (Table 39Table 34b) and low outcome density at negative contingency (Table 39e), the absence of the cue occurring with or without the outcome (cells C and D), is the same across both. The only difference across these conditions is the presence of the cue with or without the outcome (cells A and B), which is reversed for these two conditions. As such, judgements of positive contingency are affected by the increased frequency of cue-outcome pairings, cell A trials. Whilst judgements of negative contingency are affected by the increased presence of cue-no outcome pairings, cell B trials. This supports Wasserman et al. (1990) findings that cell A counts increase positive contingency judgements whilst cell B counts decrease positive contingency judgements.

When comparing the calculations for low outcome density at positive contingency (Table 39a) and high outcome density at negative contingency (Table 39f), the cue occurs with the same frequency with the outcome (cell A) and with no outcome (cell B) across the two conditions. The only difference is the occurrence of the absence of cue with or without the outcome (cells C and D), these two figures are reversed across these conditions. As such the increase in no-cue-outcome, cell C, in the high outcome density negative contingency seems to be leading to the increased negative contingency judgements. In tandem, judgements of positive contingency are affected by the increased presence of no cue-no outcome pairings, cell D trials. Whilst judgements of negative contingency are affected by

Table 39 Contingency calculations for outcome density

Table a *Low outcome density at Positive Contingency*

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C) = .5
Absent	4	16	P(O/~C) = .2

ΔP = .5-.2=.3

Table b *High outcome density at Positive Contingency*

	Outcome		
Cue	Present	Absent	
Present	16	4	P(O/C) = .8
Absent	10	10	P(O/~C) = .5

ΔP = .8-.5=.3

Table c *Low outcome density at Zero Contingency*

	Recovery		
Cue	Present	Absent	
Present	7	13	P(O/C) = .35
Absent	7	13	P(O/~C) = .35

ΔP = .35-.35=0

Table d *High outcome density at Zero Contingency*

	Outcome		
Cue	Present	Absent	
Present	13	7	P(O/C) = .65
Absent	13	7	P(O/~C) = .65

ΔP = .65-.65=0

Table e *Low outcome density at Negative Contingency*

	Outcome		
Cue	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	10	10	P(O/~C) = .5

ΔP = .2-.5= -.3

Table f *High outcome density at Negative Contingency*

	Outcome		
Cue	Present	Absent	
Present	10	10	P(O/C) = .5
Absent	16	4	P(O/~C) = .8

ΔP = .5-.8= -.3

the increased presence of no cue-outcome pairings, cell C trials. This supports the findings of Wasserman et al. (1990), that cell D trials increase positive judgements whilst cell C counts decrease positive contingency judgements. The data show that this logic extends in a coherent manner to negative contingency judgments.

An effect of scale length was found, when the scale ran from -100 to 100, participants gave fewer positive judgements, compared to when the scale ran from 0 to 100. The contingency judgements of the -100 to 100 scale more closely match the predictions of the contingency account, as the contingency judgements were closer to the programmed contingency. However, this effect of scale is not necessarily explainable by the contingency account, as there are no assumptions regarding scale length made with the contingency account, to my knowledge. Similarly, the Wasserman account cannot account for these findings either. However, since the scale effect suggests that by using a bidirectional scale, judgements are shifted closer to the programmed contingency. This may not necessarily be an issue with the contingency account but rather a methodological issue whereby participants cannot give the judgement they may want to give when using a unidirectional scale.

Overall, the contingency account is able to account for the effects of contingency found in this experiment. The findings also support the Wasserman account. However, neither account is able to explain the shift in judgements when a bidirectional scale is used instead of a unidirectional scale.

Associative account

The associative perspective can account for the judgment of contingency between a cue and outcome. As the associative strength of a cue, has been found to correlate positively with ΔP (Baker et al., 2000). When the correlation of cue with outcome is positive, its associative strength approaches λ. When the cue is unrelated to the outcomes, the correlation is zero, as such the associative strength is zero. When the cue is negatively correlated with the outcomes, the associative strength of cue becomes negative. As on no-outcome trials the λ will be zero, hence $\lambda\text{-}V^{n-1}_{Total}$ will become negative, giving rise to negative associative strength for the cue. This reflects the findings that participants give negative judgements in a negative contingency, and positive judgements in the positive contingency.

However, the effect of scale is not necessarily explainable by the associative account, as there are no assumptions regarding scale length. However, as mentioned previously, since the scale effect suggests that by using a bidirectional scale judgement are shifted closer to the programmed contingency. This may not necessarily be an issue with the associative perspective but rather a methodological issue.

Probabilistic account

The model can account for the contingency effect found, as it does predict that more extreme levels of outcome density results in higher judgments of control in positive and negative contingency conditions. The model predicts that some participants will perceive zero contingencies as non-zero, i.e. as positive or negative contingencies. Hence, an effect of scale would be expected for zero contingency conditions because participants cannot easily represent theory beliefs about a negative contingency.

6.3.5 Limitations And Future Directions

Across all the contingency conditions, the probability that the cue occurs across any block of trials is 0.5, as out of the cue occurs on 20 trials out of the 40 trials of a block. However, the probability that the cue occurs with and without the outcome is only the same in the zero-contingency condition, 0.5, and in fact is the same across both the high and low outcome density conditions. Across the positive and negative contingency conditions, the probability of the cue occurring with and without the outcome differ see Table 40 for these calculations. When manipulating the P(O) such that a contingency is positive or negative, the P(C) is also changed such that the P(C) occurring with the outcome and without the outcome differs. This difference in the probability of the cue occurring may influence judgements of participants.

Blanco et al. (2013) found that when cue density was low participants gave lower judgements in the high outcome condition, compared to when cue density was high. Participants' judgement did not vary in the low outcome condition, whether cue density was high or low, and the low-density conditions had a probability of 0.2, with the high conditions having a probability of 0.8. such the cue probability in the positive and negative contingencies in this experiment are on the low side, which may suggest that cue density may not have influenced the contingency judgments of participants. However further experimental research may be needed here to clarify the effect of low cue density on contingency judgements, in positive and negative contingencies.

The outcome density effect still appears on a bidirectional scale. However, the overall judgements are much lower and closer to programmed contingencies. The large density effects seen in the literature may be due to the use of unidirectional scale. As such, I would recommend the use of bidirectional scales, so participants are not forced into implicitly thinking the contingency must be positive. An extension to this research may see if the outcome density effect changes if the proportion of the scale changes. For example, does the outcome density effect size stay stable if a -10 to 10 scale is used instead of a -100

Table 40 Contingency calculations for cue density

Table a

Cue density at low outcome density for Positive Contingency

	Outcome	
Cue	Present	Absent
Present	10	10
Absent	4	16
	P(C/O)=.71	P(C/~O)=.38

Table b

Cue density at high at high outcome density for Positive Contingency

	Outcome	
Cue	Present	Absent
Present	16	4
Absent	10	10
	P(C/O)=.62	P(C/~O)=.29

Table c

Cue density at low outcome density for Zero Contingency

	Outcome	
Cue	Present	Absent
Present	7	13
Absent	7	13
	P(C/O)=.5	P(C/~O)=.5

Table d

Cue density at high at high outcome density for Zero Contingency

	Outcome	
Cue	Present	Absent
Present	13	7
Absent	13	7
	P(C/O)=.5	P(C/~O)=.5

Table e

Cue density at low outcome density for Negative Contingency

	Outcome	
Cue	Present	Absent
Present	4	16
Absent	10	10
	P(C/O)=.29	P(C/~O)=.62

Table f

Cue density at high at high outcome density for Negative Contingency

	Outcome	
Cue	Present	Absent
Present	10	10
Absent	16	4
	P(C/O)=.38	P(C/~O)=.71

to 100 scale, rather than having 100 options for positive or negative judgements, there will only be 10 e.g., in anchoring research, the implied precision of the anchor affects the size if anchoring effects (Switzer & Sniezek, 1991). Therefore, plausibly, people's willingness to judge something as zero may vary across these scales. As such this could further explore methodological variations under which the outcome density effect reduces and increases.

The experiment which manipulated contingencies, did not examine the outcome density effect in a negative contingency using a 0 to 100 scale. It was reasoned that it would not make logical sense to include such a condition as participants would not be able to give negative judgements. However, to fully examine the effect of scale, it would be interesting to see if participants would give lower judgements for negative contingency compared to zero and positive, even with the unidirectional scale.

6.4 Conclusion

To summarise, the length of scale can influence the contingency judgements given, such that a unidirectional positive scale, may influence participants to have more positive bias in the judgements they make, in comparison to a bidirectional scale. As such the use of a unidirectional scale may have led to the larger outcome density effects seen in previous experiments. Both the contingency and associative accounts are able to account for the contingency findings, however only the probabilistic theory is able to account for the scale related findings.

Contingency judgements are much more positive in the positive contingency compared to the zero and negative contingencies, however the outcome density effect is largest in the zero contingency conditions. In general, when outcome density is high, people make more positive association, even in the negative contingency condition.

Chapter 7: Scenarios and Prior Beliefs Pilot

I am interested in the real-world application of the contingency paradigm when assessing false belief formation. For this, a key question is *how do pre-existing beliefs influence judgements?* Two lines of previous enquiry can help to answer this question which is examined in two sections reported in this chapter and the one that follows. The first line of enquiry focusses on theoretical explanations of the formation of causal beliefs. All these accounts rely on some level of belief updating in which the prior belief held by individuals are updated based on the various sources of information or experience. Specifically, some of these explanations emphasise the importance of previous knowledge of the mechanisms linking a cause and effect. The second line of enquiry is primarily empirical and tests the impact of prior beliefs that are manipulated or formed in the lab. Both lines of enquiry are reviewed below; after which I report my investigation of how pre-existing beliefs acquired outside the lab affect contingency learning.

7.1 Theoretical accounts of belief updating in the formation of causal beliefs

Being able to examine how events co-occur underlies much of fundamental behaviours such as judging causation (Cheng, 1997), categorisation (Smith & Medin, 1981), and learning (Bandlow, 1976). As such a great deal of research has been conducted within this area. One of the most robust findings on the role of covariation of events, is the impact of prior belief and expectation on judgements (Alloy & Tabachnik, 1984; Chapman & Chapman, 1969; Peterson, 1980).

The theoretical explanations for causal association are often categorised between mechanism-based accounts and covariation based accounts (Ahn & Kalish, 2000; Michotte, 1963). The mechanisms-based accounts suggest that to be able to interpret causation previous knowledge of the mechanisms linking a cause and effect is needed. In contrast, covariation-based accounts suggest that knowledge of causal links occur due to experience of the covariation between a cause and effect. Regardless of which of these accounts may precede the other (Cheng, 1993, 1997; Perales & Catena, 2006), understanding causal associations in everyday life often relies on some level of belief updating in which the prior belief held by individuals are updated based on the various sources of information and experience (Fugelsang & Thompson, 2001; White, 1995).

Perales et al. (2007) conducted an experiment investigating two different categories of causal association accounts. A 3 (information type between subjects: covariation, mechanism, combined) by 2 (contingency between subjects: high vs low) design was used. In the covariation condition participants were given information of an empirical study which showed that 6 out of 10 subjects administered a drug suffered from a stomach ache, and 10 subjects who were not administered the drug did not suffer from stomach aches. Participants in the mechanism condition were told a certain drug generates peptides in the stomach, and that peptides are a known cause of stomach aches. In the combined conditions participants were provided both types of information. One hundred and forty-four participants made judgements of prior strength of the relationship between cause and effect, and prior confidence judgements that a causal link actually exists. Afterwards participants were presented with new covariation data on 20 new patients. In the high contingency condition participants were told 8 out 10 patients took the drug had a stomach-ache, and the other 10 patients who did not take the drug did not have a stomach-ache. In the low contingency group 2 out of 10 participants who took the drug had a stomach-ache, and the other 10 who did not take the drug do not have a stomach-ache. Participants then gave judgements of on how reliable the result of the new evidence is, along with judgments of how effective the potential drug is.

Main effects of both contingency and information type were found on the reliability of the evidence. There was no interaction effect of information type with contingency. Participants in the high contingency condition gave higher judgements of reliability than those in the low contingency group. Participants in the covariation group gave higher judgements of the evidence than both the combined and mechanism group. The findings of the study suggests that the perceived reliability of new covariation information is determined by which source of information was presented to participants, independent of the type of prior information previously held. In addition to this, prior beliefs were more likely to be changed by reliable information than unreliable information. Overall, the baselines differences in prior beliefs tend to show up in final integrative judgements of contingency, however it did not modulate the main effect of new information contingency. Overall, this suggests that prior beliefs do not interfere in the updating of belief. However, this study only focused on prior beliefs which were formed during the experimental settings, and not beliefs participants already held.

Evans et al. (2005) conducted an experiment in which they examined the role of prior beliefs on learning when the prior belief conformed with the learning. First a pilot study was conducted to determine stereotypical beliefs about a range of occupations. The results of the pilot were used to determine personality attributes which individuals believe positively and negatively predict performance in the occupation (belief link). In addition to this, the pilot was also used to find attributes which were neutral in predicting performance (belief neutral). Across two experiments participants viewed a series of screens which presented the individual occupations of people, along with selected personality attributes of everyone. After each case participants provided judgements of how well the individual would perform in that occupation and were provided feedback which would help improve their judgement. In experiment 1 the attributes presented which predicted performance were neutral, whilst the irrelevant attributes which did not improve participants' judgement were the belief-linked attributes. In experiment two, these were flipped, such that the belief-neutral attributes were irrelevant, whilst the belief-linked attributes did predict performance in the occupation. Experiment 1 found no influence of prior belief on judgements, however, experiment 2 found that people learnt to use the belief-linked attributes better, particularly when the attribute was a positive predictor. These findings suggest that an effect of belief polarity, such that holding positive beliefs results in learning positive cue attributes more easily. However, this experiment did not use the actual prior beliefs of participants, but rather the general prior beliefs, or stereotypes the population hold. As such there is a misspecification of the prior belief measure, as the individual participants may hold different prior beliefs, or prior beliefs that may not be as strong.

Whilst prior belief is often based on own experience, it can also be based on observations of others, for instance a younger sibling learning how to make decisions from their older siblings. Biele et al. (2009) investigated the role of good and bad advice prior to choice task. Those who received good advice performed better than those who received bad advice. This suggests that even receiving advice once, can influence judgements significantly due to the formation of prior beliefs.

7.2 Role of belief in active contingency paradigms

Much of the literature discussed thus far has examined the influence of prior belief across a number of different non-contingency tasks. Czupryna et al. (2021) examined the role of prior beliefs on the illusion of control effect using a modified contingency paradigm. One hundred and forty-five participants observed stock prices of different companies using simulated charts. The prices could increase or decrease in increments. Participants had to make the stock price reach the highest value by the end of the block, which had 50 increment changes. Participants could influence the stock prices by placing the cursor in a specific area on the screen, the change could be minor or major, however it was not guaranteed that a change would take place. There was a total of 10 blocks, with the base rate probability of outcome varying between 0.33, 0.5, 0.67. Each base rate appeared in 3 blocks, and the tenth block duplicated the base rate of the first blocks. At the end of each block their perceived base rate probability estimates were collected. The illusion of control was measured via two ways; as the difference between a participant's perception of their impact on the process of generating results (perceived control) and their objective influence on the results (real control). Prior belief was analysed as the carry over effect of the previous block. The control judgement participants gave for block 1, was considered the prior belief for block 2, this logic was applied to all blocks. The results suggest that the illusion of control may have two sources of biases, prior beliefs and behavioural factors. With prior beliefs there was negative correlation with real control, the higher the control, the more negative the bias in the control level assessment. Behavioural factors were positively correlated with real control, the higher the real control, the more positive the bias in the control level assessment. However, this experiment does not use a contingency paradigm, and it assesses illusion of control in an active task, rather than contingency judgements within a passive task.

Yarritu and Matute (2015) examined the influence of prior knowledge on the illusion of causality across two experiments. Previous research suggests that the illusion of causality can be reduced if participants are told when and how to respond to get an outcome (Hannah & Beneteau, 2009); are restricted to respond a certain number of times (Yarritu et al., 2014); asked to respond with a certain frequency (Byrom et al., 2015). Yarritu and Matute (2015) were interested in what occurs when participants are free to respond when they must achieve a certain outcome and hypothesised that prior expectations about the causal relationship between the cue and outcome will influence how participants respond.

In experiment 1, Yarritu and Matute's participants completed the adapted allergy paradigm, whereby they learnt the effectiveness of a fictitious medicine curing a fictitious disease (this is the active version of the passive contingency task used in this thesis). Participants could administer the medication for each trial if they wanted to and were then shown whether the patient had recovered. Half the participants were informed prior to the task that the recovery rate of those who have taken medication is 80% (medication expectation condition), whilst that the other half were told that participants who did not take the medication had a recovery rate of 80% (no cue expectation condition). Under both conditions the recovery was the same, patients will recover 80% of the time, which does match the prior information given to participants. However, due to the focus on different cells of the contingency table, one group should start the task expecting the medication to work, and the other group should expect the medication not to be required for recovery. During the task participants underwent 100 trials, for each trial they could choose to administer the medication, after which they would see whether the participants recovered. At the end of all the trials, participants were asked to rate the effectiveness of the medication on a 0 to 100 scale.

A significant difference was found between the probability of responding between the groups, such that the medication expectation participants administered medication more frequently than the no cue expectation group. Judgements of medication effectiveness were also higher in the medication effectiveness group, compared to the no cue expectation group. A mediational model was used with expectation as the independent variable, probability of cue administration as the mediator and judgment as the outcome variable. The test was significant and indicated that when the probability of cue was included, the effect of expectation disappeared, suggesting a total mediation of the probability of cue on judgement. The findings suggest that expectation influences responding, which is mediated by the responding rate of the participants.

In experiment two, Yarritu and Matute (2015) used a passive contingency task to examine this effect in further detail. One hundred and fourteen participants took part in a fully between-subjects experiment, using the same cover story as before. The probability of cue was manipulated to be high (0.8) or medium (0.5), and half the participants were presented with trials in which the cue occurred with a probability of 0.8, whilst the other half experienced trials with the cue occurring with a probability of 0.5. The probability of the outcome was 0.8 across all conditions. During the training phase, for each trial participants were asked a predictive question regarding whether the participants would recover from the illness, which was followed by the outcome. After all trials participants gave contingency judgements. There was a main effect of cue such that participants in the high cue probability condition gave higher judgements than in the medium density condition. There was no main effect or interaction effect for prior expectation. Overall, comparisons of experiments 1 and 2 suggests that prior expectations influence contingency judgements, depending on whether participants are able to control the cue.

However, the two experiments vary in their manipulation of expectation. In the first experiment participants were either informed that out of 10 patients who had taken the medication 8 recovered, or out of 10 patients who had not taken the medication 8 recovered. In the second experiment the prior expectation was manipulated by either a 0.8 or 0.5 probability of cue occurring over 100 trials, with probability of outcome being 0.8 for both conditions. The second experiment manipulates the probability of cue to be 0.8 or 0.5, whilst the first experiment manipulates the probability of cue to be 1 or 0, as such the cue manipulations are not equivalent, particularly as participants are only gaining knowledge regarding two trial types of a contingency table in the first experiment. Table 41 shows the contingency calculations for the high/low expectation conditions across the two experiments, and it shows that the probability of cue is not equivalent between experiments. The effect of prior expectation seen in experiment 1 may be due to participants in the no expectation group, actually having no expectation that the medication is needed, and never sampling any instances of the cue being effective. In the second experiment participants in the low probability cue group actually do experience instances of the cue being effective. As such the findings of the second experiment are not comparable to the first experiment, as the prior expectations set under the experimental conditions are not the same for the low/no expectation group.

Table 41 Contingency calculations for Yarritu and Matute (2015)

Table a

Yarritu and Matute (2015) Experiment 1 Expectation condition

	Outcome		
Cue	Present	Absent	
Present	8	2	P(O/C)=0.8
Absent	0	0	P(O/~C)=--
	P(C/O)=1	P(C/~O)=1	
	ΔP=1-1=0		

Table b

Yarritu and Matute (2015) Experiment 1 No Expectation condition

	Outcome		
Cue	Present	Absent	
Present	0	0	P(O/C)=--
Absent	8	2	P(O/~C)=0.8
	P(C/O)=0	P(C/~O)=0	
	ΔP=0-0=0		

Table c

Yarritu and Matute (2015) Experiment 2 High Expectation condition

	Outcome		
Cue	Present	Absent	
Present	64	16	P(O/C)=0.8
Absent	16	4	P(O/~C)=0.8
	P(C/O)=0.8	P(C/~O)=0.8	
	ΔP=0.8-0.8=0		

Table d

Yarritu and Matute (2015) Experiment 2 Medium Expectation condition

	Outcome		
Cue	Present	Absent	
Present	40	10	P(O/C)=0.8
Absent	40	10	P(O/~C)=0.8
	P(C/O)=0.8	P(C/~O)=0.8	
	ΔP=0.5-0.5=0		

7.3 Role of belief in passive contingency paradigms

Mutter et al. (2007) examined the role of age and prior beliefs on contingency judgements. In a pilot study 40 participants were asked to indicate the relationship (no relationship to perfect relationship) and direction of relationship (negative or positive) for 76 pairs of events. The pairs of events were reduced down to 12 pairs, 3 each for positive, negative, zero and unknown belief. Afterwards forty-eight participants took part in a 2 (age between subjects: young vs old) by 4 (belief based on prior belief from the pilot within subjects: positive, negative, zero, unknown) by 3 (objective contingency within subjects: positive, negative, zero) mixed factorial design experiment. Participants completed 12 contingency judgement problems, for each they provided a pre-test estimate of the contingency and gave a confidence rating for their estimate based on the pairs of events. For each problem participants were told that a scientific study had generated data on the relationship between the two events and they were going to see a series of these cases that would provide the new data. Each case consisted of statements about the co-occurrence of the events, there were a total of 24 cases per problem. After each problem another estimate of the contingency and confidence rating was given.

The pre-test contingency ratings all matched the expected belief rating, participants of the experiment gave similar belief ratings as those from the pilot. The post-test contingency estimates show a main effect of prior beliefs, participants gave contingency judgements which matched the beliefs of the problems. A main effect of objective contingency was found, participants gave more negative judgements for negative contingencies, and more positive estimates for positive contingencies. There was no significant interaction between the prior belief and contingency. Overall, this suggest that beliefs do affect contingency judgements.

Catena et al. (2008) examined the role of prior beliefs in a passive contingency task, using compound cues, across 4 experiments, using two scenarios. In the social scenario participants were asked to imagine spending a number of evenings with two fictitious people (A and B), in terms of pleasant and unpleasant evenings. In the medical scenario, participants imagined they were part of a research team investigating the side effects of two chemicals (A and B). Prior beliefs were induced by providing instructions. For the social scenario participants were told person A was fun and attractive, whilst person B was unattractive and boring. For the medical scenario, participants were told that chemical A was an antibiotic drug, whilst chemical B was a vitamin. In both scenarios, the link between A and the outcome was more plausible than the link between B and the outcome. A control group was also included for whom no prior belief was induced.

In experiments 1a and 1b, a 2 (prior beliefs between-ss: induced, not induced) 2 × (probability of outcome between-ss: high vs low) × 3 (cue between-ss: A, B, AB) between design was used with 106 participants per scenario. Participants were given information on the co-occurrence of a 2-cue compound and the outcome in contingency tables. In the high outcome condition, the probability that outcome occurred in the presence of AB was .90, and the probability of the outcome occurring in the absence of AB was also .90. In the low outcome condition, probability of the outcome occurring in the presence of AB was .90 however the probability of the outcome occurring in the absence of AB was .10. Half the participants were provided with the prior belief induction, whilst the other half were not. Participants were expected to learn to attribute the same causal inferences to both cues, as they never occurred individually. Contingency judgements were collected for the cue types, A, B and AB on a -100 to 100 scale. Main effects of contingency and cue were found, which suggested that higher judgements were given in the high contingency condition. And participants gave higher judgements for AB cue, followed by A and then the lowest to B. Interaction effects of prior beliefs × contingency, and prior beliefs × cue were found.

For the group which had no prior belief induction, judgements were higher in the high contingency condition compared to the low contingency condition across all cue types. In both levels of the contingency, judgements of A did not differ from B, and both were lower than AB. However, for the prior belief induced group, there were no significant effect of contingency. Although AB judgments were higher than A and B, judgments for A were also higher than B. Prior beliefs modulated the causal strength participants assigned to each element of the compound cue. When prior belief favours one cue, that cue is judged more causally effective.

To further explore influence of predictive value, in experiments 2a and 2b, participants were given the same information as experiments 1a and 1. They were also given information about the predictive value of the compound cue, and the individual elements of cue B. The same task as before was used. Main effects of contingency and cue were found again which replicates the previous experiment. However, no significant effect of prior belief was found. Overall, this suggest that prior beliefs have no impact on judgements when predictive information is available. Together these findings suggest that prior beliefs can affect contingency judgements. However, individuals do integrate knowledge to form judgements, and new evidence can change prior beliefs.

Yarritu et al. (2015) further examined the role of prior belief in passive contingency paradigms. One hundred forty-seven participants completed the adapted allergy paradigm used in this thesis. One group of participants were induced to have a strong prior belief that a placebo medication was effective in treating a fictitious disease, whilst the other group had a weaker prior belief induction. Both groups then observed a series of fictitious patients who always took the placebo treatment along with a second medication which was effective. Two different medication/disease names for used for the first and second phase of the experiment. The group which had a stronger prior belief had more difficulties in learning that the second treatment was effective compared to the group which had a weaker prior belief induction. This study shows how strong prior belief can affect consequent judgements.

7.4 Summary and rationale for a new investigation

The literature covered so far suggests that prior belief does influence judgements, although this effect can be modulated by other factors. However, these studies have examined prior beliefs that have been induced within the experimental setting and have then assessed the influence on judgments. I am interested in the real-world application of the contingency paradigm when assessing false belief formation. That is, how do pre-existing beliefs influence judgements. Prior beliefs formed under lab settings may hold no relevance to the individual as such, they may be discarding that prior belief upon new information. However, paradigms which are directly relevant to prior beliefs that participants may hold, could differ in how prior beliefs affect judgements, as such this is an important area to explore.

Whilst some studies have examined actual prior beliefs, these have used pilot studies to determine prior beliefs, and have then assumed that the experimental participants hold the same prior beliefs, which may not be the case (Evans et al., 2005; Mutter et al., 2007). There are two issues with this approach. The first being two participants in the same condition may hold different beliefs because they held different beliefs before coming into the lab. Secondly, two participants in the same condition may hold different beliefs because they respond differently to the belief formation manipulation.

Here, I would like to make a distinction between expectation and belief. Expectation relates to what individuals predict will happen, whether based on experience or something else, while belief is something which individuals consider to be true, perhaps even without evidence. *Essentially expectations are something I predict will happen, whilst beliefs are something I have determined to be true, even if I lack the evidence to support it.* The role of prior beliefs is of interest here for this thesis, as expectations may be changed based on further experience, however beliefs may be more resistant to change, if those are strong or extreme beliefs that one holds.

Many of the studies presented above are measuring expectation rather than belief. In the studies participants are trained to expect certain things. Afterwards participants are assessed on whether or not they can overcome the pre-trained expectations. I think it is unlikely that participants will come to have beliefs about abstract or even non-abstract cue-outcome pairings. In particular as the participants were just told about these pairings within the experimental context, and the pairings do not have personal relevance to the participants. Are the participants going to use these expectations they have been trained in, in their daily lives to make decision? Highly unlikely, as they are not relevant to them at all.

As such, many of the studies above may have weak external validity, because beliefs are something participants have confidence in, being true for them, regardless of what evidence they may be presented with. As such decisions made using beliefs participants hold are going to hold far more importance, than decisions or judgements made in an experiment about things that hold no personal relevance to participants. To say these things are the same, is a fundamental flaw.

Even if we consider expectation and beliefs to be the same, the studies are still flawed because beliefs are fundamentally something people place some level of importance on, so it impacts decision making. The cover stories/contexts used may have little relevance to participants' personal beliefs, so the way in which belief would affect judgement in real life, is fundamentally going to be different to how "belief" in an experimental study will affect decision making.

This experiment will investigate the influence of existing prior beliefs on contingency judgements. I predict that prior belief will influence judgements, such that the stronger the prior belief, the stronger the influence on the judgement. In addition to this, I predict that the outcome density effect will be robust. This experiment investigates the final research question of this thesis, "Do factors beyond methodological variations influence the outcome density effect?". The factor of focus are prior beliefs.

As the influence of existing prior belief is of interest for this study, there is a need to use different cover stories to the ones used thus far within this thesis, which matches the prior beliefs being examined. Within the contingency research several cover stories have been used, as previously outlined within chapter 2. The type of cover stories used can be categorised into active and passive cover stories. For the active cover stories participants are asked to provide some kind of response which may result in an outcome and then provide estimates of control over the outcome. For the passive paradigms participants view a series of cue-outcome pairings, and at the end provide estimates of contingency.

Some cover stories for the active paradigms are fairly simple, such as participants aiming to illuminate a light bulb (Alloy & Abramson, 1979; Gillan et al., 2014a; Jenkins & Ward, 1965; Msetfi et al., 2005), controlling a blue flash of light on a computer screen (Matute et al., 2007b), or assessing the extent to which pressing the spacebar on a computer keyboard caused the appearance of a geometric shape (Crump et al., 2007; Vallée-Tourangeau et al., 2005). Other active paradigm cover stories have a more real-world reflection, such as participants pressing a button and shooting lasers at Martian spaceships to prevent them from landing and invading earth, in a video game format (Matute et al., 2014), or participants were asked to control the stereo systems' in different parts of a house (Msetfi et al., 2013a).

Cover stories for the passive contingency tasks include abstract contingencies between geometric shapes (Crump et al., 2007), sounds and pictures (Kattner & Ellermeier, 2011), and video game formats where people assess the ability of different plants to grow (Moreno-Fernández et al., 2017; Wasserman et al., 1996). Some of the more complex cover stories within passive contingency paradigms are related to medical scenarios such as foods causing allergies (Lee & Lovibond, 2021; Shanks & Lopez, 1996; Vadillo et al., 2011a, 2013), chemical that affect bacteria (Klauer & Meiser, 2000), or the adapted allergy task where medication may cure symptoms (Barberia et al., 2019; Blanco et al., 2013; Chow et al., 2019b; López et al., 1998; Matute et al., 2011b; Yarritu et al., 2015).

To ensure participants would be familiar with the cover stories chosen, I ran a pilot study to explore the recognition and views of different cover stories. All cover stories were based on everyday scenarios, and part of the pilot work was designed to identify cover stories that would be familiar to most participants. The four cover stories which were the most recognised, had the most range in responses yet participants were most neutral on will be chosen for the experiment.

7.5 Methods

7.5.1 Stimuli and measures

An exploratory survey was used to choose the four cover stories. Twenty-four cover stories were identified (see Table 42 for the cover stories). Key criteria for identifying a potential cover were (a) it included a putative association between two variables, (b) it was deemed plausible that lay people might have a prior belief about that association (e.g., due to media reports or personal beliefs), (c) there could be individual differences in prior beliefs about the putative association (e.g., alternative views are expressed in the media).

7.5.2 Procedure

Participants took part in a 10-minute survey which was advertised on Prolific, for which they were awarded £1.25. After reading the information sheet and providing consent, demographic information such as age and gender were collected, and participants completed the following two questions. Participants were asked “have you heard of the following topics/issues/ proposal?” with “yes” and “no” options for all items in Table 42. For any questions which participants had answered “yes”, their belief for those questions were then assessed using a 7-point Likert scale, see Table 43 for full questions and Likert scale options. For the first question the cover stories were worded neutrally (e.g., privatised healthcare and wait times) to avoid biasing the participants' answers in the second question

Table 42. The 24 cover stories used in the pilot to measure participants recognition rate

Cover stories
Brexit affecting the UK farming industry
Brexit affecting the UK fishing industry
Energy price caps and energy bills
Gun control and gun related crimes
Mandatory vaccines and illness
Marijuana use and crime related to drugs
Marijuana use and health
Minimum wage and joblessness
Privatised healthcare and health associated costs
Privatised healthcare and wait times
Social media use and sociability
Tax on residential waste and overall waste production
Tax on sugar and changes in sugary products bought by consumers
Taxes on healthy foods and healthy foods bought by consumers
Taxes on petrol and diesel and use of cars
Technology and jobs
Universal basic income and benefits to society
Universal basic income and people working
Use of activated charcoal having health impacts
Use of cannabis oil and cancer
Use of detox teas having health impacts
Use of essential oils having health impacts
Use of ginger and turmeric affecting viruses
Vitamin C supplements and catching a cold

To note, these scenarios were identified in September 2022

Table 43 Question 2 cover story and Likert-scale wording measuring participants' belief

Cover story wording Likert scale wording						
1	2	3	4	5	6	7
Brexit has ... for the UK farming industry						
Brexit has … for the UK fishing industry						
Privatising healthcare has … for health associated costs						
Universal basic income has ... for those who work						
Tax on residential waste has ... for overall waste production						
Universal basic income can provide society with…						
Many benefits	Some benefits	Limited benefits	Neither benefits nor harms	Limited harms	Some harms	Many harms
Imposing energy price caps can ... energy bills						
Increased gun control can … gun crime						
Legalising marijuana use can ... drug crimes						
Creating a minimum wage … joblessness						
Social media use can ... sociability						
Tax on sugar … sugar consumption						
Taxes on healthy foods has ... health food consumption						
Taxes on petrol and diesel … use of cars						
New technology can ... jobs						
Use of ginger and turmeric … protection against viruses						
Often increases	Sometimes increases	Occasionally increases	Neither increases nor decreases	Occasionally decreases	Sometimes decreases	Often decreases
Mandatory vaccinations have...						
Marijuana use has …						
Use of activated charcoal has…						
Use of detox teas…						
Use of essential oils has…						
Many health benefits	Some health benefits	Limited health benefits	Neither benefits nor harms	Limited health harms	Some health harms	Many health harms
Vitamin C supplements can ...a cold						
Often prevent	Sometimes present	Occasionally present	Neither prevent nor cause	Occasionally cause	Sometimes cause	Often cause
Privatising healthcare can … wait times						
Often improve	Sometimes improve	Occasionally improve	Neither improve nor worsen	Occasionally worsen	Sometimes worsen	Often worsens
Use of cannabis oil…						
Often successfully treats cancer symptoms	Sometimes successfully treats cancer symptoms	Occasionally successfully treats cancer symptoms	Neither treats no worsens cancer symptoms	Occasionally worsens cancer symptoms	Sometimes worsens cancer symptoms	Often worsens cancer symptoms

on their views.

7.5.2 Participants

A total of 100 participants were recruited via Prolific for the study, of whom 25 were women. The modal age group was 26 to 30 years, within age-groups ranging from 18-20 to 65-70 years.

7.5.3 Data analyses / selection criteria

The following criteria was adopted to select 4 cover stories for the experiment. The exclusion of cover stories:

- Recognised by less than 80 people to ensure the chosen cover stories are highly recognisable to participants.
- With a median rating less than 3 or above 5 to avoid any cover stories that overall people have homogenous extreme beliefs on.
- With standard deviation under 1.5 to avoid any cover stories that people are generally in agreement on

7.6 Results and Discussion

Table 44 and Figure 18 show the participant recognition levels for each cover story, along with median and standard deviation for participants' belief. Participant recognition was based on question 1 where participants identified all the topics, they were familiar with. The median and standard deviations were calculated from the responses to question 2.

The following 5 cover stories were left over after the inclusion criteria, previously outlined was applied: marijuana use and health, privatised healthcare and wait times, technology and jobs, social media use and sociability, and energy price caps and energy bills.

Energy price caps and energy bills, whilst a valid cover story to use, in September to November 2022 when this study was going to take place, new policies on energy price caps and energy bills were being introduced, as such this was discarded to avoid a general bias in views. Participants potentially could be hyper-aware, due to extensive media coverage and the price caps drastically affecting energy bills, as such beliefs might change rapidly between the pilot and actual study as such there might no longer be a range of prior beliefs. Consequently, the other 4 cover stories were used for the experiment.

High variability in ratings for a given scenario are assumed to be due to differences in prior beliefs between participants. However, it is important to acknowledge that for some scenarios it may be difficult to provide a single judgement as there may be multiple contrasting opinions and thoughts held by individual participants. For example the use of cannabis may be beneficial for pain relief, however it may also be harmful as it can increase the risk of schizophrenia (Sharma et al., 2012). As such there may be within participant variability in prior beliefs over time.

Table 44. Participant recognition rate, median, and standard deviation

	Participant recognition (max 100)	Median[1]	Std. Deviation
Energy price caps and energy bills****	97*	4**	2.04***
Minimum wage and joblessness	94*	4**	1.38
Universal basic income and people working	93*	2	1.27
Social media use and sociability****	91*	5**	1.80***
Tax on sugar and changes in sugary products bought by Consumers	91*	5**	1.35
Brexit affecting the UK farming industry	89*	6	1.67***
Taxes on Petrol and diesel and use of cars	88*	5**	0.76
Marijuana use and health****	88*	3**	1.68***
Mandatory vaccines and illnesses	88*	2	1.89***
Universal basic income and benefits to society	88*	2	1.39
Privatised healthcare and health associated costs	87*	6	1.78***
Vitamin C supplements and catching a cold	87*	3**	0.86
Gun control and gun related crimes	85*	6	1.53***
Technology and jobs****	85*	4**	1.79***
Brexit affecting the UK fishing industry	84*	6	1.39
Privatised healthcare and wait times****	83*	3**	1.61***
Marijuana use and crime related to drugs	76	5**	1.79***
Use of detox teas having health impacts	66	3**	1.57***
Use of essentials oils having health impacts	64	3**	1.40
Use of cannabis oil affecting cancer	64	3**	0.95
Use of activated charcoal having health impacts	50	3**	1.21
Use of ginger and turmeric affecting viruses	48	2	1.03
Taxes on healthy foods and healthy foods bought by consumers	45	6	1.66***
Tax on residential waste and overall waste production	39	3**	1.70***

[1] 1-7 response scale with ‘4’ as a neutral response (no association)

* More than 80 participants recognised this cover story

** median was between 3 and 5

Figure 18. Recognition rate (left-side vertical scale), median and standard deviation for judged association (right-side vertical scale)

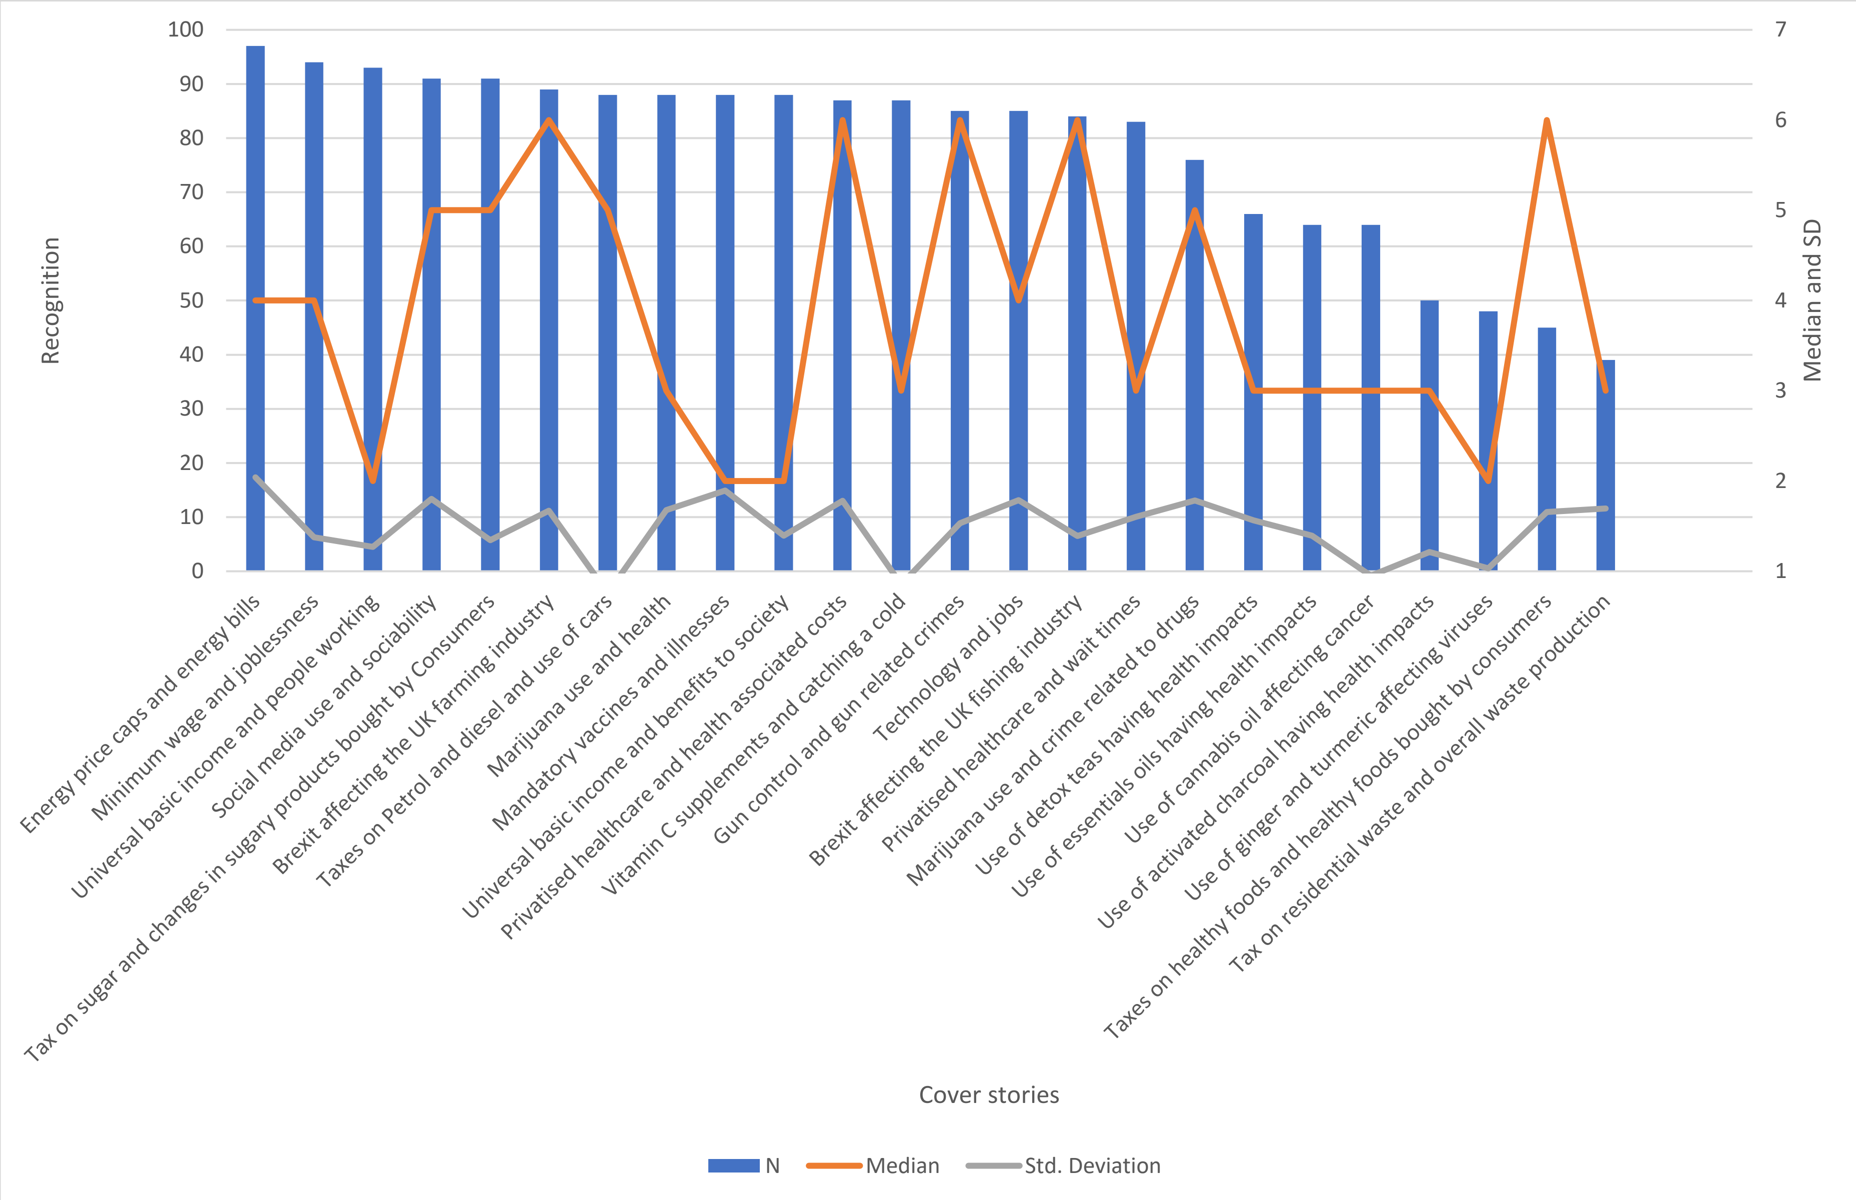

Chapter 8: Scenarios and Prior Beliefs Experiment

Prior literature (Chapter 7) shows that prior beliefs can affect judgements, such that strong prior beliefs are linked to better learning of contingencies when the observed cue-outcome association matches the prior belief. However if the observed cue-outcome association is in the opposite direction to the prior belief then learning takes longer to occur and judgements can be more biased and resistant to change (Biele et al., 2009; Yarritu et al., 2015). In addition to this, research also suggests that the effects of prior belief can be modified based on factors such as evidence of the predictive value of cues (Catena et al., 2008). However, the influences of pre-existing prior beliefs on judgements in passive contingency tasks have not been explored thus far, rather the prior beliefs are often induced within the experimental setting. As such, this experiment investigates the influences of existing prior beliefs on judgements. I predict that prior belief will influence judgements, such that stronger positive prior belief, will result in more positive judgments.

To assess the influence of prior beliefs on contingency judgements, new cover stories are required to ensure that (a) participants likely hold a belief about the association specified in the cover story that was formed prior to the study, and (b) those prior beliefs are heterogeneous. As such, the previous chapter outlined the pilot study used to select the cover stories that are used in this experiment. The majority of cover stories used within the passive contingency paradigm are related to medical scenarios such as foods causing allergies (Lee & Lovibond, 2021; Shanks & Lopez, 1996; Vadillo et al., 2011a, 2013), chemical that affect bacteria (Klauer & Meiser, 2000), or the adapted allergy task where medication may cure symptoms (Barberia et al., 2019; Blanco et al., 2013; Chow et al., 2019b; López et al., 1998; Matute et al., 2011b; Yarritu et al., 2015). As such variations in judgements influenced by cover stories will also be examined in this experiment. In sum, different cover stories have been used to study the outcome density effect in contingency learning tasks. While different cover stories are used in different studies, there is little (if any) research that formally examines variation in contingency judgment as a function of cover story.

To summarise, this experiment explores the influence of prior belief on contingency judgements. I predict that stronger positive prior belief, will result in more positive judgments. In addition to this the effect of cover story on contingency judgements are examined, with no predictions made as this is more of an exploratory aim. Finally, I predict that the outcome density effect will be robust across the new cover stories, due to its robustness in previous experiments in this thesis and the published literature which uses different scenarios. The main goal of this chapter is to examine whether the outcome density effect is robust. This chapter also examined how methodological variations in the contingency learning paradigm affect the size of the outcome density effect. In addition to this, the chapter examines whether factors (prior belief) beyond methodological variations influence the outcome density effect.

8.1 Methods

This experiment investigates the influence of prior belief on contingency judgements. It is possible that the act of asking about prior beliefs in a study may itself alter the influence of those beliefs. To gain some understanding of this possibility, we manipulate the time-point for prior belief measurement—measuring belief either at the beginning or the end of the study session. The influence of cover story on judgements in a contingency paradigm and the differences in the outcome density effect between different cover stories is also investigated.

8.1.1 Design

A 4 × 2 × 2 mixed design was used, with judgement as the dependent variable. Three factors were varied in the experimental design: cover story (within subjects: marijuana, social media, technology, healthcare) by density level (within and between subjects: high density vs. low density) by prior belief measurement (between subjects: measured prior to contingency-judgement task vs measured after task). The prior belief questions were the same as those used in the pilot (Table 45). The four cover stories chosen from the pilot were used. All participants completed 4 contingency paradigms, one for each cover story. Outcome density was manipulated to be either low (0.2) or high (0.8) while maintaining a zero contingency. Cue density was held at a medium density of 0.5, for both high and low cue. Participants completed two low density blocks and two high density blocks. The order

Table 45 Prior belief questions

Cover story wording Likert scale wording						
1	2	3	4	5	6	7
Social media use can ... sociability New technology can … jobs						
Often increases	Sometimes increases	Occasionally increases	Neither increases nor decreases	Occasionally decreases	Sometimes decreases	Often decreases
Marijuana use has …						
Many health benefits	Some health benefits	Limited health benefits	Neither benefits nor harms	Limited health harms	Some health harms	Many health harms
Privatising healthcare can … wait times						
Often improve	Sometimes improve	Occasionally improve	Neither improve nor worsen	Occasionally worsen	Sometimes worsen	Often worsens

of the density of blocks were randomised. Half of the participants completed a questionnaire on their prior beliefs on the 4 cover stories before the contingency paradigms, whilst the other half completed these questions after the contingency paradigm. See Figure 19 for participant allocation.

8.1.2 Task and stimuli

The task was adapted from the task described in previous chapters using new stimuli images (see Figure 20), cover stories and judgement questions (see appendix A) were which had been tailored to match each of our newly devised cover stories. Participants provided contingency judgements on a -100 to 100 scale. The task was built on Gorilla (Gorilla Experiment Builder, 2020), as an online task. Participants could complete the task from their location of choice.

Figure 19. Participant allocation across conditions

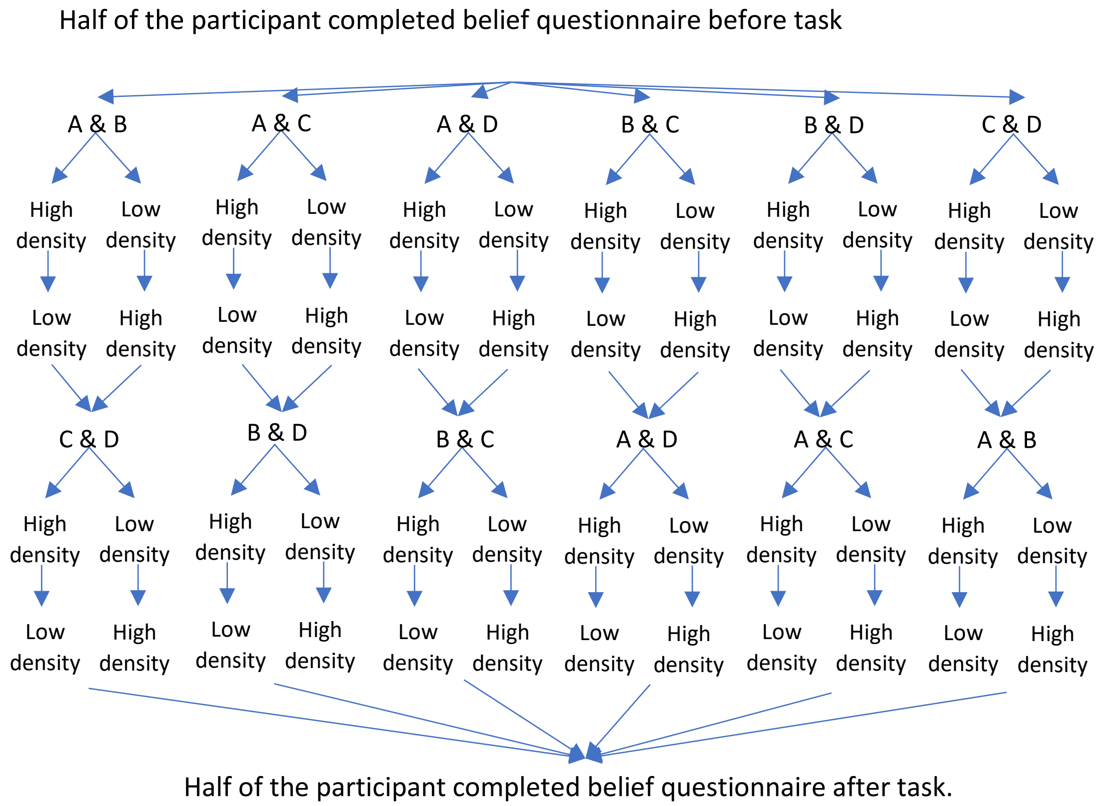

8.1.3 Procedure

The procedure was similar to the previous studies. Participants signed up on the Prolific research participants platform, then read the information sheet and completed the

Figure 20. Experimental stimuli images. Note. All the stimuli images are from Freepik.com

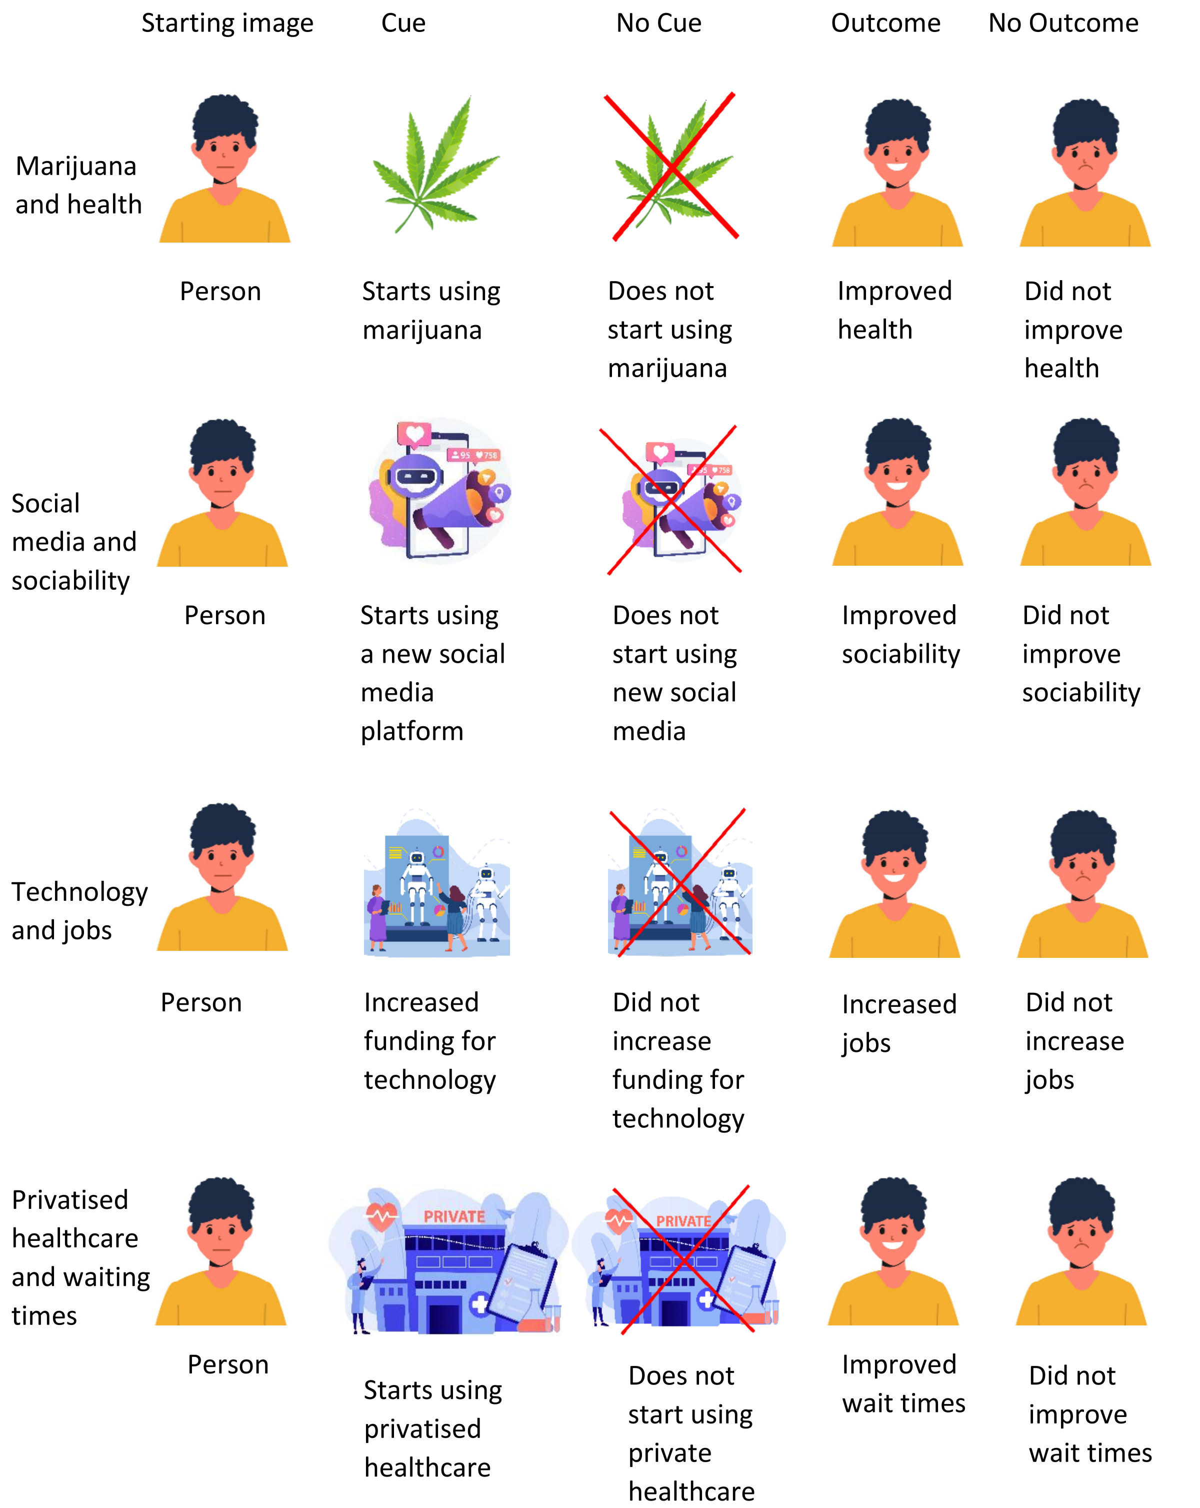

online consent form. After completing demographic questions, half of the participants completed the belief questionnaire items. Afterwards all participants took part in 4 contingency judgement tasks and reported contingency judgements on a -100 to 100 scale for each cover story. Identical timings were used for all trials. Participants were randomly allocated to the order of cover story, belief and density condition (Figure 19). Participants who had not completed the belief questionnaire yet completed it at this point. The experiment took around 40 minutes to complete.

8.1.4 Participants

A power analysis was performed to detect a medium sized effect (d=0.5), with a power of 1 – β=0.80 and Type-I error of α=.05 (two sided), the sample size required was approximately N=100. Hence a recruitment target of 200 participants is proposed to maintain the target level power within any 50/50 split of subsamples (e.g., by level of prior belief). A total of 204 participants were recruited and were paid £6.00 for completion of the study. As such the study provides an adequate sample size to test the hypotheses. There was a total of 93 women, and the modal age groups were 31-35 and 36-40 years, with age groups ranging from 18-20 to 70+.

Prior belief was coded from -3 to 3, so that a positive score indicates a positive prior belief, whilst a negative score indicates negative belief.

8.2 Results

8.2.1 Outcome density

A paired samples t-test was conducted to test the outcome density effect across all conditions. To do so, the average of the two low density responses and of the two high density responses were used. Participants gave significantly lower contingency judgements in the low-density condition *(M*=-12.52, *SD*=28.82) compared to the high-density condition *(M*=16.75, *SD*=24.21), *t*(203)=10.532, *p* < .001, Cohen's *d*=0.74. Thus, an outcome density effect was observed. It also seemed slightly smaller than the average effect detected in the previous studies of this thesis. Previous outcome densities found in this thesis ranged from 0.60 to 1.49.

8.2.2 Cover story

Four independent samples *t*-tests were conducted to test the outcome density effect within each cover story. Participants gave significantly lower judgements in the low density conditions compared to the high density conditions across all cover stories, marijuana effects *t*(202)=7.11, *p* < .001, social media effects *t*(202)=3.83, *p*< .001, technology effects *t*(202)=5.56, *p*< .001, and healthcare effects *t(*202)=7.39, *p*< .001 (Table 46).

Cover story was manipulated within participants. However, outcome density was partially within-between, as each participant did two low and two high outcome density

Table 46. Mean (SD) judgements by manipulation of cover story and density

	Marijuana effects	Social media effects	Technology effect	Healthcare effects
Mean judgement *(SD)*	2.43 *(39.62)*	-2.83 *(38.34)*	1.91 *(34.81)*	6.98 *(39.69)*
Participants in high density condition	102	102	103	101
Mean High density response *(SD)*	20.12 *(33.24)*	7.15 *(37.05)*	14.44 *(28.76)*	25.42 *(31.46)*
Mean Low density response *(SD)*	-15.26 *(37.66)*	-12.81 *(37.37)*	-10.87 *(35.92)*	-11.11 *(38.70)*
95% CI for mean difference	25.58, 45.19	9.69, 30.23	16.34, 34.28	26.78, 46.27
Outcome density effect size d	1.00	0.54	0.78	1.04
95% CI for Cohen's d	0.704,1.286	0.257,0.815	0.493,1.063	0.741,1.326

conditions. A random intercept model was used to examine these effects in more detail. Judgements was the dependent variable, with 4 (cover story: marijuana vs social media vs technology vs healthcare) by 2 (outcome density: low vs high) as the independent variable.

There was a large main effect of density $F(1, 816)=142.93$, $p < .001$, participants gave higher judgements in the high outcome density condition compared to the low outcome density condition. There was a main effect of cover story $F(3, 816)=2.78$, $p=.040$, participants gave the highest judgements in the healthcare cover story, followed by marijuana, technology, and the lowest judgements in the social media condition, see Table 46 for means and standard deviations. There was an interaction effect of cover story by density $F(3, 816)=2.67$, $p=.047$; in the high-density condition participants gave judgements from high to low in the order of healthcare, marijuana, technology, social media, which reflects the main effect of cover story. However, for the low-density condition, the more positive judgement was given to the technology cover story, followed by healthcare, social media, with the most negative judgements given to marijuana. Another way to characterise this interaction is that the size of the outcome density effect varied across the scenario, being medium-sized ($d=0.5$) for social media effects and large ($d=1.0$) for healthcare and marijuana effects.

8.2.3 Time of prior belief measurement

Time of prior belief measurement and cover story

Four two-way ANOVAs were conducted to examine the effect of outcome density and time of belief measurement on judgement in each cover story by manipulation of cover story and density, by time of belief measurement. There was a main effect of density type across all cover stories. There was no main effect of time of belief measurement, nor an interaction effect of density by time of belief measurement, see Table 47. Based on these analyses, it is implausible that any effect of measuring prior belief is anything other than a small effect.

Time of prior belief measurement and prior belief

To test if the contingency task influenced participants general beliefs about the association, four one-way ANOVAs were conducted with time of prior belief measurement (before contingency task vs after contingency task) as the independent between subject factor. The dependent variables were the 4 prior beliefs collected for marijuana, social

Table 47 Four two-way ANOVAs to examine effect of density and time of belief measurement on contingency judgement, for each cover story. F(1, 200)=

	F	*p*	ηp2
Marijuana cover story			
Time of belief measurement	0.01	0.937	0.00
Density	50.68	<.001	0.20
Density * Time of belief measurement	1.15	0.286	0.01
Social media cover story			
Time of belief measurement	0.25	0.618	0.01
Density	14.52	<.001	0.07
Density * Time of belief measurement	3.14	0.078	0.02
Technology cover story			
Time of belief measurement	1.45	0.23	0.01
Density	30.87	<.001	0.13
Density * Time of belief measurement	0.63	0.429	0.00
Healthcare cover story			
Time of belief measurement	2.16	0.144	0.01
Density	54.93	<.001	0.22
Density * Time of belief measurement	0.09	0.763	0.00

media, technology, and healthcare questions for all participants. There was no main effect of time of belief measurement for any of the scenarios (Table 48).

Overall, measured (‘prior’) belief did not differ significantly between the two measurement times. As such the contingency task has not substantially influenced the stated prior beliefs of participants. Due to a lack of main effects of time of prior belief measurement, this variable will not be used for further analysis.

8.2.4 Prior belief effects

Figure 21 shows the spread of prior beliefs responses across the four scenarios. Overall participants seem to have more positively skewed prior beliefs for the marijuana and privatised healthcare scenarios, the median response was 1 for marijuana, technology, and healthcare, and 0 for social media. With the possible exception of the marijuana scenario, the study has successfully identified scenarios for which there was considerable heterogeneity for prior beliefs.

Prior belief, cover stories, and contingency judgements.

Pearson’s correlations were conducted for the correlation between prior beliefs and judgements. The judgements were split by low density and high-density conditions, see Table 49. Overall, the prior beliefs for marijuana, social media, and technology, show a

Table 48 Four one-way ANOVAs to examine effect of contingency task on belief, for each cover story. F(1, 202)=

	F	*p*	η_p^2
Marijuana belief	3.21	.075	0.02
Social media belief	0.82	.367	0.01
Technology belief	0.12	.734	0.00
Healthcare belief	0.18	.671	0.00

significant positive correlation with the high-density condition of the respective contingency

Figure 21. Spread of prior belief across scenarios

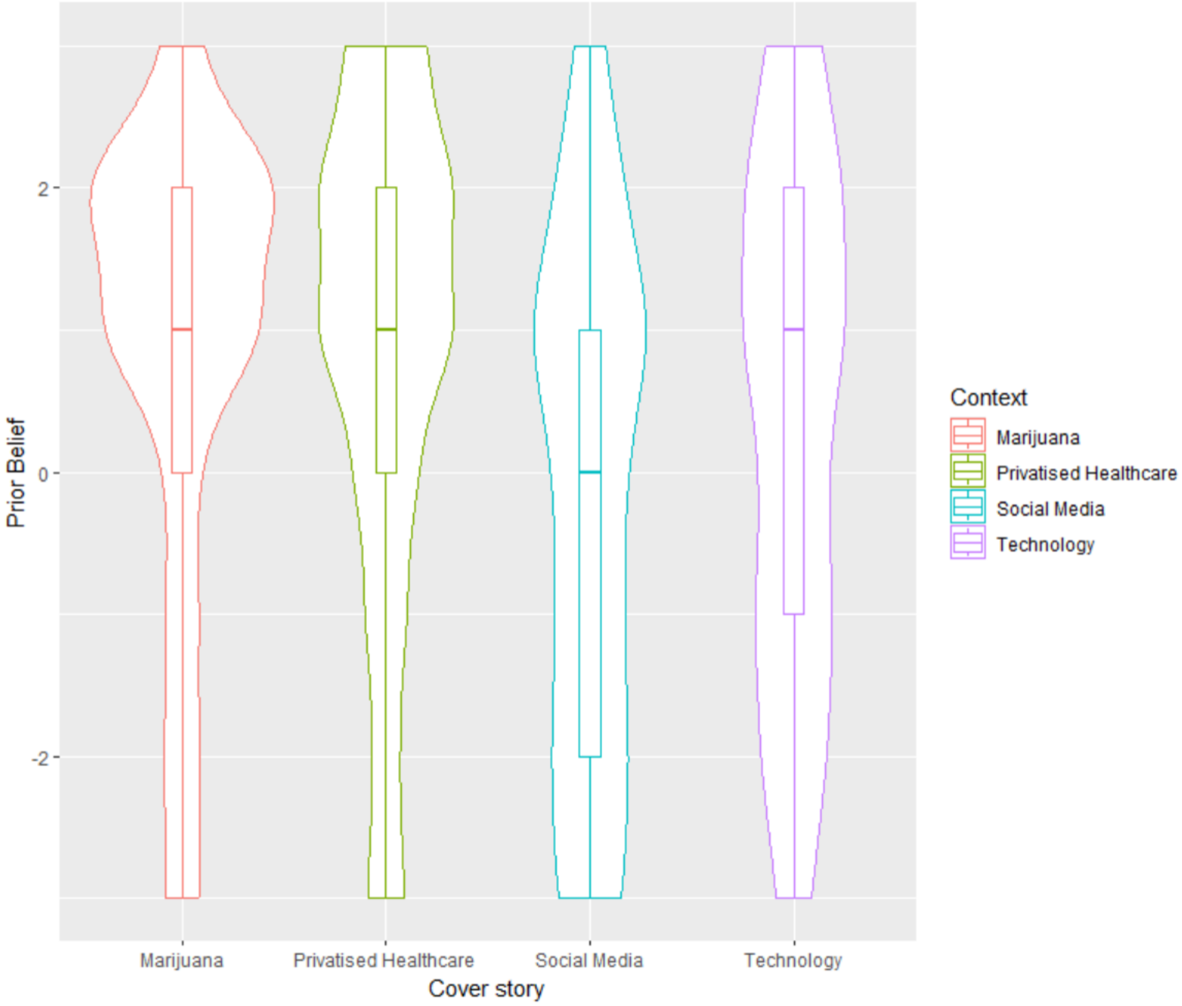

Table 49.Correlations between prior beliefs and judgements (n=102)

Density condition	Prior beliefs			
	Marijuana	Media	Technology	Private
Marijuana high	0.273**	-	-	-
Marijuana low	-0.085	-	-	-
Media high	-	0.437**	-	-
Media low	-	0.013	-	-
Technology high	-	-	0.256**	-
Technology low	-	-	-0.023	-
Private high	-	-	-	0.162
Private low	-	-	-	-0.021

**. Correlation is significant at the 0.01 level (2-tailed).

condition. Stronger prior beliefs for these scenarios are associated with more positive contingency judgements in the high outcome density condition.

A multiple regression was run to predict contingency judgement for each cover story, based on density type, prior belief, density type x prior belief (Table 50 and Figure 22). For the marijuana cover story, the regression was significant $F(3, 200)=19.89$, $p < .001$, adjusted $R^2=.218$. The level of density, and the interaction between density and prior belief significantly predict contingency judgements (Table 50a). For the social media cover story, the regression was significant $F(3, 200)=12.35$, $p < .001$, adjusted $R^2=.144$. The level of density, prior belief, and the interaction between density and prior belief significantly predict contingency judgements (Table 50b). For the technology cover story, the regression was significant $F(3, 200)=12.28$, $p < .001$, adjusted $R^2=.143$. The level of density significantly predicts contingency judgements (Table 50c). For the healthcare cover story, the regression was also significant $F(3, 200)=18.92$, $p < .001$, adjusted $R^2=.209$. The level of density significantly predicts contingency judgements (Table 50d).

Table 50. Multiple regression for each cover story, by density, prior belief, and density x prior belief interaction

Table a Multiple regression for Marijuana cover story

	B	95% CI	*β*	*t*	*p*
Density	14.46	[9.02, 19.91]	0.37	5.24	<.001**
Prior Belief	1.92	[-1.07, 4.90]	0.08	1.26	.208
Density x prior belief interaction	3.83	[0.84, 6.82]	0.18	2.53	.012**

Note. R^2_{adj}=.218 (N=204, p <.001). CI=confidence interval for *B*

** significant at the 0.05 level

Table b Multiple regression for Social media use cover story

	B	95% CI	*β*	*t*	*p*
Density	9.89	[4.94, 14.85]	0.26	3.94	<.001**
Prior Belief	4.51	[1.77, 7.24	0.21	3.25	<.001**
Density x prior belief interaction	4.24	[1.50, 6.97]	0.20	3.05	.003**

Note. R^2_{adj}=.144 (N=204, p <.001). CI=confidence interval for *B*

** significant at the 0.05 level

Table c Multiple regression for Technology cover story

	B	95% CI	*β*	*t*	*p*
Density	12.06	[7.53, 16.58]	0.35	5.26	<.001**
Prior Belief	1.78	[-0.71, 4.26]	0.09	1.41	.160
Density x prior belief interaction	2.24	[-0.25, 4.73]	0.12	1.78	.077

Note. R^2_{adj}=.143 (N=204, p <.001). CI=confidence interval for *B*

** significant at the 0.05 level

Table d Multiple regression for Privatised healthcare cover story

	B	95% CI	*β*	*t*	*p*
Density	16.49	[10.97, 22.01]	0.42	5.89	<.001**
Prior Belief	1.43	[-1.49, 4.35]	0.06	0.97	.334
Density x prior belief interaction	1.87	[-1.05, 4.79]	0.09	1.27	.207

Note. R^2_{adj}=.209 (N=204, p <.001). CI=confidence interval for *B*

** significant at the 0.05 level

Figure 22 Scatter plots of beliefs by contingency judgements for each cover story. Regressions lines by density level

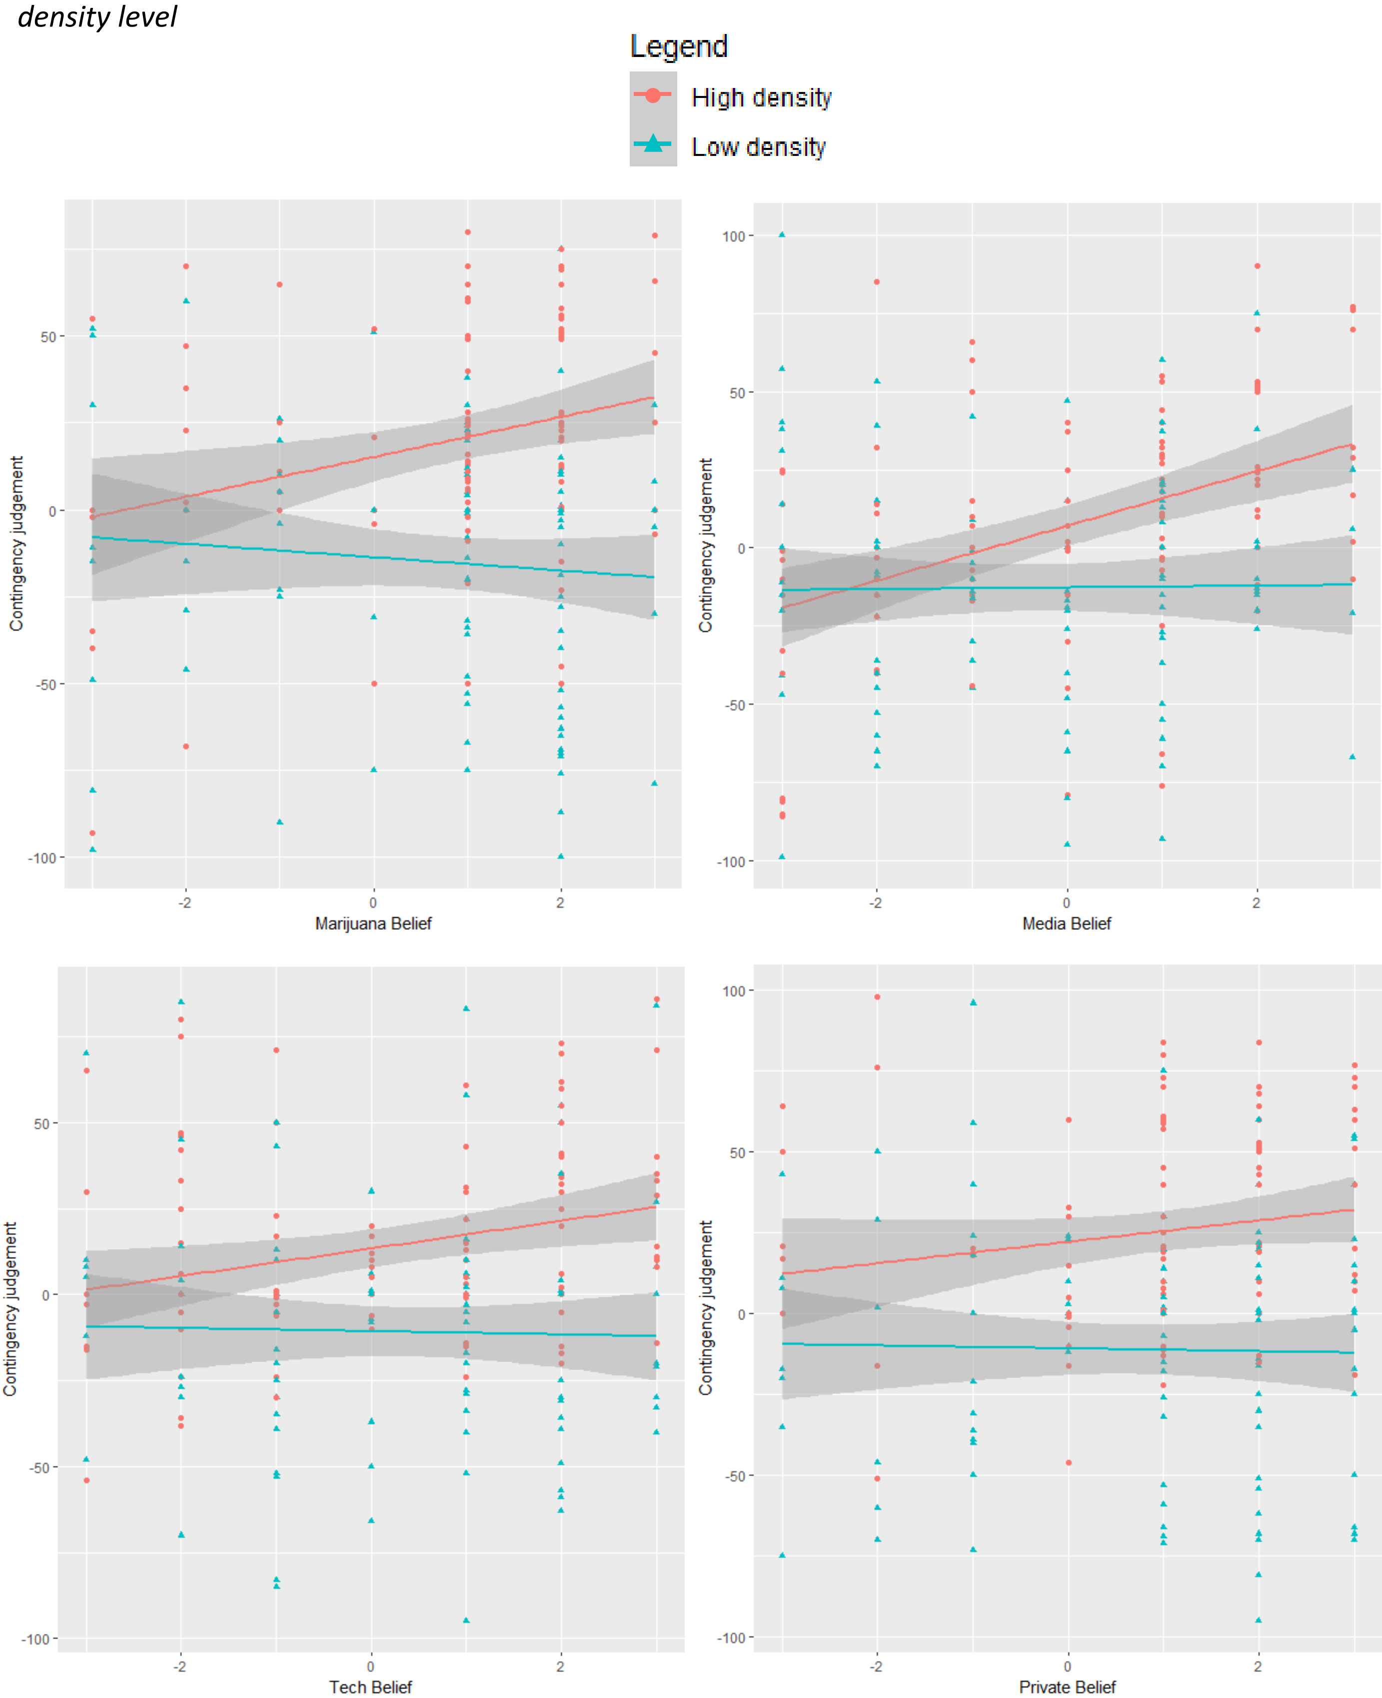

8.3 Discussion

This study used a contingency paradigm to investigate whether prior beliefs and cover stories influence judgements. I hypothesised that prior belief would influence judgements, such that people with stronger positive prior beliefs will have higher judgements. In addition to this, I also predicted that there will be a difference in judgments across the 4 different cover stories and that there would be an effect of outcome density. The study found support for the influence of effect of cover story and outcome density, and partial support for the effect of prior belief.

8.3.1 Outcome density effects

As expected, the outcome density effect was robust, and appeared in all conditions across the experiment. Participants consistently gave higher contingency judgements in the high outcome density condition compared to the low outcome density condition, replicating the findings of the previous experiments in this thesis. It is interesting that in the low outcome density conditions the mean is negative, indicating that the perceived contingency is negative, when the programmed contingency was zero. In addition to this, the mean of the high-density condition indicates that participants did not experience a highly positive contingency.

8.3.2 Cover story effects

As expected, a cover story effect was also found. Overall participants gave the highest judgements in the healthcare condition, followed by marijuana, technology, and the lowest judgements in the social media condition. This aligns with prior beliefs: in the healthcare and marijuana scenario, prior beliefs are generally positive, whereas in the social media scenario participants held mixed views. However, cover story did have an interaction effect with outcome density, such that for the high outcome density condition, the same pattern is followed. However, for the low outcome density condition participants gave the highest judgement in the technology cover story, followed by healthcare, social media, and marijuana. This suggests that the high and low outcome density conditions are not experienced in the same way across the cover stories. In particular in the low outcome density condition, the mean contingency rating given was negative across all cover stories, however, it was most negative in the marijuana setting. As such these findings suggest that judgements of contingency depend on density type and cover story, with the size of the outcome density effect varies, being largest in the healthcare and marijuana scenarios and smallest in the social media scenario. For the social media cover story, the mean judgement across conditions was negative, suggesting that participants perceived a negative contingency overall for this cover story.

8.3.3 Time of Prior Belief measurement

The time of belief measurement was manipulated to see if the contingency task influenced belief. Also, time of belief measurement was manipulated to see if the contingency judgements varied depending on when prior beliefs were collected. Potentially by collecting prior belief before the contingency judgements, participants may be primed to think about the scenarios, which could influence their contingency judgements. However, time of prior belief measurement was not significant across either of these analyses, suggesting that the contingency task itself does not influence the pre-existing beliefs of participants, and participants were not influenced by thinking about the beliefs regarding the cover stories before the task.

8.3.4 Prior belief effects

Prior belief and scenario

There was a main effect of prior belief across the scenario found, which suggested that participants vary in their beliefs across the four scenarios of marijuana, social media, technology, and healthcare. This was expected, as in the pilot participants also varied in their beliefs across this scenario. Interestingly the spread of responses was different across the four scenarios, with some more positively skewed (marijuana and healthcare) whilst others were more evenly distributed (social media and technology).

Prior belief, cover stories, and contingency judgements

Across all the cover stories a main effect of density was found suggesting that regardless of cover stories, participants overall gave higher contingency judgements in the high outcome density condition compared to the low outcome density condition. I had predicted an effect of prior belief; however, this was only partially supported. Prior belief of the marijuana, media, and technology scenarios, positively correlated with the high-density condition of the respective cover stories. This suggests that those with stronger prior beliefs gave higher judgements in the high-density condition. However, this effect only held true for the social media cover story with the inferential analysis.

For the marijuana and social media cover stories there was also an interaction of prior belief by density, suggesting that for both cover stories, the contingency judgements participants made, are dependent on their prior belief and the density type. The interaction indicates that for the low outcome density participants generally give lower contingency judgements regardless of prior belief. However, for the high outcome density condition, prior belief affects contingency judgements, such that more positive prior beliefs result in more positive contingency judgements, and more negative prior beliefs, results in more negative contingency judgements. Although this effect seems more the case for very positive or very negative prior belief. This is particularly interesting as those who held a more negative prior belief, perceived a negative contingency in the high-density condition, thus far in this thesis, the high-density condition has always been perceived as positive and never less positive than the low outcome density condition, which is the case here. This supports the hypothesis that extreme prior beliefs influence contingency judgements and the findings of the prior literature (Catena et al., 2008; Czupryna et al., 2021; Yarritu & Matute, 2015). These results imply a kind of preparedness to learn or mislearn in line with expectation. The results also suggest that for these cover story, when the prior belief is so strong, the outcome density bias can be overcome, and participants can perceive a negative contingency, even with a high outcome density condition. The literature thus far has shown that participants pretty much always perceive a positive contingency in a high outcome density, even when the contingency is zero. However, these findings suggest that strong prior beliefs can overcome this effect.

This effect was found for the social media cover story, this may be due to the fact that the outcome density effect itself was much smaller for this cover story (0.54), compared to the other cover stories, and the only cover story for which the mean contingency judgement was negative. Although it was also found for the marijuana cover story which has a larger effect size (1.00). These findings together suggest that cover stories can affect the contingency judgments participants made, particularly if the cover stories are related to strong prior beliefs that participants may hold. As such the strong outcome density effects seen in prior studies, may be due to the lack of strong prior beliefs participants hold regarding those cover stories. Many of the stories used within contingency literature which uses passive paradigms use abstract scenarios (Crump et al., 2007; Kattner & Ellermeier, 2011) or the real life scenarios used such as growing plants (Moreno-Fernández et al., 2017; Wasserman et al., 1996) or fictitious medication treatment (Barberia et al., 2019; Blanco et al., 2013; Chow et al., 2019b; Klauer & Meiser, 2000; Lee & Lovibond, 2021; Matute et al., 2011b; Shanks & Lopez, 1996; Vadillo et al., 2011a; Yarritu & Matute, 2015). All of these are scenarios that participants may not care about or hold well-defined beliefs about. In addition to this, with the medical cover stories where patients are provided with some form of medication, participants may implicitly hold the belief that medications are more likely have some outcome than no outcome, based on life experiences, as such those contexts may results in more positive judgements even in zero contingencies.

8.3.5 Theoretical Evaluation

Contingency account

The contingency account does not explicitly consider these factors within the account. If the prior beliefs have developed from prior real-life experience of the four cell types of the contingency tables, then these may be taken into account when participants are making their contingency judgements. For example, if participants hold strong beliefs that social media use does not increase sociability, due to past experiences, this will count as cell B trials (cue-no outcome, with outcome being sociability). As cell B trials decrease positive judgements, this may reduce the overall positivity of the judgements. In addition to this, if past experiences are included within the contingency calculation, this changes the contingency from the pre-programmed zero contingency to a negative contingency. Table 51 demonstrates the effect of prior belief on Δ*P.*

I found that people with a prior belief favouring a positive association between marijuana use and health gave more positive contingency judgments for this association. Using the framework from Wasserman, this would mean that it is as if they gave higher weight to cell A and D relative to other cell counts. In contrast, people with a prior belief for the negative association between social media use and sociability gave more negative contingency judgments for this association. Using the framework from Wasserman, this would mean that it is as if they gave greater weight to cell C and B relative to other cell counts.

It is not always feasible to judge what people's prior experiences or beliefs are, nor the extent to which this may influence judgements. In the tasks participants are not explicitly asked to simply focus on the task at hand and to not incorporate trials or experience that could have occurred prior to the experiment in their judgement. Asking participants to ignore their prior experince or beliefs may be hard, as implicit biases are hard

Table 51 Influence of prior beliefs on contingency calculations

Outcome density at Zero Contingency

	Recovery		
Medication	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	4	16	P(O/~R) = .2
	ΔP = .2-.2=0		

Outcome density at Zero Contingency with past experiences accounted for

	Recovery		
Medication	Present	Absent	
Present	16	16	P(O/C) = .5
Absent	16	4	P(O/~c) = .8
	ΔP = .5-.8= -0.3		

to overcome, as such this may not be feasible. As such prior beliefs may always influence judgements.

Associative account

Similar to the previous reasoning given, if the prior beliefs have developed from prior real-life experience of cue and outcome occurrences, then these could carry over into the experimental paradigm. The model is based on the concept of associative history of cues (i.e., the prior training that they have received). As such participants may already have some level of associative strength between cue and outcome. The experimental paradigm could then increase or decrease the associative strengths. As such an effect of prior belief is possible to account for with the associative perspective.

Probabilistic account

Whilst the account does not explicitly consider the role of prior beliefs. If prior beliefs have developed from prior real-life experience with the cue and outcome occurrences, then it would be reasonable to assume that cue-outcome pairings observed in the experimental setting may provide further opportunity for learning

8.3.6 Limitations and future direction

The cover stories chosen for this experiment, were the options which participants on-average had more neutral beliefs on from the pilot. This was a choice made to avoid uniform belief scenarios. However, by only choosing scenarios that were more neutral, the effect of extreme prior beliefs cannot be examined in further detail. As such further research on the effect of cover stories which people may hold strong positive or negative prior beliefs on would be good direction to take this research further. A large outcome density effect is expected when most people have extreme positive beliefs. A smaller (maybe small, maybe zero, perhaps even reversed) outcome density effect is expected when people have extreme negative beliefs. Currently this study shows that strong negative prior belief can overcome the typically seen positive bias in a zero contingency high outcome density condition. It would be interesting to see if strong positive prior beliefs can overcome a programmed negative contingency, and to what extend this may occur. In addition to this a two-part study where the beliefs are elicited a week in advance, perhaps embedded in a series of other foil judgments could also be conducted.

However, in the wider context, this role of belief on judgements, and that prior belief can overcome programmed judgements may suggest that intervention which targets those with strong negative prior belief, may be resistant to changing their belief. As such this line of research could go much further. Other factors such as the source of belief, beliefs on associated topics, personal experience etc. could be examined. In addition to this, further exploration of how strong positive or negative prior beliefs influence judgements across zero, positive and negative judgements could be examined.

Experiment 6 found evidence of an effect of cover story. That is, the cover story used can affect the contingency judgements participants give, as such cover stories may lead to biases. An overall limitation of the methodology of this thesis is the use of the medication cover story across 6 of the 7 experiments. Participants may be slightly biased in that people tend to have an implicit belief that medications work or have some effect. This may explain the general positive shift effect seen across experiments. An important extension of this thesis was to include a 6th experiment introducing four novel cover stories. This allowed me to examine the possible impact of prior belief on contingency judgment in a more comprehensive manner that has previously been done (to the best of my knowledge). As such this experiment has shown that the cover story chosen for an experimental paradigm is important. However, by changing the cover stories, the wording of the scales and the contingency judgement question itself had to be changed, which could influence the judgments given by participants. However, it would be worthwhile to further explore cover stories, particularly cover stories which are more neutral in nature and to see if there are cover stories for which participants give judgement of zero contingency.

8.4 Conclusion

To summarise, the cover stories used within contingency paradigms can influence the contingency judgements given. In addition to this, pre-existing prior belief can also influence contingency judgements, depending on the cover story used. As such when designing experiments that may use a more "real-world" scenario, it is important to consider the influence of pe-existing prior beliefs. The outcome density effect is robust across all conditions, although the size of the effect varies. However, the size of the outcome density effect may be a function of the cover stories used.

Chapter 9: Affect

As I am interested in the real-world application of the contingency paradigm when assessing false belief formation, a key factor to consider is, *how does affect state influence judgements?*

In this chapter I will summarise the literature on the depressive realism effect. This effect suggests that estimates of control in contingency paradigms differ between depressed and non-depressed participants, such that depressed participants do not overestimate their estimates of control, whilst non-depression participants do. This is interesting as this could help identify the mechanism which leads people to overestimate control, in zero contingencies. Whilst all of the illusion of control work focusses on active tasks, it is unclear if these effects can be duplicated in passive tasks which are of interest for this thesis. And if they can be duplicated, then the theories used to explain the illusion of control, may shed light on the overestimate of contingency judgements seen in passive tasks. Several theories have been provided to explain this effect such as the motivational hypothesis, and learned helplessness, which provide explanations in terms of behaviours. Whilst Msetfi et al's (2013) work on the effect of ITI suggests a non-behavioural mechanism.

However, there are also many studies which have failed to replicate this effect. As such it is unclear if a) this is a robust effect b) whether this effect actually relates to depression or to another affective state. For example, anxiety and depression are highly morbid disorders (Mineka et al., 1998; Pine et al., 1998; Vos et al., 2017; Wittchen et al., 2011). Hence the depressive realism effect might not necessarily be about depression but may be driven by anxiety. Similarly stress and depression are comorbid (Wheatley, 1997), with long term stress thought to result in depression (van Praag, 2004).

Alternatively, depression may be linked to negative affectivity (Danhauer et al., 2013). Negative affect is thought to be related to depression, but does differ, as people can experience one and not the other. However negative affect tends to be highly correlated with depression. Rumination has been considered to be factor which links negative affect and depression (Iqbal & Dar, 2015). Hence negative affect may be the underlying affective state which leads to the depressive realism affect. Although research suggest that clinically depressed participants have lower positive affect scores compared to healthy controls (Suslow et al., 2019). As such the depressive realism effect may be due to lower or even a lack of positive affective state.

Those with depression have been found to have symptoms of suppressed or over anger (Busch, 2009; Busch et al., 2004; Sahu et al., 2014). Difficulties in controlling anger is considered by some to play a role in the onset and persistence of depression (Luutonen, 2007). And there have been positive associations found between levels of anger and hostility with depression (Riley et al., 1989).

Overall, it is unclear if the depressive realism effect can be replicated in passive contingency tasks. If it can be replicated, then the overestimation of contingency judgements seen in this thesis, and the literature, may be partially explainable by the theories which explain the overestimation of control in non-depressed participants. Also, it is uncertain if the depressive realism effect is actually due to depressive state, or another affective state that depression may be comorbid with. As such this chapter outlines the depressive realism literature, along with considerations of the other affective states. The depressive realism effect is then examined in a passive contingency task.

9.1 Depressive realism

One strong direction the contingency research has taken, which helps to answer this question, is towards the influence of depressive state on contingency judgements. Alloy and Abramson (1979) found that individuals with depression, had more accurate judgements of the relationship between their actions and outcomes compared to non-depressed individuals. Specifically, these researchers found that non-depressed individuals believed they had higher control over outcomes in the high outcome density compared to the low outcome density condition, although this effect varied depending on the particular experimental condition. For depressed individuals, however, judgments of control over an outcome were similar between high and low outcome densities. For each condition depressed participants accurately featured low perceived control, as such there was no evidence of the illusion of control in depressed participants. This effect is known as depressive realism (Alloy & Abramson, 1988).

Alloy and Abramson’s studies suggest that affect might shape how cues and outcomes are cognitively processed. Particularly, negative affect such as depression can allegedly remove the density effect. Alloy and Abramson (1979) explain these findings in terms of the modified learned helplessness hypothesis (Miller & Seligman, 1975). This is the concept that when people become exposed to stressful or aversive situations repeatedly, and the situation stops independent of their own actions, those events are classified as being uncontrollable. The learned helpless hypothesis also suggests there are motivational, cognitive, and associative deficits. The motivational deficit in the hypothesis suggests that people have less motivation to initiate responses, whilst the cognitive or associative deficit refers to the difficulty in learning future response-outcome contingencies. As such they learn to not even try changing such situations, hence “learned helplessness”. In addition to this people who consistently learn that responses and outcomes are non-contingently related become depressed. Altogether the learned helplessness hypothesis suggest that people with depression have poor performance on instrumental tasks, partly due to an associative or cognitive, and motivational deficits. The learned helplessness hypothesis then provides the prediction that those with depression will underestimate their degree of control over the outcomes, whilst non-depressed participants will overestimate their degree of control.

However, this does not explain the Alloy and Abramson findings. While depressed participants were accurate in their ratings of the contingency, the non-depressed participants did not show consistent illusions of control across all experimental conditions. As such Alloy and Abramson proposed that those with depression are characterised by not having expectations of control, and this only interferes with the initiation of responses which links to the motivational deficit of learned helplessness. However, the lack of expectation in control does not interfere with the perception of the response-outcome relationship which links to the associative deficit in helplessness. As such the modified learned helplessness hypothesis suggests that people with depression perform worse on instrumental tasks than non-depressed people, as they fail to form the responses that increase the probability of an outcome, rather than due to those with depression being unable to differentiate between the effect their responses have on the outcome.

Alloy and Abramson also provide an alternative explanation for the findings based within social psychology, in a form of misattribution as people generally attribute causality to self when succeeding in a task, and causality to the environment or situation when failing in a task (Feather, 1969; Fitch, 1970; Streufert & Streufert, 1969; Weiner et al., 1971; Wortman et al., 1973). With these findings it is suggested that people are using motivation to enhance their self-esteem by attributing success to self, and by attributing failures to the environment, this avoids lowering or damaging their self-esteem (Bradley, 1978; Miller, 1978; Miller & Ross, 1975). This reasoning is applied to contingency judgements, such that non-depressed participants did show the illusion of control when outcomes occurred, however did not show the effect when the outcomes did not occur to avoid lowering their self-esteem. As such depressed participants who believe they cannot control the lack of outcomes must then believe this lack of control is due to external factors rather than internal factors. If depressed participants are assumed to not be motivated to care about their self-esteem, their accuracy of contingencies follows as they do not care about maintaining or enhancing their self-esteem, whether the outcome occurs or not, or if the outcome is bad or good. This suggests that depression is often linked to low self-esteem (Sowislo & Orth, 2013). As such Alloy and Abramson present a motivational hypothesis that differential motivation for maintaining self-esteem may explain the difference in findings between depressed and non-depressed participants.

9.1.1 Further support and alternative explanations

Msetfi et al. (2005) further explored depressive realism and intertrial interval (ITI) in two experiments using action-outcome contingency judgment tasks. They used a 2 (mood: depressed vs non-depressed) × 2 (ITI length: short vs long) × 2 (possibility of outcome in ITI:

presence of lightbulb vs blank screen) × 2 (gender) fully between-subjects design. One hundred and twenty-six participants judged the extent to which their pressing of a button controlled the onset of a light, on a 0 to 100 scale, after 40 trials. They used a high outcome density zero contingency condition where the probability of the onset of the light was .75 , regardless of whether the button was pressed. Depressed participants did not differ in their judgements of control between the short and long ITI conditions, whereas non-depressed participant gave significantly higher judgements in the long ITI condition compared to the short ITI condition.

In the second experiment Msetfi et al. (2005) used a 2 (mood) × 2 (ITI length) × 2 (outcome density) × 2 (gender) fully between subjects design to further explore the ITI effect. An additional condition of low (0.25) vs high outcome (0.75) density was added, as previously only a high-density condition was used. Ninety-six participants provided judgements of control. The results suggested that there were no effects of ITI in the low outcome density condition for any group. However, in the high outcome density condition an effect of ITI was found. Non-depressed participants significantly differed in their judgements between short and long ITI conditions, whilst the depressed participants did not. In addition to this, the results also suggested that when ITI was short there was no outcome density effect, whereas when the ITI was long there were outcome density effects.

Overall, the two experiment together suggest that the depressive realism effect is a function of outcome density and the intertrial interval length (ITI). The depressive realism effect was only found when participants were in the long ITI condition, that is they were exposed to long periods of no action-no outcome trials within the experimental context. When these ITI or these no action-no outcome trials are included in the contingency calculation it can change the overall experienced contingency of participants in such a way that is very different to the programmed contingency. Table 52 shows that when extra no action-no outcome trials are added, which are in effect ITis, the probability of outcome occurring in the absence of cue, decreases. Overall, this increases the contingency. Msetfi et al., suggest that non-depressed participants take into account ITIs when giving contingency judgements, as such showing increased judgments of control which more accurately reflects the experienced contingency when considering the ITIs. Meanwhile depressed participants do not take into account the ITI to the same extent. As such depressive realism may be due to reduced contextual processing instead of increased accuracy of contingency judgement or negative expectation. Hence Msetfi et al. (2005) provides an alternative explanation for the depressive realism effect and suggests that the effect is consistent with the cognitive distortion view of depression. They argue that symptoms of depression include cognitive symptoms of difficulties in concentration and attention (American Psychiatric Association, 2022). The ITI results in extended waiting periods which could lead to problems in maintaining attention on the task. As such the contextual information of the trial may not be processed accurately by depressed participants compared to non-depressed participants, overall leading to the depressive realism effect.

Table 52 Zero contingency calculations with and without intertrial intervals

Zero Contingency with no intertrial interval

	Outcome		
Action	Present	Absent	
Present	15	5	P(O/C)=.75
Absent	15	5	P(O/~C)=.75

ΔP=.75-.75=0

Zero Contingency with intertrial interval

	Outcome		
Action	Present	Absent	
Present	15	5	P(O/C)=.75
Absent	15	5+5	P(O/~C)=.60

ΔP=.75-.60=0.15

Further support for the role of ITIs in contingency judgements comes from Msetfi et al's (2007) experiment which investigated the influence of ITIs on judgements across zero, positive and negative contingencies, between depressed and non-depressed participants. In a 2 (contingency: zero, positive) by 2 (length of intertrial interval: short, long) by 2 (mood: nondepressed, depressed) by 2 (sex: female, male) fully between-subjects design 96 participants provided judgements of control. They found a significant main effect of contingency; participants gave higher judgements in the positive contingency compared to the zero contingency. Similarly, an interaction effect of ITI by mood indicated that when ITI was long non-depressed participants gave higher judgements in both zero and positive contingencies compared to depressed participants. This replicated previous research and suggests that the ITI effect with depressive realism extends to positive contingencies as well.

In a follow up experiment with 48 participants, the same effect was examined in the negative contingency condition (Msetfi et al., 2007). An interaction effect of mood by ITI was found, suggesting that non-depressed participants gave higher judgements of control in the long ITI condition compared to the short ITI. The depressed participants gave more negative judgements in the long ITI condition.

9.1.2 Non-replication of the Depressive Realism Effect

Whilst the above studies have shown the depressive realism effect and the various conditions under which it may occur, many studies have failed to replicate this effect (Bryson et al., 1984; Dobson & Pusch, 1995; Kapcli & Cramer, 1999). Dev et al. (2022) attempted to replicate the original Alloy and Abramson study to see if the depressive realism effect could be found, across two experiments using a similar study design. The study used a 4 (inventory to diagnose depression-current (IDD-C)) x 3 (contingency condition: high outcome density in zero contingency, low outcome density in zero contingency, high outcome density in positive contingency) x 2 (gender) between subjects' design. The IDD-C was used to split participants into 4 groups based on minimal, mild, moderate, and severe depressive symptoms.

On hundred thirty-six undergraduate students and 246 participants were recruited from Mturk who completed an adapted version of the Alloy and Abramson (1979) task. Across 40 trials, participants could choose to press a button after which a lightbulb may appear. After every 10 trials participants provided judgements of control on a 0 to 100 scale, half the participants were asked about the lightbulb appearing, whilst the other half were asked about the lightbulb not appearing. In addition to this, participants were randomly assigned to a positive contingency, a zero contingency with low outcome density, or a zero contingency with a high outcome density condition. Participants were pre-screened according to the Inventory to Diagnose Depression – Current, to give four levels of depression severity. As the original depressive realism effect was found in women, gender was also included as a factor. Compared to the original Alloy and Abramson (1979) task the positive contingency condition was added, and participants depressive traits were measured using the IDD-C.

Main effects of contingency were found, such that participants' judgements significantly differed between contingency conditions. A main effect of depression was found; however, the results suggest that increased depression severity results in overestimation of judgements of control compared to lower levels of depression. This contradicts the earlier findings on the depressive realism effect.

9.2 Anxiety

Anxiety and depression are highly comorbid disorders (Mineka et al., 1998; Pine et al., 1998; Vos et al., 2017; Wittchen et al., 2011). Hence there are concerns that the depressive realism effect might not necessarily be about depression, but may be driven by anxiety, which is supported by the Dev et al. (2022) study. As such to clarify this effect, further examination of anxiety is needed.

Dev et al. (2022) also collected anxiety measures from the participants, in particular anxiety was measured using the anhedonic depression and anxious arousal subscales of the mood and anxiety symptoms questionnaire (MASQ) (Wardenaar et al., 2010). Greater anxious arousal led to increased control bias and had a significant interaction with contingency type. The control bias was greatest in the zero contingency conditions compared to the positive contingency. Within the zero contingency there was greater control bias with the high outcome density condition. This suggest that anxiety may influence the illusion of control, with the greatest effect under zero contingency conditions.

Obsessive compulsive disorder (OCD) is thought to be characterised by anxiety (American Psychiatric Association, 2022). OCD is a condition whereby patients suffer from obsessive thoughts which are distressing, such as harm befalling loved ones, which results in repeated actions they feel compelled to act on (American Psychiatric Association, 2022). There is debate regarding the reasoning for the compulsions in OCD, with some claiming these are attempts to reduce anxiety and gain control over the threats (Carr, 1974; Foa & Kozak, 1986; McFall & Wollersheim, 1979) due to a lack of perceived control in their life (McLaren & Crowe, 2003; Moulding & Kyrios, 2006). However others suggest that compulsions are due to the belief of those with OCD that they have excessive control over or sense of power (Salkovskis et al., 1998).

Gillan et al. (2014) investigated the role of perceived control in patients with obsessive -compulsive disorder (OCD) using an illusion of control paradigm, used by Alloy and Abramson (1979). Twenty-six control participants were matched to 26 patients with OCD, who completed the illusion of control paradigm, which included both high and low outcome density conditions. Depression measures were also collected for all participants.

Main effects of OCD condition and outcome density were found however there was no interaction effect of OCD condition by outcome density. Participants gave higher judgements of control in the high outcome density condition compared to the low-density condition. Also, participants with OCD provided lower estimates of control in both the low and high outcome density conditions, than the control participants. There was no effect of depression on control. The results of the Gillan et al. (2014) study suggest that OCD patients have reduced sense of control; however, the study also found no relationship between the illusion of control and depressive symptoms. A limitation highlighted in this study, was the lack of comparison to patients with generalised anxiety disorder (GAD) who present with positive symptoms, as this could help to see whether the reduced control seen in OCD patients is due to anxiety, or the other features of OCD. As such anxiety is another affective state that is of interest to the current study.

9.3 Positive and negative affect

Negative affect and depression are thought to be related, however people can experience one and not the other (Danhauer et al., 2013). Negative affect tends to be highly correlated with depression. Research suggest that clinically depressed participants have lower positive affect scores compared to health controls (Suslow et al., 2019). As such the depressive realism effect may be due to lower or even a lack of positive affective state.

Novovic et al. (2012) conducted an experiment on the influence of positive and negative affect on the illusion of control on fifty-four participants. Both positive and negative affect was measured at three-time points; two weeks prior to the task, immediately before and after the illusion of control task. They found that all three measurements of the positive affect schedule were significantly correlated with the increased illusion of control, whilst none of the negative affect subscales were correlated. Overall, this study provides evidence for positive affect influencing contingency judgements in active paradigms. This study also suggests that negative affect is not relevant in the prediction of control over uncontrollable situations, which may seem contrary to the evidence shown thus far on depressive realism. The hierarchical model by Watson provides a possible explanation for these results (Watson, 2005). The model suggest that depression consists of a mixture of symptoms from both negative and positive affect. From negative affect, depression includes states of high distress and high negative emotions, whilst from positive affect states of low joy of life, general loss of motivation and enthusiasm, and ability to enjoy every day routines are characteristic of depression. As such the results of the Novovic et al. (2012) suggest perhaps that depression is acting on the illusion of control effect, through low positive affect, rather than through high negative affect. As such low affect or the absence of depression increases the illusion of control, due to higher positive affect.

Overall, this study suggests that participants who are high in positive affect may be more susceptible to believe that a relationship exists between cause and effect. This indicates that the depressive realism effect is due to low positive affect, rather than negative affect. Interestingly the study suggests that low positive affect leads to accurate judgements, whilst increased positive affect leads to overestimation of judgements. Positive affect have been found to lead to an increased belief that causes and effects are linked compared to negative affect (Clore & Huntsinger, 2007).

9.4 Anger

Anger as an emotional state is known to influence people's perceptions and formation of their decision. In addition to this it may guide their behaviour whilst they are angry. Those with depression have been found to have symptoms of suppressed or over anger (Busch, 2009; Busch et al., 2004; Sahu et al., 2014). Difficulties in controlling anger is considered by some to play a role in the onset and persistence of depression (Luutonen, 2007). And there have been positive associations found between levels of anger and hostility with depression (Riley et al., 1989).

Anger makes people optimistic about their own chances of success (Fischhoff et al., 2005) in risky situations in comparison to neutral or fearful emotional states. When participants were presented with Kahneman and Tversky's disease problem (Tversky & Kahneman, 1981) and required to choose between a gamble or a sure option, regardless of the framing those with an angry disposition tended to make risk seeking choices whilst those with a fearful disposition chose risk averse choices (Lerner & Keltner, 2001). Moreover, angry participants are more likely to make stereotypic judgements in comparison to participants who are more neutral or ‘sad’ (Bodenhausen et al., 1994). Similarly, participants who scored highly on anger measures had a greater reliance on heuristic cues. Intuitively it may be thought that anger as a negative emotional state would lead to more pessimistic choices and expectation, however the literature suggests that anger as an emotional state induces optimism and risk seeking. Increased anger may lead to increased relative sensitivity to irrelevant contextual cues, as there may be greater reliance on heuristic cues. As such during the high-density conditions although there is no relationship between the cue and outcome, the increased frequency of cue-outcome pairing may bias judgements. Anger is considered to be a negative emotion however it could lead to greater outcome density bias, this is an interesting affect state to explore further.

9.5 Summary

To summarise, the research within the contingency literature indicates that depression may influence biases in learning related to false beliefs (e.g., Alloy & Abramson, 1979). To date, most of this work has focused on the impact of negative mood (Msetfi et al., 2005) compared to neutral mood, and how this may influence contingency judgements. However, the influence of other forms of affect, such as anxiety, or anger, or positive affect have been studied less. Research from other areas of cognitive psychology indicates that other forms of affect can influence judgement in general. Overall, this suggests that contingency judgements may vary as a function of individual differences in affective state.

In addition to this, the research within contingency judgements has focused on the illusion of control, an active form of a contingency paradigm, which does not fully translate to the passive contingency paradigm of interest for this thesis. In a book chapter Alloy et al.

(1985) described three experiments which used passive contingency paradigms to explore the depressive realism effect. Across these studies by Alloy et al. (1985) no depressive realism effects were found, however these studies were only briefly presented in the chapter and the details on the data and method have not been provided comprehensively in journal articles.

The question is thus how well findings from studies that operationalized negative affect as depression generalize to other forms of negative affect. To assess the impact of affect on judgements I ran a well powered study which tested whether individual differences in dispositional affect shapes judgment. Specifically, testing that people who feel relatively affectively positive will be more prone to believing in false beliefs due to their relative insensitivity to contextual cues. In contrast, those who have a negative affect (e.g., depression, anxiety, stress, and negative affect) will be less prone to believing false beliefs. In this study I propose to assess the generality of the influence of affect on the formation of false beliefs. To this end, the study used a series of validated affect measures, including the Depression Anxiety Stress Scales (Lovibond et al., 1995), Positive and Negative Affect Schedule—expanded form (Watson & Clark, 1994), and Buss-Perry Aggression Questionnaire (Buss & Perry, 1992).

9.6 The current study

The contingency literature suggests that affect, such as depression may influence biases in learning, particularly that depressed participants provide a more accurate judgement of contingencies compared to non-depressed participants (Alloy & Abramson, 1979; Msetfi et al., 2005). In addition to this research also suggests that positive affect may exaggerate the contingency judgements participants provide. (Novovic et al., 2012). As such this experiment investigated the hypothesis that affect influences judgement in zero contingency, however the research in this area has focused on the influence of affect in active contingency paradigm, the illusion of control, rather than the passive contingency tasks.

This experiment had two aims. Firstly, to see if the depressive realism effect could be replicated, along with the effect of positive affect on contingency judgment, using passive contingency tasks. Secondly this experiment aims to explore how other types of affect influences' the outcome density effect as the focus of the literature has been mainly on depression, with one study on positive and negative affect. The hypothesis tested is that people who feel affectively relatively positive will be more prone to believing in false beliefs due to their relative insensitivity to contextual cues was tested. In contrast, those who have a negative affect (e.g., depression, anger, anxiety, negative affect) will be less prone to believing false beliefs. In this study I propose to assess the generality of the influence of affect on the formation of false beliefs. The main goal of this chapter is to examine whether the outcome density effect is robust. In addition to this, the chapter examines whether factors (affect) beyond methodological variations influence the outcome density effect.

9.7 Methods

9.7.1 Design

Judgement of effectiveness of medication in treating illness was the dependent variable. Outcome density (within-ss: high vs. low density) was varied in the experimental design. And natural differences in affect were also measured as another predictor variable.

9.7.2 Materials

The same medical treatment paradigm outlined in chapter 2 was used. Contingency judgement in one high and one low outcome density conditions was collected for all participants. The affect measures outlined below were used as the predictor variables and were self-report questionnaires.

Depression, Anxiety, and Stress

Depression, anxiety and stress were measured using the depression anxiety stress scales (DASS) (Lovibond et al., 1995) scale. The DASS is a 42-item self-report instrument designed to measure the three related negative affect states of depression, anxiety and tension/stress. The DASS uses a 4-point scale, ranging from 0=did not apply to me at all, to 3=applied to me very much, or most of the time over the past week. There are 14 items for each subscale, which are summed to give a maximum score of 42 per subscale. Both the Beck's depression inventory (BDI) (Beck et al., 1961) and the DASS has been used across the affect contingency literature (Cavus & Msetfi, 2016; Msetfi et al., 2015). The DASS depression scale shows correlations of 0.74 with the BDI, and the DASS anxiety scale shows correlations of 0.81 with the Beck Anxiety Inventory (BAI) (Lovibond & Lovibond, 1995). The BDI differs from the depression subscale of the DASS, as the BDI includes items such as insomnia and weight loss, which are not measures of affect. The DASS has good internal consistency (α=0.89) (Akin & Çetın, 2007). The DASS was chosen over the BDI as the DASS also includes other negative states of stress and anxiety, which are of interest to this study, and allows more measures to be tested with fewer items, to reduce participant fatigue effects.

Positive and negative affect

The PANAS-X is 20 Item scale measuring affect with 10-item subscales, each measuring negative and positive affect (Watson & Clark, 1994). The PANAS uses a 5-point scale ranging from 1=very slightly or not all to 5=extremely. The scores for the 10 items in each subscale are summed for overall positive and negative affect scores, which gives a maximum score of 50 for either subscale. The subscales of positive and negative have good internal consistency (α=0.90, α=0.91 respectively), and good levels of test-retest reliability (Positive Affect=0.80, Negative Affect=0.76) (Serafini et al., 2016).

Anger and hostility

Anger and hostility was measured using the Buss-Perry Aggression Questionnaire (BPAQ) (Buss & Perry, 1992). The BPAQ is a 29-item self-report questionnaire designed to measure four factors of aggression: anger, hostility, verbal aggression, and physical aggression. The BPAQ uses a 5-point scale, ranging from 1=extremely uncharacteristic of me to 5=extremely characteristic of me. The measures for verbal and physical aggression were not included in the analyses as these are not forms of affect. Anger has 7 items whilst hostility consist of 8 items, which are summed to give a maximum score of 35 and 40 respectively. The subscales of anger and hostility show internal consistency of α=0.73-0.64, α=0.73-0.78, respectively (Hornsveld et al., 2009).

The following scales were also included in the study however are not part of the main analyses for the purposes of this thesis. The Santa Clara Strength of Religious Faith Questionnaire (SCSORF) (Plante & Boccaccini, 1997). The SCSORF is a 10-item self-report questionnaire designed to assess strength in religious faith. The gratitude questionnaire

(McCullough et al., 2002) is a 6-item self-report scale measuring gratitude. The Rosenberg self-esteem scale (RSES) (Rosenberg, 2006), is a 10-item self-report questionnaire measuring global self-worth by measuring both positive and negative feelings about the self.

9.7.3 Procedure

Participants were recruited through crowdsourcing platform *Prolific.co* and the Department of Psychology research participation pool at King's College London (SONA). Upon sign-up, participants were given a link to the study. Participants logged into the experiment, read the information sheet, provided consent to participate, and answered demographic questions (gender, age, ethnicity, and ability to communicate in English). General instructions followed. Participants then answered the questionnaires on affect. Following this participants completed the contingency task used previously. However, participants only took part in one high and one low density condition, both of which were completed in the same session. The task took on average 24.5 minutes to complete, and at the end of the task a short debrief was given. Participants were given either participation credits or £3.13.

9.7.4 Participants

An a priori power analysis was conducted using G*power (Faul et al., 2007) to determine the minimum sample size required to test the study hypothesis. Results indicated the required sample size to achieve 80% power for detecting a small effect of 0.1 at a significance criterion of $\alpha=.05$, is $N=151$ for multiple linear regression with 7 predictors. I aimed for 170 participants to allow for participant withdrawal and errors. The obtained sample size of $N=189$ is adequate to test the study hypothesis. More participants were recruited than the expected sample size due to the recruitment methods utilised. The credit system used allowed for more participants to complete the study by the cut off before researchers could end the study.

Twenty participants were excluded from analysis due to completing the experiment under 20 minutes. The task itself, without the questionnaire items takes a minimum of 16 minutes to complete due to the programmed trials, as such completing the experiment is under 20 minutes is considered too quick to have fully read each questionnaire item and answered appropriately. A total of 169 participants were included in the analysis. Of the 169 participants 118 were female, with the modal age group between 18 to 20, and the age groups ranges between 18 and 20 to 71+.

9.7.5 Data analyses

One item on the BPAQ was reversed coded for the anger subscale. The items in each subscale were summed together to create the overall measure for each subscale, a higher score on each subscale measure indicates individuals are highly characteristic of that measure. To see whether affect predicts the outcome density effects, the difference between the low and high outcome density judgements were calculated and used as an outcome variable. Correlations were run across all affect measures with the outcome density measures. To test the effect of affect on judgements, the predictors were split into three groups, and each group was regressed with the three outcome variables (high, low and difference in outcome density). As such a total of nine regression models were run. The first three models entered depression as a predictor variable, and then entered anxiety and stress. The second three models entered positive and negative affect as the predictors. The final three models entered anger and hostility as the predictors.

9.8 Results

9.8.1 Descriptives

See Table 53 and Figure 23 for the means and standard deviation for each affect measure. The internal reliability of all the scales were high, and the predictor variables are non-normally distributed.

High outcome density has a significant positive correlation with anger and stress. Low outcome density has a significant positive correlation with stress, anxiety, and positive affect. The difference in outcome density has a significant positive correlation with anger. Table 54 shows the correlations across variables.

A paired samples *t*-test was conducted to determine the effect of outcome density. The results indicated a significant difference between the high outcome density (*M*=45.98,

Table 53 The means, standard deviations, Cronbach's α of the subscales

Participants (*n*=169)	Mean (*SD*)	Actual range	Possible range	Items per subscale	Cronbach's α
BPAQ Anger	17.07 *(5.58)*	7 - 31	7 – 35	7	.832
BPAQ Hostility	23.53 *(6.23)*	8 - 37	8 - 40	8	.795
DASS Anxiety	7.52 *(7.22)*	0 -37	0 - 42	14	.900
DASS Depression	12.97 (10.16)	0 – 42	0 - 42	14	.955
DASS Stress	14.04 *(9.29)*	0 -42	0 - 42	14	.934
PANAS Negative	22.67 *(8.14)*	10 - 46	10 -50	10	.892
PANAS Positive	29.57 *(7.68)*	10 – 48	10 - 50	10	.889

SD=28.59) and low outcome density conditions *(M*=18.25, *SD*=16.38). The 95% confidence interval of the difference between the means ranged from [23.89 to 31.58] and did indicate a difference between the means of the samples. The effect size was *d*=1.19.

9.8.2 Depression, Anxiety and Stress

Multiple regressions were run to predict contingency judgement (low, high, difference) for depression, stress, and anxiety, see Table 55 for full regression models. The first model entered depression as a predictor variable, and then entered anxiety and stress.

Figure 23 Affect scores with interquartile ranges and median

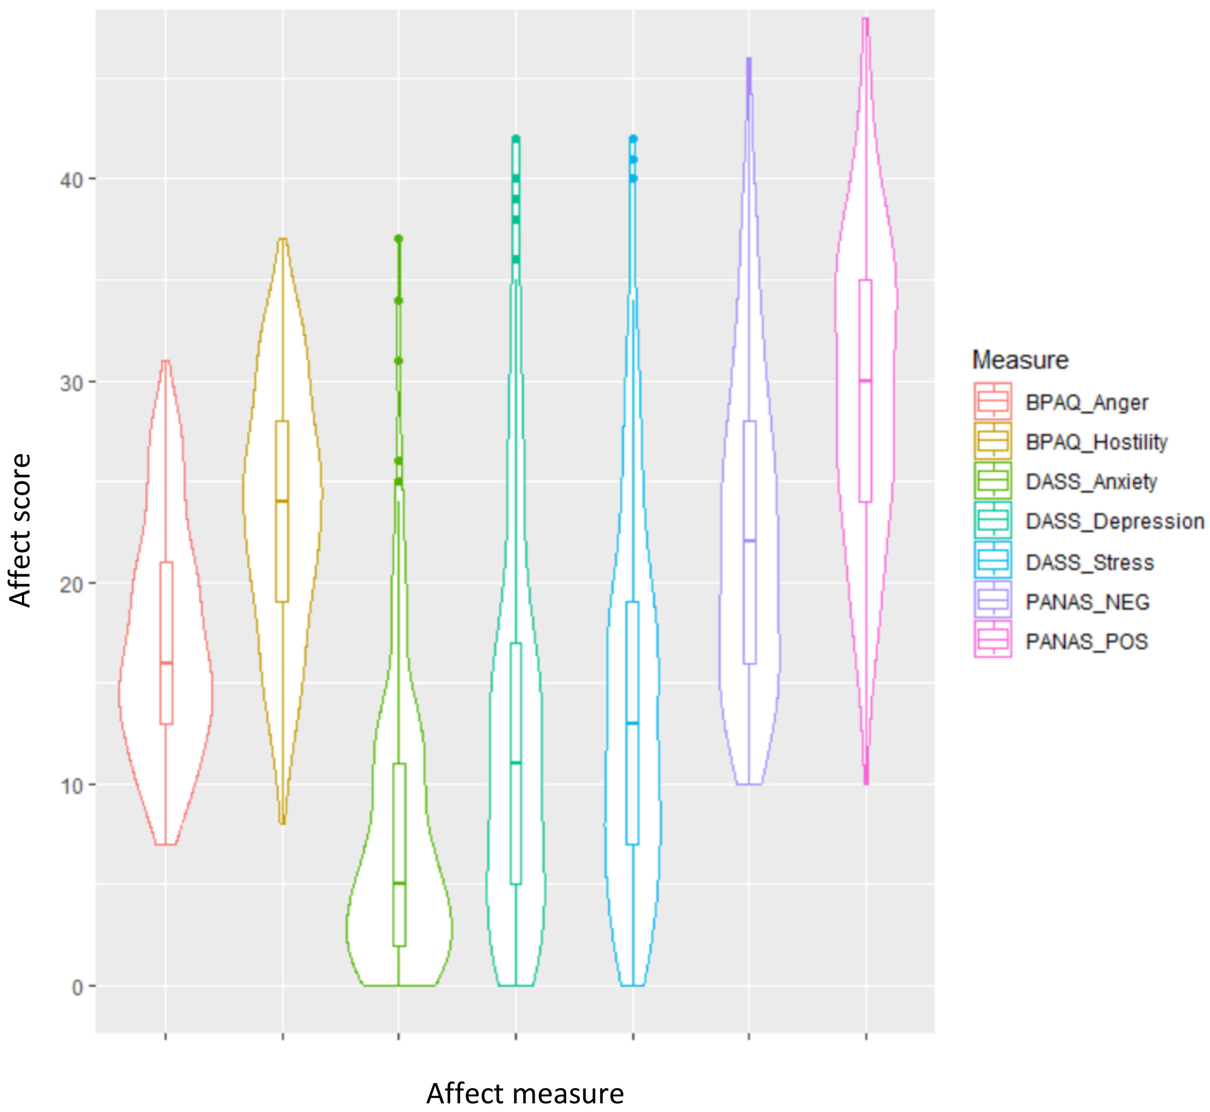

For low outcome density judgements the first regression model was not significant, *F*(1, 167)=0.59, *p*=.443, adjusted R^2=-.002. The second regression model with stress and anxiety added was significant, *F*(3, 165)=3.09, *p*=.029, adjusted R^2=.036. The change in R^2 was significant, *F*(2, 165)=4.32, *p*=.015. The coefficient of anxiety was a significant predictor of judgement in the low outcome density condition, the higher participants scored on the anxiety measure, the higher the contingency judgement given in the low outcome density

Table 54 Correlations for the affect variables and the outcome density conditions

	BPAQ Anger	BPAQ Hostility	DASS Stress	DASS depression	DASS anxiety	PANAS Positive	PANAS Negative
High outcome density	.20**	.09	.15**	.05	.09	.06	.08
Low outcome density	.08	.00	.15**	.06	.20**	.16**	.15
Outcome density difference	.17**	.01	.08	.02	-0.03	-.03	0.00
BPAQ Anger		.53**	.45**	.29**	.29**	-.20**	.40**
BPAQ Hostility			.48**	.46**	.41**	-.39**	.49**
DASS Stress				.75**	.79**	-.21**	.69**
DASS depression					.64**	-.51**	.60**
DASS anxiety						-.21**	.75**
PANAS Positive							-.17*

**. Correlation is significant at the 0.01 level (2-tailed).

*. Correlation is significant at the 0.05 level (2-tailed).

condition. Depression and stress were not significant predictors of contingency judgements in the low outcome density condition.

Table 55 Regression models for depression, anxiety and stress predicting contingency judgements

Table a Multiple regression for low outcome density judgements predicted by depression, stress, and anxiety

Independent variables	R^2	ΔR^2	R^2 change	B	95% CI	*β*	*t*	*p*
Model 1	.004	-.002	.004					
Depression				0.09	[-0.14, 0.33]	0.06	0.77	.443
Model 2	.053	.036	.050*					
Depression				-0.22	[-0.57, 0.13]	-0.14	-1.23	.221
Stress				0.08	[-0.39, 0.55]	0.05	0.33	.745
Anxiety				0.57	[0.06, 1.09]	0.26	2.19	.030**

Table b Multiple regression for high outcome density judgements predicted by depression, stress, and anxiety

Independent variables	R^2	ΔR^2	R^2 change	B	95% CI	*β*	*t*	*p*
Model 1	.003	-.003	.003					
Depression				0.15	[-0.28, 0.58]	.054	.694	.489
Model 2	.034	.017	.032					
Depression				-0.38	[-1.02, 0.27]	-0.13	-1.15	.252
Stress				0.93	[0.07, 1.80]	0.30	2.14	.034**
Anxiety				-0.25	[-1.20, 0.70]	-0.06	-0.51	.609

Table c Multiple regression for the difference in outcome density judgements predicted by depression, stress, and anxiety

Independent variables	R^2	ΔR^2	R^2 change	B	95% CI	*β*	*t*	*p*
Model 1	.000	-.005	.000					
Depression				.59	[-0.35, 0.47]	.02	.27	.775
Model 2	.029	.011	.028					
Depression				-0.16	[-0.77, 0.46]	-0.06	-0.51	.612
Stress				0.86	[0.04, 1.67]	0.29	2.07	.040**
Anxiety				-0.82	[-1.72, 0.08]	-0.22	-1.80	.074

*R^2 change significant at the 0.05 level
** coefficient significant at the 0.05 level

For high outcome density judgements, the first regression model was not significant, $F(1, 167)=0.48$, $p=.489$, adjusted $R^2=-.003$. The second regression model with stress and anxiety added was also not significant, $F(3, 165)=1.96$, $p=.122$, adjusted $R^2=.017$. The change in R^2 was not significant, $F(2, 165)=2.70$, $p=.070$. The coefficient of stress was a significant predictor of judgement in the high outcome density condition, the higher participants scored on the stress measure, the higher the contingency judgement given in the high outcome density condition. Depression and anxiety were not significant predictors of contingency judgements in the high outcome density condition.

For the difference between outcome density judgements the first regression model was not significant, $F(1, 167)=0.78$, $p=.775$, adjusted $R^2=-.005$. The second regression model with stress and anxiety added was not significant, $F(3, 165)=1.64$, $p=.182$, adjusted $R^2=.011$. The change in R^2 was not significant, $F(2, 165)=2.42$, $p=.092$. The coefficient of stress was a significant predictor of difference between judgements in the high and low outcome density conditions. The higher participants scored on the stress measure, the larger the difference between judgements in the conditions. Depression and anxiety were not significant predictors of the difference in contingency judgements.

Overall, these models have failed to find evidence of depression predicting contingency judgements. Evidence was found for stress predicting contingency judgements in the high-density condition, and the difference between conditions. Evidence for an effect of anxiety was found for predicting judgements in the low-density condition.

9.8.3 Positive and Negative affect

Multiple regressions were run to predict contingency judgement (low, high, difference) for positive affect, and negative affect, see Table 56 for full regression models. The model entered positive affect and negative affect as the predictors.

For low outcome density judgements, the regression model was significant, $F(2, 166)=4.93$, $p=.008$, adjusted $R^2=.045$. The coefficient of positive affect was a significant predictor of judgement in the low outcome density condition, the higher participants scored on the positive affect measure, the higher the contingency judgement given in the low outcome density condition. The coefficient of negative affect was a significant predictor of

Table 56 Regression models for positive and negative affect predicting contingency judgements

Table a Multiple regression for low outcome density judgements predicted by positive and negative affect

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	0.56	0.45					
Positive affect			0.39	[0.08, 0.70]	0.19	2.46	0.015**
Negative affect			0.35	[0.05, 0.64]	0.18	2.34	0.020**

Table b Multiple regression for high outcome density judgements predicted by positive and negative affect

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	.012	.012					
Positive affect			.284	[-0.30, 0.86]	.08	0.97	.332
Negative affect			.322	[-0.22, 0.87]	.09	1.17	.243

Table c Multiple regression for the difference in outcome density judgements predicted by positive and negative affect

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	.001	-.011					
Positive affect			-.10	[-0.65, 0.45]	-0.03	-0.37	.714
Negative affect			-.02	[-0.54, 0.49]	-0.01	-0.09	.926

** coefficient significant at the 0.05 level

judgement in the low outcome density condition, the higher participants scored on the negative affect measure, the higher the contingency judgement given in the low outcome density condition.

For high outcome density judgements, the regression model was not significant, $F(2, 166)=1.00$, $p=.372$, adjusted $R^2=.000$. Neither positive nor negative affect were significant predictors of judgements in the high-density condition.

For the difference between outcome density judgements the regression model with negative affect added was not significant, $F(1, 166)=1.64$, $p=.934$, adjusted $R^2=-.011$. Neither positive nor negative affect were significant predictors of the difference in judgements between the high and low outcome density conditions.

Overall, evidence was found for positive affect and negative affect predicting contingency judgements in the low outcome density condition.

9.8.4 Anger and Hostility

Multiple regressions were run to predict contingency judgement (low, high, difference) for anger, and hostility, see Table 57 for full regression models. The model entered anger and hostility as the predictors.

For low outcome density judgements, the regression model was not significant, $F(2, 166)=0.77$, $p=.467$, adjusted $R^2=-.003$. Neither anger nor hostility predicted judgements in the low outcome density condition.

For high outcome density judgements, the regression model with hostility added was significant, $F(2, 166)=3.49$, $p=.033$, adjusted $R^2=.029$. The coefficient of anger was a significant predictor of judgement in the high outcome density condition. The higher participants scored on the anger measure, the higher the contingency judgement. Hostility was not a significant predictor.

For the difference between judgements in the density conditions the regression model was not significant, $F(2, 166)=2.34$, $p=.100$, adjusted $R^2=.016$. Anger and hostility were not significant predictors.

Table 57 Regression models for anger and hostility affect predicting contingency judgements

Table a Multiple regression for low outcome density judgements predicted by anger and hostility

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	.009	-.003					
Anger			0.32	[-0.2, 0.83]	0.11	1.24	.218
Hostility			-0.15	[-0.61, 0.30]	-0.06	-0.66	.509

Table b Multiple regression for high outcome density judgements predicted by anger and hostility

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	.040	.029					
Anger			1.08	[0.17, 1.98]	0.21	2.34	.021**
Hostility			-.081	[-0.90, 0.73]	-0.02	-0.20	.845

Table c Multiple regression for the difference in outcome density judgements predicted by anger and hostility

Independent variables	R^2	ΔR^2	B	95% CI	*β*	*t*	*p*
Model	.027	.016					
Anger			.76	[-0.11, 1.62]	0.16	1.73	.086
Hostility			.07	[-0.70, 0.85]	0.02	0.18	.855

** Coefficient significant at the 0.05 level

Overall, evidence was found for anger predicting contingency judgements in the high outcome density condition. No evidence was found that hostility predicted contingency judgements.

9.9 Discussion

The study used a learning paradigm to investigate whether affective states relate to the development of false beliefs. I hypothesised that people who feel affectively positive (e.g., happy) will be more prone to believing in false beliefs, and those who score higher on the depression measure will provide more accurate contingency judgements. The study found partial support for these hypotheses. In the low outcome density condition judgements were predicted by anxiety, positive affect, and negative affect. In the high outcome density condition judgements were predicted by stress and anger. The difference in judgements was predicted by stress. Neither depression nor hostility predicted any of the outcome variables, as such this study failed to find evidence of the depressive realism affect in passive contingency paradigms. Overall the outcome density effect was a robust effect.

9.9.1 Outcome density

The outcome density effect was found in this experiment, replicating the previous experiments of this thesis. Msetfi et al. (2005) found evidence that there is no outcome density bias in short ITI conditions, replicating the work of Allan and Jenkins (1980) and Wasserman et al. (1993), and that long ITI is required for an outcome density effect. ITIs are fairly well research mechanisms in the active contingency literature, particularly in relation to depression (Msetfi et al., 2007, 2013b). However, the current study did find an outcome density effect. When designing outcome density experiments, potentially unintended ITIs were considered, hence for all experiments within this thesis, the presence of cue and outcome, and absence of cue and outcome are defined as individual images as such even though there may be intertrial intervals, these should not be defined as no cue – no outcome pairings for the experiments within this thesis. As such the outcome density bias has been replicated here with the absence of intertrial intervals which are seen as no cue-no outcome trials in the literature. Technically, the intertrial interval within this experiment follows the outcome or no outcome image, whereby participants see a "next patient" phrase along with a next button they are required to click**.** However Msetfi et al. (2005) study used the illusion of control, an active version of the contingency paradigm, as such for the passive version used in this thesis, the outcome density effect can be seen with short/ no ITIs.

9.9.2 Depression realism

The depressive realism effect would predict no correlation with low outcome density judgements, and a negative correlation with high outcome density condition judgements. However, depression was not correlated with any of the judgement variables.

Whilst the current study used a within-subjects design to reduce noise between participants in the outcome density conditions, in addition to this the study design was well powered and controlled, the depressive realism effect was not found in this study. There was no significant relationship between depression and judgements in the outcome density conditions.

Previous studies on depressive realism have recruited participants who were pre-screened and scored highly on the respective depression measure used, which was not the case here. To note, the lack of effects seen on the measures may be due to a lack of variation in affect measures, such that there are not enough participants who score highly on those measures. However Msetfi et al. (2015) also used the DASS, but the average scores were 4.65, 3.33, and 5.46 on the depression, anxiety and stress scales respectively in their study, which are much lower than the average scores in the current study of 12.97 (depression), 7.52 (anxiety), and 14.04 (stress). This may suggest that overall, this sample happened to be more depressed, anxious, and stressed, compared to other samples.

A meta-analysis of the depressive realism effect investigated 75 experiments across 7305 participants, and found that generally the effect is quite small (Moore & Fresco, 2012). Both depressed and non-depressed participants showed illusions of control, non-depressed participants showed a bigger illusion of control effect. The meta-analysis also found that methodological variations were a moderator of the depressive realism effect, such that experiment which used self-report rather than clinical interviews and has a lack of objective standard of reality were more likely to find effects of depressive realism. As such it may be that the depressive realism effect is an effect of certain methodological circumstances.

Many of the contingency paradigm experiments examining affect have dichotomised or categorised depression, and used ANOVAs and t-tests to examine group difference and found an effect of depressive realism (Alloy & Abramson, 1988; Msetfi et al., 2007, 2015). As such differences in result could depend on whether analysis uses a continuous variation or dichotomised categories of depression. However Dev et al. (2022) used categories of depression and found no effect of depressive realism. For this experiment, the affect variables were not dichotomised as power to detect an effect would be reduced. In addition to this, there is no reason to assume there is some underlying dichotomy to the affect variables, nor that it would exist at the median. A regression analysis was used here, which essentially breaks down the total variance of the data into different parts and verify these sub variances using a F test, just as an ANOVA would do. As such there is no need to dichotomise the affect measures.

It may be that the depressive realism effect is only seen within clinical populations. However more recent studies on the depressive realism effect in clinical populations have also failed to replicate this effect. Venkatesh et al. (2018) examined the effect with clinically depressed participants along with control participants who were randomly allocated to either a rumination or distraction conditions and then completed a contingency task and provided judgements of control. Neither group differed on their judgements of control, although all participants showed the illusion of control effect.

Passive vs active tasks for depressive realism effect

However the lack of depressive realism effects in this experiment have replicated the lack of depressive realism seen by Alloy et al. (1985) in passive contingency tasks. Fundamentally the depressive realism effect is about the control participants believe they have over an outcome. Within passive tasks participants have no control over the cue nor outcome. As such the hypotheses and models put forward to explain the overestimation of control of participants in active tasks, may not be applicable to the overestimation of contingency seen in passive tasks. The depressive realism effect may be an effect that is only seen within active contingency paradigms, and not all types of contingency paradigms.

9.9.3 Anxiety, Positive and Negative affect

An effect of anxiety on judgement, would predict no correlation with low outcome density judgements, and a negative correlation with high outcome density condition judgements. However, anxiety was positively correlated with judgements in the low outcome density and had no correlation with judgements in the high outcome density or difference in judgements between density.

An effect of negative affect on judgement, would predict no correlation with low outcome density judgements, and a negative correlation with high outcome density condition judgements. Negative affect was not correlated with judgements in either the high or low outcome density conditions.

An effect of positive affect on judgement, would predict positive correlations of positive affect of both outcome density conditions. Positive affect is positively correlated with low outcome density, suggesting at low outcome density, more affectively positive participants gave higher judgements in the low outcome density condition. There was no correlation between positive affect and judgements in the high outcome density condition.

Anxiety, negative, and positive affect significantly predicted contingency judgements in the low outcome density conditions, but not the high-density condition. This suggest that participants who are more anxious, or high in positive or negative affect overestimate contingency judgements compared to participants who are less anxious when the probability of outcome is low. This provides partial support for the Dev et al. (2022) experiment that there is an association between anxiety and contingency judgements, and Novovic et al. (2012) for an association between positive affect and contingency judgements.

Anxiety, negative, and positive affect did not significantly predict contingency judgements in the high-density condition, or the difference between densities. This may be due to the general overestimation of contingencies when outcome probability is high across all participants as seen in the previous experiments of this thesis. This suggests that at high outcome density the overestimation of contingency judgements occurs regardless of how anxious, or positively or negatively one may feel. However, at low outcome density, the more anxious, or positive or negatively affective someone may be the more they may overestimate the contingency judgements. As such those who are less affective on these measures may be better at estimating the contingency at low outcome probabilities.

9.9.4 Stress

An effect of stress on judgement, would predict no correlation with low outcome density judgements, and a negative correlation with high outcome density condition judgements. Stress was positively correlated with judgements in both the high and low outcome density conditions.

Stress was a significant predictor of judgements in the high outcome density conditions, and the difference between outcome density conditions. This suggests that when the probability of outcome is high, those who are more stressed overestimate the contingency judgements. This is interesting as the prior literature and the studies of this thesis have already shown that participants generally overestimate the contingency when outcome density is high. However, these results imply that stressed participants overestimate contingency judgements even more, as such they show an even greater outcome density bias.

9.9.5 Anger

An effect of anger on judgement, would predict positive correlations of anger with all density conditions. Anger was positively correlated with the high-density condition, suggest that participants scoring highly on the anger measure gave more positive judgements. Anger was also correlated with the difference in outcome density judgements, suggesting those scoring highly on the anger measure, gave judgements that differed more between the high and low outcome density conditions.

Anger was a significant predictor of judgements in the high outcome density conditions. This suggests that when the probability of outcome is high, those who are angrier overestimate the contingency judgements.

9.9.6 Arousal effect on contingency judgement

Overall, the results suggest that increased affect across anger, stress, anxiety, positive and negative affect, effect contingency judgements. The previous literature on depressive realism, may have suggested that negative emotions, that have negative valence results in more accurate perception of contingencies. However, the results of this study do not support that, and instead indicates that the valence of the affect measures are perhaps not that important. Rather the arousal one experiences seems to influence judgements. For each of the significant affect measures the higher they scored on the measures the more positive the contingency judgement. Also, both positive and negative affect are negatively correlated yet both were significant predictors of contingency judgements. Overall, those who were affectively more neutral, or perhaps have lower levels of arousal have more accurate perception of contingencies. Those who were more aroused, regardless of whether the arousal was in terms of positive or negative affect, were more likely to judge contingencies more positively, than those less aroused. This is quite speculative, as I have not run any studies on arousal to see whether there is an influence on contingency judgements. As such a future direction to take this research, may be to explore the effect of arousal.

9.9.7 Theoretical Evaluation

Contingency account

Contingency theory, and Wasserman's (1990) extension do not consider the influence of other factors, as such neither can account for these findings. I found that people in a higher affective state judge contingencies to be more positive. Using the framework from Wasserman (1990), this would mean that it is as if the state of affect increases weight given to cells A and D relative to other cell counts. It may be possible that people who feel more affective whether positive or negative, weight the cells differently to those who are affectively neutral. For example, those who score highly on affect measures may place even greater importance or pay more attention to cell A trials, or the occurrence of outcomes.

Associative account

Whist the associative perspective does not include the role of affect in the model, as such it would not be able to account for these findings. However previously it was assumed that differential salience could occur for cue and outcome. As such a further assumption could be made that the perceived salience for cue and outcome could vary between individuals. If this is the case, then perceived salience could possibly vary due to affect.

Those may feel more affectively, may have greater salience for the cue and/or outcome, in which case they would be expected to give more positive judgements of contingency, as found in this experiment.

Overall, the Rescorla-Wagner model does not necessarily account for the role of affect. However, it could potentially account for it.

Probabilistic account

The model does not include the role of affect in the model, as such it would not be able to account for these findings.

9.9.8 Limitations and Future Directions

This experiment was primarily exploratory and, as such, further research on this finding could be carried out. For example, a within-subjects experiment in which participants complete both a passive and active contingency task, in which affect is also measured could be conducted. Such a study could assess how judgements differ between tasks, and if affect differentially influences contingency judgements between tasks as well. This would shed further light on how the affect may vary between tasks. However, it is still unclear what aspect of affect is influencing contingency judgements, as such further exploration of the different aspects of affect, such as arousal could be examined. Hence a study could be run which uses a short-term mood manipulation to identify whether such changes in mood will influence judgements. Participants could complete a guided autobiographical recall task to make them briefly feel slightly elevated and slightly flattened in mood, to see if the effects seen here can be replicated with induced mood.

9.10 Conclusion

A robust effect of outcome density found. No depressive realism effect was found, however there was considerable individual differences in the affective states, as such the conditions necessary to find individual difference were available. These findings may be due to certain effects not occurring with the passive contingency judgment task. The previous findings of affect influencing judgements of control may be related to the task type used. Hence methodological variations within the illusion of control, such as ITI length may result in effects such as depressive realism.

Overall, these findings suggests that certain types of affective states may influence contingency judgements. Stress and anger are associated with more positive judgements when outcome probability is high, leading to even larger outcome density effects, as the difference between judgements in the low-high density conditions are even greater. Increased anxiety, positive, and negative affect increase contingency judgements in the low outcome density condition. This suggests that participants who are low on these measures are more accurate at judging contingencies, and consequently be less likely to form false beliefs.

Chapter 10: Discussion

Over 7 experiments discussed across 7 empirical chapters, I explored the conditions from which false beliefs form, focusing primarily on the effects of outcome density and cue density. I examined how this cognitive and learning process is influenced by methodological variations of the contingency learning paradigm. To further understand this process, I examined the influence of belief and affect on the contingency judgements. To test the generalisability of the findings to more real-world settings, this effect was examined using new cover stories, bridging the gap between the experimental context and the real-world.

In this final chapter, I will discuss the overall findings in relation to theory and consider empirical implications, along with discussing an integrative analysis which considers the size of every density effect in this thesis. This chapter also discusses the strengths and limitations of this thesis, along with providing suggestions for future research.

10.1 Aims and Summary of Findings

In today's world the prevalence of fake news has become increasingly problematic, due to the arising false beliefs (Ecker et al., 2022). Simply exposing individuals to facts has been proven to largely fail in eradicating such false beliefs (Schwarz et al., 2007). As such, a clear understanding of the conditions under which false beliefs are formed is needed to develop effective countermeasures against fake news.

Many false beliefs are false causal beliefs; these may encourage courses of action that do not result in the hoped-for benefits. Causal beliefs arise from individuals identifying causality, the process through which one event causes an effect, and the effect is dependent on the cause to occur. Causal relationships suggest that cause and effect are related, even if it may be temporary, and that the outcome follows the cause and does not precede it. People are predisposed to form causal beliefs and do so from a very early age (Sobel & Kirkham, 2006). However the inferences made about the underlying mechanisms of causality do not always match the causal belief about cause and effect, often described as a cognitive illusion (Sawa, 2009). Studying the mechanisms through which causal beliefs and the associated reasonings are made gives further understanding to issues that arise from cognitive illusions such as pseudoscience and superstitious thinking.

Examining these mechanisms becomes particularly important when the effect of forming such false beliefs become dangerous for health when individuals start to rely on alternative forms of medicine due to such false beliefs (Baker & Rojek, 2020). The contingency learning account provides a framework through which these mechanisms can be examined.

The aim of this thesis was to identify conditions under which false beliefs may form, particularly false beliefs formed due to high outcome density, as a high density of information has been linked to increased belief in fake news (Hills & Menczer, 2020). Three research questions were explored:

1. Are the cue and outcome density effects robust effects?
2. Do methodological variations of the contingency paradigm affect the size of the cue and outcome density effects?

3. Do factors beyond methodological variations influence the outcome density effect?

10.1.1 Are the cue and outcome density effects robust effects?

The thesis reports evidence that the outcome density effect is a robust effect, which occurs across different conditions. The hypothesis that belief in an association between two events will increase, with higher frequency of alleged cause and observed outcomes, even if there is no actual relationship between the two, was supported.

Experiment 1 (Chapter 3) examined the size of the outcome density effect in the laboratory and online using zero contingency paradigms, to see if the effect is due to the settings used in previous experiments. The findings suggest that the outcome density effect is present in both settings and that the contingency paradigm can be used effectively online. This supports previous findings (Matute et al., 2007b). The findings also confirm that the outcome density effect is a large effect; participants gave higher judgements of contingency when the probability of outcome was higher compared to when the probability of the outcome was lower.

Experiment 2 (Chapter 4) examined cue density effects in comparison to outcome density effects. I replicated the findings for outcome and cue density effects in contingency learning from Blanco et al. (2013), as well as the outcome density findings from Experiment 1 in online settings. My finding shows that the outcome density effects are larger than cue density effects on average. The relative difference in size appears to depend on other factors. Thus, the findings suggest that there is an effect of both cue and outcome density, however the cue effect seems to be smaller than the outcome density.

Overall, the outcome density effect was found to be a robust effect, observed consistently across all experiments of this thesis. People tend to judge contingencies between cue and outcome to exist when there is a higher probability of outcome, even when there is no true contingency. This density effect is greater for outcomes than cues.

10.1.2 Do methodological variations of the contingency paradigm affect the size of the cue and outcome density effects?

The second research question was how the cue and outcome density effect varied according to methodological variations. The thesis found evidence to suggest that cue and outcome density are robust effects, however the size of the effect varies across methodological manipulations.

Experiments 3 and 4 (Chapter 5) examined whether the order in which cues and outcomes are shown influence contingency judgements and consequently the outcome density effect. Whilst logically a cause must precede an outcome, sometimes people may observe or learn of an outcome before they see or hear its possible causes. For example, a person may start to experience symptoms of covid-19, and may have to think back to the past days to identify the possible cause for virus transmission. In this pair of experiments the order in which cue and outcome were presented was manipulated within the experiments, and across the experiment either cue density or outcome density were manipulated. As such the effect of cue-outcome presentation order on the cue and outcome density effects were examined. The findings suggest that event presentation order makes little or no difference to the judgement of medication effectiveness or to the size of the outcome density and cue density effect. Whilst the outcome density effect was larger than the cue density effect, replicating the findings of the second experiment, across the experiments, the size of effects for cue and outcome density were relatively consistent. The finding suggests that contingency judgments do *not* rely on events being encountered in their 'logical' order with the putative cause preceding the outcome.

Experiment 5 (Chapter 6) examined the outcome density effect across positive and negative contingencies in comparison to zero contingency conditions. The experiment also examined how judgments differed between unidirectional and bidirectional scales. Main effects of scale type and contingency were found. The findings here suggest that the 0 to 100 scale may encourage participants to provide more positive contingency judgements due to a lack of other options, as the scale provides 100 options for a positive judgement, and only one option for a non-positive judgement. In contrast the -100 to 100 scale offers 101 options for non-positive judgements as well. By offering a unidirectional scale which runs positively, participants may implicitly believe their judgement has to be positive, regardless of what they may perceive, and as such provide a positive judgement. However, when considering the high outcome density conditions for the zero contingency with the new scale, a large proportion of participants still detect a positive contingency, indicating a general tendency to make positive contingency judgements when there is an increased probability of the outcome. As such this effect may just be exaggerated with the use of the 0 to 100 scale.

Across the three contingencies, positive, negative and zero, the outcome density effect was strongest in the zero-contingency condition, and weakest in the negative contingency condition. The contingency calculations were symmetrical for the positive and negative contingencies, however the effects seen are not symmetrical. Whilst the effect sizes for the mean difference of positive contingency – zero contingency (0.84) and negative contingency – zero contingency (0.77) are both large and relatively close, the mean judgements are not equidistant from 0 as may be expected based on the symmetrical calculations. Rather all the mean judgements seem to be shifted more positively. In sum, participants seem to be making associations that are, on average, more positive than the contingency would indicate but which are further shifted in a positive direction in the presence of high outcome density. Importantly, by including both positive and negative contingencies and using a new scale to accommodate both directions of contingency, I was able to demonstrate that increasing the outcome density makes judgements more positive rather than making them stronger.

Experiment 6 (Chapter 8) examined if the size of the outcome density effect varied between different cover stories. Different cover stories have been used to study the outcome density effect in contingency learning tasks. While different cover stories are used in different studies, there has been little research that formally examines variation in contingency judgment as a function of cover story. The findings suggest that the outcome density effect does vary across different cover stories. This indicates that when examining effects using contingency paradigms that cover story used should be chosen with care, to extract general principles about human contingency learning which may be hampered by restricted choice of scenario.

Overall, these experiments have explored the different methodological conditions under which the outcome density effect may vary. The order in which cue and outcome shown do not seem to affect the density effect. However, the scale type, contingency, and cover stories do affect the outcome density effect. This suggests that these are different conditions under which stronger or weaker beliefs may be formed.

10.1.3 Do factors beyond methodological variations influence the outcome density effect?

The final research question diverged from the pure contingency paradigm to examine if the factors of affect and prior beliefs influence the outcome density effect.

Experiment 6 (Chapter 8) examines the influence of prior beliefs on the outcome density effect. Many of the prior experiments on prior belief in the contingency literature had manipulated expectation, as such in this study participants' actual beliefs were measured to examine the influence of belief on contingency judgments. The findings suggest that the extent to which prior beliefs impact contingency judgements depends on the cover story used. As such, it is important to consider the influence of pre-existing prior beliefs both within experimental contexts and in the real world. In the experimental settings, pre-existing beliefs may influence judgements differentially for participants, and unless accounted for, may affect the validity of findings. Similarly in real-world settings, interventions which aim to reduce false beliefs, should be aware of the potential effects of pre-existing beliefs.

Experiment 7 (Chapter 9) examined the influence of affect on contingency judgement. Within the contingency literature there has been a great deal of research on the depressive realism effect. The effect suggests that participants who are more depressed tend to give more accurate contingency judgements compared to non-depressed participants. This experiment found no evidence of the depressive realism effect. Overall, the findings suggest that certain types of affective states may influence contingency judgements. Stress and anger are associated with more positive contingency judgements when the outcome probability is high, leading to even larger outcome density effects, as the difference between judgements in the low-high density conditions are even greater. Increased anxiety, positive, and negative affect increase contingency judgements in the low outcome density condition. This suggests that participants who are low on these measures are more accurate at judging contingencies. This may potentially be related to levels of arousal, such that those who are high on arousal may make more biased positive contingency judgments. As such further research to test the robustness of these findings, and to tease out the possible explanations would be recommended.

10.1.4 Summary

Overall, this thesis provides evidence that false beliefs are reliably formed when the probability of outcome is increased, even when there is no causal link between the cue and outcome. However, the extent of the strength of false beliefs may vary according to the conditions under which they are presented. In addition to this, the prior beliefs individuals hold may overcome the outcome density effect, as such be less likely to form false beliefs. Also, people who are affectively more neutral may give more accurate contingency judgment, consequently, be less likely to form false beliefs.

10.2 Methodological and Empirical contributions, and Limitations and Future Directions

Here the implications, limitations, and future directions of the findings of this thesis will be discussed. The interpretations of individuals findings have been discussed in the respective empirical chapters.

10.2.1 General contributions

To examine the effect of trial types across experiments, a form of meta-analysis was conducted across the experiments of this thesis. The mean contingency judgement of each condition within each experiment was correlated with the proportion of cell A, B, C, and D trial types per condition. Partial correlations were also run to control for the scale type used, unidirectional and bidirectional, these are shown in Table 58. Finally, the relative frequency of A, B, C, and D were also correlated with each other as shown in Table 59. Results indicate the mean contingency judgements were strongly positively correlated with cell A trials across conditions. This suggests that as the number of type A trials increases, contingency judgments become more positive. Mean contingency judgements were not significantly

Table 58 Correlation between mean contingency judgement and relative frequency of trial types across all experiments, and trial types, according to the contingency table, df=34. Correlations of trial type by mean contingency judgement controlling for scale type is also shown, df=33

	Trial type			
	A	B	C	D
Mean contingency judgement	.60**	-.49**	.26	-.37*
Mean contingency judgement controlling for scale type used	.82**	-.67**	.40*	-.55**

** Correlation is significant at the 0.01 level (2-tailed)

* Correlation is significant at the 0.05 level (2-tailed)

Table 59 Correlations between relative frequency of trial types across all experiments, and trial types, according to the contingency table, df=34.

	Trial type			
	A	B	C	D
A		-.60**	.50**	-.90**
B			-.90**	.50**
C				-.60**

** Correlation is significant at the 0.01 level (2-tailed)

* Correlation is significant at the 0.05 level (2-tailed)

correlated with the proportion of type C trials. Both type B and D trials were negatively correlated with mean contingency judgements, suggesting that as the proportion type B and D trials increase, contingency judgements decrease. As such this shows that relative to the proportion of trials A and B, C and D receive less weighting, based on the strength of the correlation. However due to the high correlations between these variables as seen in Table 58, further differentiation cannot be determined. Interestingly when controlling for scale type (unidirectional and bidirectional scales used), the correlations have all become stronger, and the mean contingency judgement becomes significantly correlated with trial type C.

However, Wasserman (1990) found that people tend to weigh the cells of a contingency table (Table 60) (reproduced from Chapter 1) according to, A > B > C > D for judgements of data presented in contingency tables. As such people tend to weigh cell A trials the most followed by cell B and then C trials, and cell D trials are given the least weighting. In addition to this, it was also suggested that the trials of cells A and D lead to increased positive judgements whilst cell C and B counts decrease contingency judgements.

The findings of the above correlation also found that increasing the proportion of type A trials increases the positivity of judgements, whilst increasing the proportion of type B trials decreases judgements. This support Wasserman's (1990) findings. However, this thesis also found that increasing proportion of type D trials decreased judgment which does not support Wasserman's (1990) findings. This may be due to the majority of experiments in

Table 60 Visual presentation of the possible outcome in a contingency paradigm

	Outcome	
Cue	Present	Absent
Present	A	B
Absent	C	D

this thesis utilising zero contingency paradigms, to ensure the following criteria were met for the zero-contingency paradigm: cue held at a probability of 0.5 across conditions, P(O/R) =P(O~/R), equal number of trials for both P(O/R) and P(O~/R), the trials proportion of A and C were always the same, as was the trial proportion for B and D. This then forces a negative correlation between A and D, and between B and C cells, making it difficult to assess the weight for D independently of A; and to assess the weight for B independently of C. As such the correlations seen here are likely due the methodological variation.

However, within Wasserman's (1990) series of experiments, one study did find that participants weighed trial A and B higher, than trials C and D overall, which is generally supported by the findings of this thesis. In addition to this, in the series of experiments by Wasserman, participants were presented with contingency tables, as such descriptions of the contingency were provided. In the experiments across this thesis, participants were exposed to trial by trial information of the contingency. As such using different methodologies, similar results have been found. Overall supporting the findings of Wasserman. To further confirm this effect, studies which utilise zero continency calculation which are not symmetrical could be explored. An experiment which switches the cell which participants are asked about could be one line of enquiry here. Currently participants are asked about cell A trials, as such would judgements differ if participants are asked about the other cell types. The contingency questions are always worded in terms of presence of cue and presence of outcome. It would be interesting to see if judgements differ if questions are worded in terms of the other 4 cells, see Table 61.

Another direction to take this research is to examine further methodological variations of the contingency paradigm. A general limitation of the methodology of this thesis is the use of the contingency paradigm itself. The task itself was not varied, the same number of trials, the same timing etc were used throughout. This allows one to see the different effects arise across the same paradigm. It would be useful to examine further variations of the paradigm, for example using fewer trials, to see if the outcome density effect can be seen with half the number of trials.

This thesis has provided a clear insight into the biases that arise around contingency judgements using a consistent paradigm. There are other biases which may lead to false beliefs, as outlined in chapter 1. Adopting a similar consistent and methodological approach taken in this thesis could be helpful in understanding these biases.

10.2.2 The outcome density effect across experiments

Looking across the studies in this thesis, the evidence suggests that the outcome density effect is a robust effect. See Table 62 for comparisons across all experiments of this thesis. Participants consistently give higher contingency judgements when the probability of the outcome is high compared to when it is low, in a zero contingency. As such people tend to judge a more positive contingency between cue and outcome, forming a false or biased

Table 61 Possible changes to judgement question

	Outcome	
Cue	Present	Absent
Present	A - To what extent have patients ***recovered*** **<u>with medication</u>**?	B - To what extent have the patients not ***recovered*** **<u>with medication</u>**?
Absent	C - To what extent have patients ***recovered*** **<u>without medication</u>**?	D - To what extent have patients ***not recovered*** **<u>without medication</u>**?

Bolded italics show the cue

<u>Bolded underlined</u> shows the outcome

Table 62 Summary of experiments, with all manipulations, showing the mean contingency judgements, standard deviations and sample sizes across all conditions. The manipulations of each conditions are shown along with the average cue and outcome density per

Study	Condition	Mean Contingency judgment	SD	N	Effect size (*d*) between low and high outcome density conditions	Density probability	Location	Density type	Order of cue and outcome	Scale used	Contingency	Cover story	Average outcome density	Average cue density	ΔP	Proportion of trial types per condition			
																A	B	C	D
1	1	19.57	16.95	58	1.36	Low	Lab	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	2	47.4	23.49	58	1.36	High	Lab	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10
	3	25.45	20.88	58	0.85	Low	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	4	44.83	24.53	58	0.85	High	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10
2	1	18.61	16.85	85	1.49	Low	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	2	43.74	16.89	85	1.49	High	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10
	3	31.29	23.95	85	0.54	Low	Online	Cue	Cue-outcome	Unidirectional	Zero	Illness	.50	.20	.00	.10	.10	.40	.40
	4	43.45	20.8	85	0.54	High	Online	Cue	Cue-outcome	Unidirectional	Zero	Illness	.50	.80	.00	.40	.40	.10	.10
3	1	17.52	14.94	107	1.17	Low	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	2	42.93	26.96	107	1.17	High	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10
	3	20.83	15.65	107	1.18	Low	Online	Outcome	Outcome-cue	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	4	44.7	23.99	107	1.18	High	Online	Outcome	Outcome-cue	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10
4	1	35.7	20.33	97	0.55	Low	Online	Cue	Cue-outcome	Unidirectional	Zero	Illness	.50	.20	.00	.10	.10	.40	.40
	2	46.55	19.1	97	0.55	High	Online	Cue	Cue-outcome	Unidirectional	Zero	Illness	.50	.80	.00	.40	.40	.10	.10
	3	31.58	23.23	97	0.69	Low	Online	Cue	Outcome-cue	Unidirectional	Zero	Illness	.50	.20	.00	.10	.10	.40	.40
	4	46.87	21.31	97	0.69	High	Online	Cue	Outcome-cue	Unidirectional	Zero	Illness	.50	.80	.00	.40	.40	.10	.10

5	1	32.84	18.59	44	0.6	Low	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.35	.50	.00	.18	.33	.18	.33
	2	44.7	21	44	0.6	High	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.65	.50	.00	.33	.18	.33	.18
	3	-5.14	34.38	87	0.82	Low	Online	Outcome	Cue-outcome	Bidirectional	Zero	Illness	.35	.50	.00	.18	.33	.18	.33
	4	22.54	33.39	87	0.82	High	Online	Outcome	Cue-outcome	Bidirectional	Zero	Illness	.65	.50	.00	.33	.18	.33	.18
	5	48.11	18.04	44	0.84	Low	Online	Outcome	Cue-outcome	Unidirectional	Positive	Illness	.35	.55	.30	.25	.25	.10	.40
	6	63.41	18.41	44	0.84	High	Online	Outcome	Cue-outcome	Unidirectional	Positive	Illness	.65	.45	.30	.40	.10	.25	.25
	7	26.23	41.23	44	0.42	Low	Online	Outcome	Cue-outcome	Bidirectional	Positive	Illness	.35	.55	.30	.25	.25	.10	.40
	8	43.09	39.61	44	0.42	High	Online	Outcome	Cue-outcome	Bidirectional	Positive	Illness	.65	.45	.30	.40	.10	.25	.25
	9	-20.84	36.14	43	0.32	Low	Online	Outcome	Cue-outcome	Bidirectional	Negative	Illness	.35	.45	-.30	.10	.40	.25	.25
	10	-9.53	34.15	43	0.32	High	Online	Outcome	Cue-outcome	Bidirectional	Negative	Illness	.65	.55	-.30	.25	.25	.40	.10
6	1	-15.26	37.66	102	1	Low	Online	Outcome	Cue-outcome	Bidirectional	Zero	Marijuana	.20	.50	.00	.10	.40	.10	.40
	2	20.12	33.24	102	1	High	Online	Outcome	Cue-outcome	Bidirectional	Zero	Marijuana	.80	.50	.00	.40	.10	.40	.10
	3	-12.81	37.37	102	0.54	Low	Online	Outcome	Cue-outcome	Bidirectional	Zero	Social media	.20	.50	.00	.10	.40	.10	.40
	4	7.15	37.05	102	0.54	High	Online	Outcome	Cue-outcome	Bidirectional	Zero	Social media	.80	.50	.00	.40	.10	.40	.10
	5	-10.87	35.92	103	0.78	Low	Online	Outcome	Cue-outcome	Bidirectional	Zero	Tech	.20	.50	.00	.10	.40	.10	.40
	6	14.44	28.76	103	0.78	High	Online	Outcome	Cue-outcome	Bidirectional	Zero	Tech	.80	.50	.00	.40	.10	.40	.10
	7	-11.11	38.7	101	1.04	Low	Online	Outcome	Cue-outcome	Bidirectional	Zero	Healthcare	.20	.50	.00	.10	.40	.10	.40
	8	25.42	31.46	101	1.04	High	Online	Outcome	Cue-outcome	Bidirectional	Zero	Healthcare	.80	.50	.00	.40	.10	.40	.10
7	1	18.25	16.383	169	1.19	Low	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.20	.50	.00	.10	.40	.10	.40
	2	45.98	28.586	169	1.19	High	Online	Outcome	Cue-outcome	Unidirectional	Zero	Illness	.80	.50	.00	.40	.10	.40	.10

Please note a gradient of green (high) to red (low) has been used on any column with numerical information, to easily show differences in values.

belief, even when there is no relationship between the cue and outcome. Overall, this may suggest a form of predisposition to weigh the occurrence of outcome highly.

Across the studies, the size of the outcome density effect varies from 0.6 to 1.49 in size, as seen in Table 62 which shows effect sizes and mean contingency judgements across all conditions of the experiments in this thesis. Interestingly, the largest outcome density effects are seen in experiment 1, experiment 2 (conditions 1 and 2), experiment 3, and experiment 7, with outcome density ranging from 0.85 to 1.49. In these conditions, a unidirectional scale was used with probabilities of outcome at 0.2 and 0.8. In Experiment 5, condition 1 and 2, which use the unidirectional scale, the probabilities of outcome were 0.36 and 0.65, giving an effect size of 0.6. This is nearly half the other equivalent condition such as experiment 2, conditions 1 and 2. Overall, this suggests that the outcome density effect can be reduced and depending on the manipulations there can be variability in the outcome density effect. However further exploration, which systematically explores how contingency judgements vary at different probabilities of outcome, and how this may affect the overall outcome density effect. That is, do contingency judgements in zero contingency increase linearly with increased outcome probability?

A potential explanation for some of the results in this thesis, that I have not explicitly tested, is the possible effect of intertrial interval (ITI). Msetfi et al. (2005) found that varying the programmed ITI can affect the *experienced* cue-outcome contingency. The ITI is similar to a cell D trial; both cue and outcome are absent, thus fluctuations in perception of the ITI might influence the perceived proportion of type D trials, consequently influencing perceived contingency. Overestimations of the contingency between events in a high density with long ITI condition can be explained by the contingency learning account because the additional ITI period adds exposure to "no cue and no outcome," changing the overall contingency. Table 63 shows how the contingency table would change if ITIs were factored into the calculation. Including ITIs in the contingency calculations, the low outcome zero contingency condition becomes mildly positive, and the high outcome zero contingency condition becomes moderately positive. As such, incorporation of ITIs may explain the outcome density effect.

However, in the passive contingency experiments of this thesis, the cue and outcome presentation durations and ITI are precisely controlled, with trial periods and ITIs distinctly demarcated. The presence and absence of cues and outcome are defined by 4 separate images. This is very different to the type D trials programmed in an active task, such as those used by Msetfi et al., where the distinction between trials and the ITI is less clear. Therefore, it is less clear why participants would incorporate the ITI into contingency judgements in a passive task. Further, the role of the ITI should be identical across manipulations of cue and outcome density, and thus incorporation of the ITI does not

Table 63 Contingency calculation of Experiment 1, with and without ITIs

Table A Low outcome density contingency calculation

	Outcome		
Cue	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	4	16	P(O/~C) = .2

ΔP = .2-.2=0

Table B Low outcome density contingency calculation with ITI

	Outcome		
Cue	Present	Absent	
Present	4	16	P(O/C) = .2
Absent	4	56	P(O/~C) = .07

ΔP = .2-.07=0.13

Table C High outcome density contingency calculation

	Outcome		
Cue	Present	Absent	
Present	16	4	P(O/C) = .8
Absent	16	4	P(O/~C) = .8

ΔP = .8-.8=0

Table D High outcome density contingency calculation with ITI

	Outcome		
Cue	Present	Absent	
Present	16	4	P(O/C) = .8
Absent	16	54	P(O/~C) = .23

ΔP = .8-.23=0.57

naturally explain the difference in contingency judgments observed between manipulations of cue density and outcome density.

10.2.3 The cue density effect across experiments

Cue density effects were found across Experiments 2 and 4. The size of cue density effect were relatively similar across the experiments. Outcome density effects were larger than the cue density effects across experiments, 2, 3, and 4, see Table 64. Interestingly when comparing across the contingency responses, for both the cue and outcome density manipulation the high outcome density conditions seem to be fairly similar. However, the responses for the low-density conditions, are much lower for the outcome density manipulation compared to the cue density manipulation. This may suggest that at high density, similar weighting is given to both cues and outcomes. At lower density, lower weighting is given to the outcome compared to the cue. As such the outcome density effect

Table 64 Comparison of means (SD), CIs and effect size for cue and outcome density across experiments 2, 3 and 4

	Outcome Density			Cue density		
	Experiment 2	Experiment 3		Experiment 2	Experiment 4	
		Cue-Outcome	Outcome-Cue		Cue-Outcome	Outcome-Cue
High density response	43.74 *(16.89)*	42.93 *(26.96)*	44.70 *(23.99)*	43.45 *(20.80)*	46.55 *(19.10)*	46.87 *(21.31)*
Low density response	18.61 *(16.85)*	17.52 *(14.94)*	20.83 *(15.65)*	31.29 *(23.95)*	35.70 *(20.33)*	31.58 *(23.23)*
95% CI for mean difference	[19.41, 30.85]	19.77, 31.05	19.03, 28.70	[6.54, 17.77]	6.07, 15.61	9.63, 20.95
Density effect size *d*	1.49	1.17	1.18	0.54	0.55	0.69

may be driven by the lower density conditions. This may also suggest that density or probability of cue, does not necessarily affect the weighting given to cues. As the given contingency judgements when cue is manipulated are relatively similar.

The fact that cue and outcome density effects are different in size does seem to suggest that they different in some aspect. This may potentially be in terms of how much attention people pay to the stimuli or perhaps the salience of the events. In fact Experiments 2, 3, 4 shows that the density effects differ in size between cue and outcome, and even when the order of these are switched, such as the outcome is presented first, the outcome density effect is always larger. This does seem to indicate that outcomes and cues are perceived differently.

The relationship between contingency judgements and cue density could be further examined to see if across various levels of cue probability, do contingency judgements stay fairly similar. in addition to this, cue density effects were not examined in positive and/or negative contingencies, which may shed further light on how cue effects may influence judgements.

10.2.4 Information presentation

Within this thesis only single cue-outcome pairs are shown to participants within a block of trials. However, in the real world it is likely that people will observe multiple different cues and outcomes, which was not really examined here. However, the contingency literature in general has examined the effects of multiple cues and outcomes, although perhaps not with the same factors examined here, hence this may be an extension to this work (De Houwer & Beckers, 2003; López et al., 1998).

Another factor to consider is that in real life there may be a temporal gap between a cue and outcome, and in fact the gap may not be consistent. This could potentially influence the associations or contingencies people form between events, as memory could be a large factor at play here.

For participants, the approach taken in this thesis, means that they can learn contingencies in a fairly simple to understand way. The use of pictures makes things easier. The approach allows for some level of noise to enter, as participants for all experiments took part in the experiments from their home location, emulating how participants may be exposed to information on social media. It was not a fully controlled lab study, with no noise. Whilst this could affect the results, it was a more realistic setting.

Fatigue effects or boredom are major concerns for experiments like this, where participants view a series of trials. There are 40 trials in each block. To reduce participants not paying attention between each trial they have to manually click on a next button. However, it may be that participants do stop paying attention, however if they did completely stop paying attention, then we would not expect to see any differences in effect or judgments between the high- and low-density conditions, however, this density effect is present throughout all the experiment, suggesting that participants are paying attention throughout the entire experiment. It is likely that participants are using central routes of processing, as they are told what each image is.

Active vs passive paradigm

All the experiments in this thesis have utilised passive paradigms, to reflect the way in which participants may see information in their daily lives, for example on social media.

However realistically many people are engaging with the information they receive, for example by sharing information. The active side of human behaviour was not explored in this thesis, however it has been examined across many experiments in the contingency literature (Gillan et al., 2014b; Novovic et al., 2012; Yarritu et al., 2014). A series of experiments in which participants take part in a mixture of passive and active paradigms may be more reflective of real life scenarios. However prior to that further research on the differences in judgements between the passive and active paradigms may be useful to pinpoint how or why contingency judgements differ across the paradigms.

10.2.5 Interventions

A potential application of the findings of this thesis, can be in the use of interventions in preventing or reducing the formation of false beliefs which may results in negative health consequences. This research has highlighted the importance of several factors as such prior belief, and how people may be resistant to belief change due to prior belief. Also, the influence of affect on belief formation has been highlighted by the findings in Chapter 9. These factors could be incorporated into policy regarding these types of interventions, to ensure the interventions have a higher probability of working, on more of the population. Previous research has already shown that often just telling people what to look out for, or what to do is not very effective (Lewandowsky et al., 2012), as such these findings could be utilised to make more effect interventions.

The intervention could perhaps be in the form of a game whereby, participants complete two blocks of trials of cue-outcome pairings in a contingency paradigm. In one block there could be no manipulation, however in the second block, perhaps the cover story or affect of the participants could be manipulated. Afterwards the difference in judgement could be shown to participants to highlight how affect or prior belief and other factors can influence judgements.

The focus of the research has been on the general population, with the participants being of the general population. No specific clinical group were examined. Potentially the findings could be extended into clinical areas, whereby belief formation is of importance, for example in anxiety, however this has not been a main aim of the thesis.

10.2.6 Generalisability

The findings of this thesis are likely fairly generalisable, because the samples were recruited from a reasonably diverse population (e.g., greater dispersion in age than is found in most student populations). The only restrictions applied were that participants had to be, over the age of 18, and due to ethics, that participants were primarily based in the UK. The scenarios used, are not particularly country specific. In addition to this, the paradigm uses pictures to represent different events, to minimise confusing from specific wording. However cross-cultural differences may exist, whereby the pictures could be interpreted differently in different cultures or different values of belief may exist which would mean these pictures cannot be used universally. For example, cultures that may rely more on alternative medicine or holistic herbal medicines may have different views on the medication scenario, and as such there may be external influences on the judgements they may make. To explore this possible influence of prior belief was explored in this thesis.

10.3 Theoretical Contributions

Three theoretical accounts were examined to see how well they can explain the results in the thesis, each listed in Table 65.

The contingency account (Allan, 1980) is a normative account of contingency belief, which states that belief should correspond with the difference in the probability that two events co-occur, P (Event 1|Event 2), and the probability that one of these events occurs in isolation, P (Event 1|No Event 2). The contingency account would predict judgements of contingency to not differ in zero contingency conditions, even if the probability of cue or outcome is manipulated. The contingency account was able to partially or with assumptions explain the findings of experiments 5 and 6. For experiment 5, the findings in regard to judgement differences between zero, negative and positive contingencies were predicted by the contingency account. However, the differences in judgements between the 0 to 100

Table 65 Summary of how well the theoretical accounts explain the findings of the experiments of this thesis.

	Contingency Account	Contingency account with Wasserman Extension	Associative Account	Probabilistic Account
Experiment 1	No	Yes	Yes	Yes, with assumptions
Experiment 2	No	Yes	Yes, with assumptions	No
Experiment 3 & 4	No	Yes, with assumptions	No	No
Experiment 5	Partially	Partially	Partially	Partially
Experiment 6	Yes, with assumptions	Yes, with assumptions	Yes, with assumptions	Yes, with assumptions
Experiment 7	No	No	No	No

scale and -100 to 100 scale were not explainable by the account. The findings of experiment 6 regarding the effect of prior belief could be explain by the contingency account if it is assumed that prior beliefs are considered prior experience of the 4 types of contingency trial types which reinforce the contingency belief. The contingency account could not explain the findings of experiments 1, 2, 3, 4, and 7, as such the account is a good model to explain outcome and cue density effects in general.

Wasserman (1990) provided an extension to the contingency account and suggested that people tend to weigh the cells of a contingency according to, A > B > C > D when making judgements of data presented in contingency tables. This suggests that the trial types presented by a contingency table are not weighed equally or experienced equally by individuals. The contingency account with the Wasserman (1990) extension was able to explain, with some assumptions, all the findings except for the findings of experiment 7. The cue and outcome density effects found in experiments 1 and 2 were explaining by the extension, however for experiment 3 and 4, if it is assumed that the order of cue and outcome does not matter than the findings of this experiment are also predicted. Overall, the contingency account with the Wasserman extension is one of the better accounts in being able to explain the findings of this thesis.

The associative perspective focusses on the formation and strengthening of cue and outcome associations, in particular it considers the associative strength of two events, the conditioned stimulus (CS) and unconditioned stimulus (US) (Rescorla and Wagner, 1972). Strength of connection is determined by how often the events are paired together, how often events occur in isolation and how often both are absent. It is this mechanism that allows the model to be sensitive to CS–US contingencies in animal and human learning. The associative account would predict differences in judgements of association until the associative strength between cue and outcome reaches the asymptote. The associative account was able to partially or with assumptions explain the findings of experiments 1, 2, 5, and 6. It could not explain the findings of experiments 3, 4 and 7. The outcome density effects of experiment 1 was predicted by the associative account, and the cue density effects of experiment 2 could be predicted if the presence and absence of the cue are accounted for as well. However associative accounts assumes that cues are followed by outcomes, as such the findings of experiments 3 and 4 cannot be explained. The finding of experiment 5 that judgement differences across zero, negative and positive contingencies were predicted by the associative account. However, the judgment differences between the 0 to 100 scale and -100 to 100 scale were not explainable by the account. The findings of experiment 7, that prior beliefs influence contingency judgements could be explained by the account if it is assumed that prior beliefs are developed from prior real-life experience of cue and outcome occurrences, and these could carry over into the experimental paradigm. Overall the associative perspective is able to account for cue and outcome density effects, however when the role of affect or interchangeability of cue-outcome order is considered, the account falls short in its predictions.

The probabilistic model is a rule-based model, which suggests that individuals reason about cause-effect relationships by recalling all the different cue-outcome pairings and calculating a similar to Δp (Cheng & Novick, 1990). Individuals compare situations that differ due to the presence of a target cue, whether the cue is present or absent. By comparing these situations inferences about the relationship between cue and outcome can be learnt. This model would predict there to be no difference in judgements in zero contingency conditions. The account was able to partially or with assumptions explain the findings of experiments 1, 5, and 6. It could not explain the findings of experiments 2, 3, 4 and 7. The outcome density effect of experiment 1 was predicted by the account, if it is assumed that a zero contingency condition is misperceived by people as being slightly positive or negative. However, the cue density effects, and cue-outcome presentation order effects seen across experiments 2, 3 and 4 are not explainable by the account. Even if it assumed that zero contingencies are misperceived to be positive or negative, the difference in size between cue and outcome density effects cannot be accounted for, and the account assumed cues are followed by outcomes. The model can account for the contingency effect found, as it does predict that more extreme levels of outcome density results in higher judgments of control in positive and negative contingency conditions. The model predicts that some participants will perceive zero contingencies as non-zero, i.e. as positive or negative contingencies. Hence, an effect of scale would be expected for zero contingency conditions because participants cannot easily represent theory beliefs about a negative contingency. It predicts that a 0-100 scale, compared to a -100 to 100 scale, will artificially bias responses in a zero-contingency condition even in the absence of demand characteristics. Some participants will perceive a positive contingency (and will give a non-zero response), others will perceive some negative contingency and are not permitted to give a negative response. Even if the participants feel that they should respond with zero (even though they perceive some contingency) the mean will be above zero. Overall the account is able to explain outcome density effects along with scale and contingency effects, however it cannot explain the other effects such as cue density effects examined in this thesis.

In summary, the contingency account by itself cannot explain the results of this thesis, as it would suggest that at zero contingency differences between contingency judgment at high and low outcome density should not be found. However when the Wasserman (1990) weighting of trials types are also taken into account, the contingency account is better able to explain the results of this thesis. The area in which the contingency account and the Wasserman extension fail to completely predict the findings of this thesis, is where affect is considered. Similarly, neither the Rescorla-Wagner model nor the probabilistic cannot explain all the results of this thesis. Overall, the contingency account with the Wasserman extension has been able to best predict the findings of this thesis.

10.4 Concluding Summary

My interest in this project started with worries around vaccine hesitancy and consequent impact on population health. Further reading and research led me to explore the literature around belief formation and in particular false belief formation. I wanted to further examine the specific factors which led false beliefs, so better interventions could be formed which target false belief formation in the context of fake news. Across 7 experiments, I have found the outcome density to be a robust effect, increased probability of outcome in turn increases the positivity of contingency judgements. This effect can be increased and decreased through various methodological variations, however consistently it seems people give judgements which are more positive than the programmed contingency would indicate. This suggests that people may be pre-disposed to form positive contingencies or relationships between events. Findings of the thesis also show the importance of belief, and how belief can overcome the programmed contingencies, which may have implications for intervention which focus on false belief reduction. The findings of the thesis also suggest an influence of affect on contingency judgements, such that increased levels of affect leads to increased positive judgements. The findings of this PhD work link the fields of affect and beliefs with contingency learning and advance the current knowledge and understanding of contingency judgements. Also, the thesis aimed to examine the current theoretical frameworks in respect to the findings.

References

Ahn, W., & Kalish, C.W. (2000). The Role of Mechanism Beliefs in Causal Reasoning. In *Explanation and cognition* (pp. 199–226). MIT Press.

Akin, A., & Çetın, B. (2007). The Depression Anxiety and Stress Scale (DASS): The study of validity and reliability. *Kuram ve Uygulamada Eğitim Bilimleri*, *7*, 260–268.

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, *15*(3), 147–149. https://doi.org/10.3758/BF03334492

Allan, L. G., & Jenkins, H. M. (1983). The effect of representations of binary variables on judgment of influence. *Learning and Motivation*, *14*(4), 381–405. https://doi.org/10.1016/0023-9690(83)90024-3

Allan, L. G., Siegel, S., & Tangen, J. M. (2005). A signal detection analysis of contingency data. *Learning & Behavior*, *33*(2), 250–263. https://doi.org/10.3758/BF03196067

Allcott, H., & Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, *31*(2), 211–236. https://doi.org/10.1257/jep.31.2.211

Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology. General*, *108*(4), 441–485. https://doi.org/10.1037/0096-3445.108.4.441

Alloy, L. B., & Abramson, L. Y. (1988). Depressive realism: Four theoretical perspectives. In *Cognitive processes in depression* (pp. 223–265). The Guilford Press.

Alloy, L. B., Abramson, L. Y., & Kossman, D. (1985). The Judgment of Predictability in Depressed and Nondepressed College Students. In F. R. Brush, J. Bruce Overmier, & F. R. Brush (Eds.), *Affect, Conditioning, and Cognition (PLE: Emotion)*. Psychology Press.

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, *91*, 112–149. https://doi.org/10.1037/0033-295X.91.1.112

Althubaiti, A. (2016). Information bias in health research: Definition, pitfalls, and adjustment methods. *Journal of Multidisciplinary Healthcare*, 211. https://doi.org/10.2147/JMDH.S104807

American Psychiatric Association. (2022). *Diagnostic and statistical manual of mental disorders: DSM-5^{TM}, 5th ed*. American Psychiatric Publishing, Inc. https://doi.org/10.1176/appi.books.9780890425787

Attribution to Upklyak - Freepik.com. (2022). Set of cannabis production and equipment icons. Freepik.com. https://www.freepik.com/free-vector/set-of-cannabis-production-and-equipmenticons_16793902.htm

Attribution to Vectorjuice - Freepik.com. (2022). Artificial intelligence in social media abstract concept illustration. Freepik.com. https://www.freepik.com/free-vector/artificial-intelligence-in-social-media-abstract-concept-illustration_12291060.htm

Attribution to Pch.vector - Freepik.com. (2022). Cartoon tech team people creating robots in lab. Machine or hardware engineering, scientist with computer, new invention flat vector illustration. Technology, development, science concept for banner. Freepik.com. https://www.freepik.com/free-vector/cartoon-tech-team-people-creating-robots-in-lab-machine-or-hardware-engineering-scientist-with-computer-new-invention-flat-vector-illustration-technology-development-science-concept-for-banner_24644156.htm

Attribution to Vectorjuice - Freepik.com. (2022). Doctor pointing at private healthcare center with medical services. Private healthcare, private medical services, health care center concept. Pinkish coral bluevector vector isolated illustration. Freepik.com. https://www.freepik.com/free-vector/doctor-pointing-at-private-healthcare-center-with-medical-services-private-healthcare-private-medical-services-health-care-center-concept-pinkish-coral-bluevector-vector-isolated-illustration_11666946.htm

Attribution to Felicities - Freepik.com. (2022). Young man character creation design for Animation cartoon flat design. Freepik.com. https://www.freepik.com/free-vector/young-man-character-creation-design-for-animationcartoon-flat-design_14475313.htm

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, *32*(8), 1369–1378. https://doi.org/10.3758/BF03206327

Baker, A. G., Vallée-Tourangeau, F., & Murphy, R. A. (2000). Asymptotic judgment of cause in a relative validity paradigm. *Memory & Cognition*, *28*(3), 466–479. https://doi.org/10.3758/BF03198561

Baker, S. A., & Rojek, C. (2020). The Belle Gibson scandal: The rise of lifestyle gurus as micro-celebrities in low-trust societies. *Journal of Sociology*, *56*(3), 388–404. https://doi.org/10.1177/1440783319846188

Bandlow, R. J. (1976). Theories of Learning, 4th Edition. By Ernest R. Hilgard and Gordon H. Bower. Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1975. *NASSP Bulletin*, *60*(400), 134–134. https://doi.org/10.1177/019263657606040024

Barberia, I., Vadillo, M. A., & Rodríguez-Ferreiro, J. (2019). Persistence of Causal Illusions After Extensive Training. *Frontiers in Psychology*, *10*. https://doi.org/10.3389/fpsyg.2019.00024

Batterink, L., & Neville, H. (2011). Implicit and Explicit Mechanisms of Word Learning in a Narrative Context: An Event-related Potential Study. *Journal of Cognitive Neuroscience*, *23*(11), 3181–3196. https://doi.org/10.1162/jocn_a_00013

Beck, A. T., Ward, C. H., Mendelson, M., Mock, J., & Erbaugh, J. (1961). An inventory for measuring depression. *Archives of General Psychiatry*, *4*, 561–571. https://doi.org/10.1001/archpsyc.1961.01710120031004

Berggren, N., Jordahl, H., & Poutvaara, P. (2017). The right look: Conservative politicians look better and voters reward it. *Journal of Public Economics*, *146*, 79–86. https://doi.org/10.1016/j.jpubeco.2016.12.008

Biele, G., Rieskamp, J., & Gonzalez, R. (2009). Computational Models for the Combination of Advice and Individual Learning. *Cognitive Science*, *33*(2), 206–242. https://doi.org/10.1111/j.1551-6709.2009.01010.x

Blanco, F., Barberia, I., & Matute, H. (2014). The Lack of Side Effects of an Ineffective Treatment Facilitates the Development of a Belief in Its Effectiveness. *PLOS ONE*, *9*(1), e84084. https://doi.org/10.1371/journal.pone.0084084

Blanco, F., Matute, H., & Vadillo, M. A. (2011). Making the Uncontrollable Seem Controllable: The Role of Action in the Illusion of Control. *Quarterly Journal of Experimental Psychology*, *64*(7), 1290–1304. https://doi.org/10.1080/17470218.2011.552727

Blanco, F., Matute, H., & Vadillo, M. A. (2013). Interactive effects of the probability of the cue and the probability of the outcome on the overestimation of null contingency. *Learning & Behavior*, *41*(4), 333–340. https://doi.org/10.3758/s13420-013-0108-8

Bodenhausen, G., Sheppard, L., & Kramer, G. (1994). Negative Affect and Social Judgment: The Differential Impact of Anger and Sadness. *European Journal of Social Psychology*, *24*, 45–62. https://doi.org/10.1002/ejsp.2420240104

Bradley, G. W. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question. *Journal of Personality and Social Psychology*, *36*, 56–71. https://doi.org/10.1037/0022-3514.36.1.56

Bryson, S. E., Doan, B. D., & Pasquali, P. (1984). Sadder but wiser: A failure to demonstrate that mood influences judgements of control. *Canadian Journal of Behavioural Science / Revue Canadienne Des Sciences Du Comportement*, *16*, 107–119. https://doi.org/10.1037/h0080795

Burke, P. F., Dowling, G., & Wei, E. (2018). The relative impact of corporate reputation on consumer choice: Beyond a halo effect. *Journal of Marketing Management*, *34*(13–14), 1227–1257. https://doi.org/10.1080/0267257X.2018.1546765

Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: “Gambler’s fallacy” versus ”hot hand“. *Psychonomic Bulletin & Review*, *11*(1), 179–184. https://doi.org/10.3758/BF03206480

Busch, F. N. (2009). Anger and depression. *Advances in Psychiatric Treatment*, *15*(4), 271–278. https://doi.org/10.1192/apt.bp.107.004937

Busch, F. N., Rudden, M., & Shapiro, T. (2004). *Psychodynamic treatment of depression* (p. 204). American Psychiatric Publishing, Inc.

Buss, A. H., & Perry, M. (1992). The aggression questionnaire. *Journal of Personality and Social Psychology*, *63*(3), 452–459. https://doi.org/10.1037//0022-3514.63.3.452

Byrom, N. C., Msetfi, R. M., & Murphy, R. A. (2015). Two pathways to causal control: Use and availability of information in the environment in people with and without signs of depression. *Acta Psychologica*, *157*, 1–12. https://doi.org/10.1016/j.actpsy.2015.02.004

Camerer, C. (1989). Bubbles and Fads in Asset Prices. *Journal of Economic Surveys*, *3*(1), 3–41. https://doi.org/10.1111/j.1467-6419.1989.tb00056.x

Carr, A. T. (1974). Compulsive neurosis: A review of the literature. *Psychological Bulletin*, *81*, 311–318. https://doi.org/10.1037/h0036473

Carroll, R. (2015, January 17). Too rich to get sick? Disneyland measles outbreak reflects anti-vaccination trend. *The Guardian*. https://www.theguardian.com/us-news/2015/jan/17/too-rich-sick-disneyland-measles-outbreak-reflects-anti-vaccination-trend

Catena, A., Maldonado, A., Perales, J. C., & Cándido, A. (2008). Interaction between previous beliefs and cue predictive value in covariation-based causal induction. *Acta Psychologica*, *128*(2), 339–349. https://doi.org/10.1016/j.actpsy.2008.03.005

Cavus, H. A., & Msetfi, R. M. (2016). The effects of cognitive load during intertrial intervals on judgements of control: The role of working memory and contextual learning. *Acta Psychologica*, *171*, 47–56. https://doi.org/10.1016/j.actpsy.2016.09.005

Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, *74*, 271–280. https://doi.org/10.1037/h0027592

Chen, D. L., Moskowitz, T. J., & Shue, K. (2016). Decision Making Under the Gambler's Fallacy: Evidence from Asylum Judges, Loan Officers, and Baseball Umpires. *The Quarterly Journal of Economics*, *131*(3), 1181–1242. https://doi.org/10.1093/qje/qjw017

Cheng, P. W. (1993). Separating Causal Laws from Casual Facts: Pressing the Limits of Statistical Relevance. In D. L. Medin (Ed.), *Psychology of Learning and Motivation* (Vol. 30, pp. 215–264). Academic Press. https://doi.org/10.1016/S0079-7421(08)60298-4

Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405. https://doi.org/10.1037/0033-295X.104.2.367

Chow, J. Y. L., Colagiuri, B., & Livesey, E. J. (2019). Bridging the divide between causal illusions in the laboratory and the real world: The effects of outcome density with a variable continuous outcome. *Cognitive Research: Principles and Implications*, *4*. https://doi.org/10.1186/s41235-018-0149-9

Claesson, A. (2019). *COMING TOGETHER TO FIGHT FAKE NEWS: LESSONS FROM THE EUROPEAN APPROACH TO DISINFORMATION*. https://csis-website-prod.s3.amazonaws.com/s3fs-public/NewPerspectives_APRIL2019_Claesson.pdf

Clark, R. E. (2003). Fostering the work motivation of individuals and teams. *Performance Improvement*, *42*(3), 21–29. https://doi.org/10.1002/pfi.4930420305

Clore, G. L., & Huntsinger, J. R. (2007). How emotions inform judgment and regulate thought. *Trends in Cognitive Sciences*, *11*(9), 393–399. https://doi.org/10.1016/j.tics.2007.08.005

Clotfelter, C. T., & Cook, P. J. (1993). The 'Gambler's Fallacy' in Lottery Play. *Management Science*, *39*(12), 1521–1525. https://www.jstor.org/stable/2633068

Collins, D. J., & Shanks, D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition*, *30*(7), 1138–1147. https://doi.org/10.3758/BF03194331

Croson, R., & Sundali, J. (2005). The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos. *Journal of Risk and Uncertainty*, *30*(3), 195–209. https://doi.org/10.1007/s11166-005-1153-2

Crump, M. J. C., Hannah, S. D., Allan, L. G., & Hord, L. K. (2007). Contingency judgements on the fly. *Quarterly Journal of Experimental Psychology*, *60*(6), 753–761. https://doi.org/10.1080/17470210701257685

Czupryna, M., Kubińska, E., Markiewicz, Ł., & Tyszka, T. (2021). On the modified measure of illusion of control. *Journal of Behavioral and Experimental Economics*, *90*, 101648. https://doi.org/10.1016/j.socec.2020.101648

Dahl-Monroe, T. (2017). *Giving Beautiful People an Unwarranted Break: Physical Attractiveness and Perceptions of Crime Severity*. /paper/Giving-Beautiful-People-an-Unwarranted-Break-%3A-and-Dahl-Monroe/2ff8969a3fa208dba44f7482b9511dcf17b03f36

Danhauer, S. C., Legault, C., Bandos, H., Kidwell, K., Costantino, J., Vaughan, L., Avis, N. E., Rapp, S., Coker, L. H., Naughton, M., Naylor, C., Terracciano, A., & Shumaker, S. (2013). Positive and negative affect, depression, and cognitive processes in the Cognition in the Study of Tamoxifen and Raloxifene (Co-STAR) Trial. *Neuropsychology, Development, and Cognition. Section B, Aging, Neuropsychology and Cognition*, *20*(5), 10.1080/13825585.2012.747671. https://doi.org/10.1080/13825585.2012.747671

Danielle M. Karazinov. (2018). *Learning about cues that prevent an outcome: Conditioned inhibition and differential inhibition in human predictive learning*. https://journals.sagepub.com/doi/10.1080/02724990344000033

De Houwer, J., & Beckers, T. (2002). A Review of Recent Developments in Research and Theories on Human Contingency Learning. *The Quarterly Journal of Experimental Psychology Section B*, *55*(4b), 289–310. https://doi.org/10.1080/02724990244000034

De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology Section B: Comparative and Physiological Psychology*, *56*(4), 345–357. Scopus. https://doi.org/10.1080/02724990244000296

Dennis, M. J., & Ahn, W. K. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. *Memory & Cognition*, *29*(1), 152–164. https://doi.org/10.3758/bf03195749

Department for Digital, Culture, Media & Sport. (2021). *Minister launches new strategy to fight online disinformation*. GOV.UK. https://www.gov.uk/government/news/minister-launches-new-strategy-to-fight-online-disinformation

Dev, A. S., Moore, D. A., Johnson, S. L., & Garrett, K. T. (2022). Sadder ≠ Wiser: Depressive Realism Is Not Robust to Replication. *Collabra: Psychology*, *8*(1), 38529. https://doi.org/10.1525/collabra.38529

Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, *5*. https://www.frontiersin.org/articles/10.3389/fpsyg.2014.00781

Dienes, Z. (2019). How Do I Know What My Theory Predicts? *Advances in Methods and Practices in Psychological Science*, *2*(4), 364–377. https://doi.org/10.1177/2515245919876960

Dobson, K. S., & Pusch, D. (1995). A test of the depressive realism hypothesis in clinically depressed subjects. *Cognitive Therapy and Research*, *19*(2), 179–194. https://doi.org/10.1007/BF02229693

Dorsey-Palmateer, R., & Smith, G. (2004). Bowlers' Hot Hands. *The American Statistician*, *58*(1), 38–45. https://doi.org/10.1198/0003130042809

Ecker, U. K. H., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L. K., Brashier, N., Kendeou, P., Vraga, E. K., & Amazeen, M. A. (2022). The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, *1*(1), Article 1. https://doi.org/10.1038/s44159-021-00006-y

Eder, A. B., Fiedler, K., & Hamm-Eder, S. (2011). Illusory correlations revisited: The role of pseudocontingencies and working-memory capacity. *Quarterly Journal of Experimental Psychology*, *64*(3), 517–532. https://doi.org/10.1080/17470218.2010.509917

Efran, M. G. (1974). The effect of physical appearance on the judgment of guilt, interpersonal attraction, and severity of recommended punishment in a simulated jury task. *Journal of Research in Personality*, *8*(1), 45–54. https://doi.org/10.1016/0092-6566(74)90044-0

Ernst, H. M., Kuhlmann, B. G., & Vogel, T. (2019). The Origin of Illusory Correlations: Biased Judgments Converge With Inferences, Not With Biased Memory. *Experimental Psychology*, *66*(3), 195–206. https://doi.org/10.1027/1618-3169/a000444

European Commission. (2010). *Special Eurobarometer 340: Science and Technology* (340). European Commission. https://data.europa.eu/data/datasets/s806_73_1_ebs340?locale=en

Evans, J. St. B. T., Clibbens, J., & Harris, A. (2005). Prior belief and polarity in multicue learning. *The Quarterly Journal of Experimental Psychology Section A*, *58*(4), 651–665. https://doi.org/10.1080/02724980443000133

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175–191. https://doi.org/10.3758/BF03193146

Feather, N. T. (1969). Attribution of responsibility and valence of success and failure in relation to initial confidence and task performance. *Journal of Personality and Social Psychology*, *13*, 129–144. https://doi.org/10.1037/h0028071

Fiedler, K. (2000). Illusory correlations: A simple associative algorithm provides a convergent account of seemingly divergent paradigms. *Review of General Psychology*, *4*(1), 25–58. https://doi.org/10.1037/1089-2680.4.1.25

Fiedler, K., & Freytag, P. (2004). Pseudocontingencies’. *Journal of Personality and Social Psychology*, *87*(4), 453–467. https://doi.org/10.1037/0022-3514.87.4.453

Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review*, *116*(1), 187–206. https://doi.org/10.1037/a0014480

Fiedler, K., Kutzner, F., & Vogel, T. (2013). Pseudocontingencies: Logically Unwarranted but Smart Inferences. *Current Directions in Psychological Science*. https://doi.org/10.1177/0963721413480171

Fiedler, K., Walther, E., Freytag, P., & Stryczek, E. (2002). Playing Mating Games in Foreign Cultures: A Conceptual Framework and an Experimental Paradigm for Inductive Trivariate Inference. *Journal of Experimental Social Psychology*, *38*(1), 14–30. https://doi.org/10.1006/jesp.2001.1483

Fischhoff, B., Gonzalez, R. M., Lerner, J. S., & Small, D. A. (2005). Evolving Judgments of Terror Risks: Foresight, Hindsight, and Emotion. *Journal of Experimental Psychology: Applied*, *11*(2), 124–139. https://doi.org/10.1037/1076-898X.11.2.124

Fitch, G. (1970). Effects of self-esteem, perceived performance, and choice on causal attributions. *Journal of Personality and Social Psychology*, *16*, 311–315. https://doi.org/10.1037/h0029847

Foa, E. B., & Kozak, M. J. (1986). Emotional processing of fear: Exposure to corrective information. *Psychological Bulletin*, *99*(1), 20–35.

Fugelsang, J. A., & Thompson, V. A. (2001). Belief-Based and covariation-based cues affect causal discounting. *Canadian Journal of Experimental Psychology / Revue Canadienne de Psychologie Expérimentale*, *55*, 70–76. https://doi.org/10.1037/h0087354

Gibson, J. L., & Gore, J. S. (2016). Is He a Hero or a Weirdo? How Norm Violations Influence the Halo Effect. *Gender Issues*, *33*(4), 299–310. https://doi.org/10.1007/s12147-016-9173-6

Gilden, D. L., & Wilson, S. G. (1995a). On the Nature of Streaks in Signal Detection. *Cognitive Psychology*, *28*(1), 17–64. https://doi.org/10.1006/cogp.1995.1002

Gilden, D. L., & Wilson, S. G. (1995b). Streaks in skilled performance. *Psychonomic Bulletin & Review*, *2*(2), 260–265. https://doi.org/10.3758/BF03210967

Gillan, C. M., Morein-Zamir, S., Durieux, A., Fineberg, N. A., Sahakian, B. J., & Robbins, T. W. (2014a). Obsessive–compulsive disorder patients have a reduced sense of control on the illusion of control task. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00204

Gillan, C. M., Morein-Zamir, S., Durieux, A., Fineberg, N. A., Sahakian, B. J., & Robbins, T. W. (2014b). Obsessive–compulsive disorder patients have a reduced sense of control on the illusion of control task. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00204

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

Gorilla Experiment Builder. (2020). Gorilla. https://gorilla.sc/

Gorton, H. C., & Jarvis, K. (1999). The effectiveness of vitamin C in preventing and relieving the symptoms of virus-induced respiratory infections. *Journal of Manipulative and Physiological Therapeutics*, *22*(8), 530–533. https://doi.org/10.1016/s0161-4754(99)70005-9

Government Communication Service. (2022). *The Wall of Beliefs*. https://gcs.civilservice.gov.uk/wp-content/uploads/2022/09/Wall_of_Beliefs_-publication.pdf

Green, B., & Zwiebel, J. (2017). The Hot-Hand Fallacy: Cognitive Mistakes or Equilibrium Adjustments? Evidence from Major League Baseball. *Management Science*, *64*(11), 5315–5348. https://doi.org/10.1287/mnsc.2017.2804

Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, *82*, 1–20. https://doi.org/10.1037/h0076157

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, *12*(4), 392–407. https://doi.org/10.1016/S0022-1031(76)80006-6

Hamilton, D. L., & Sherman, S. J. (1989). Illusory Correlations: Implications for Stereotype Theory and Research. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotyping and Prejudice: Changing Conceptions* (pp. 59–82). Springer. https://doi.org/10.1007/978-1-4612-3582-8_3

Hannah, S. D., & Beneteau, J. L. (2009). Just tell me what to do: Bringing back experimenter control in active contingency tasks with the command-performance procedure and finding cue density effects along the way. *Canadian Journal of Experimental Psychology = Revue Canadienne De Psychologie Experimentale*, *63*(1), 59–73. https://doi.org/10.1037/a0013403

Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, *135*(4), 555–588. https://doi.org/10.1037/a0015701

Hemilä, H., & Chalker, E. (2013). Vitamin C for preventing and treating the common cold. *The Cochrane Database of Systematic Reviews*, *1*, CD000980. https://doi.org/10.1002/14651858.CD000980.pub4

Hendricks, D., Patel, J., & Zeckhauser, R. (1993). Hot Hands in Mutual Funds: Short-Run Persistence of Relative Performance, 1974–1988. *The Journal of Finance*, *48*(1), 93–130. https://doi.org/10.1111/j.1540-6261.1993.tb04703.x

Hills, T., & Menczer, F. (2020). Information Overload Helps Fake News Spread, and Social Media Knows It. *Scientific American*, *323*(6), 54–61. https://doi.org/10.1038/scientificamerican1220-54

Hornsveld, R. H. J., Muris, P., Kraaimaat, F. W., & Meesters, C. (2009). Psychometric Properties of the Aggression Questionnaire in Dutch Violent Forensic Psychiatric Patients and Secondary Vocational Students. *Assessment*, *16*(2), 181–192. https://doi.org/10.1177/1073191108325894

Iqbal, N., & Dar, K. A. (2015). Negative affectivity, depression, and anxiety: Does rumination mediate the links? *Journal of Affective Disorders*, *181*, 18–23. https://doi.org/10.1016/j.jad.2015.04.002

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs: General and Applied*, *79*(1), 1–17. https://doi.org/10.1037/h0093874

Kapcli, E. G., & Cramer, D. (1999). Judgement of control revisited: Are the depressed realistic or pessimistic? *Counselling Psychology Quarterly*, *12*(1), 95–105. https://doi.org/10.1080/09515079908254081

Kattner, F., & Ellermeier, W. (2011). Does evaluative learning rely on the perception of contingency? Manipulating contingency and US density during evaluative conditioning. *Experimental Psychology*, *5*, 391–399. https://doi.org/10.1027/1618-3169/a000108

Keren, G. B., & Wagenaar, W. A. (1985). On the psychology of playing blackjack: Normative and descriptive considerations with implications for decision theory. *Journal of Experimental Psychology: General*, *114*(2), 133–158. https://doi.org/10.1037/0096-3445.114.2.133

Klauer, K. C., & Meiser, T. (2000). A Source-Monitoring Analysis of Illusory Correlations. *Personality and Social Psychology Bulletin*, *26*(9), 1074–1093. https://doi.org/10.1177/01461672002611005

Landy, D., & Sigall, H. (1974). Beauty is talent: Task evaluation as a function of the performer's physical attractiveness. *Journal of Personality and Social Psychology*, *29*(3), 299–304. https://doi.org/10.1037/h0036018

Lawson, R. G. (1968). *Order of Presentation as a Factor in Jury Persuasion*. 35.

Lee, J. C., & Lovibond, P. F. (2021). Individual differences in causal structures inferred during feature negative learning. *Quarterly Journal of Experimental Psychology*, *74*(1), 150–165. https://doi.org/10.1177/1747021820959286

Lerner, J. S., & Keltner, D. (2001). Fear, anger, and risk. *Journal of Personality and Social Psychology*, *81*, 146–159. https://doi.org/10.1037/0022-3514.81.1.146

Levine, L. J., Prohaska, V., Burgess, S. L., Rice, J. A., & Laulhere, T. M. (2001). Remembering past emotions: The role of current appraisals. *Cognition and Emotion*, *15*(4), 393–417. https://doi.org/10.1080/02699930125955

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest: A Journal of the American Psychological Society*, *13*(3), 106–131. https://doi.org/10.1177/1529100612451018

López, F. J., Shanks, D. R., Almaraz, J., & Fernández, P. (1998). Effects of trial order on contingency judgments: A comparison of associative and probabilistic contrast accounts. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(3), 672–694. https://doi.org/10.1037/0278-7393.24.3.672

Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, *37*(11), 2098–2109. https://doi.org/10.1037/0022-3514.37.11.2098

Lovibond, P. F., & Lovibond, S. H. (1995). The structure of negative emotional states: Comparison of the Depression Anxiety Stress Scales (DASS) with the Beck Depression and Anxiety Inventories. *Behaviour Research and Therapy*, *33*(3), 335–343. https://doi.org/10.1016/0005-7967(94)00075-u

Lovibond, S. H., Lovibond, P. F., & Australia, P. F. of. (1995). *Manual for the depression anxiety stress scales* (2nd ed). Sydney, N.S.W. : Psychology Foundation of Australia. https://trove.nla.gov.au/work/30421447

Luutonen, S. (2007). Anger and depression—Theoretical and clinical considerations. *Nordic Journal of Psychiatry*, *61*(4), 246–251. https://doi.org/10.1080/08039480701414890

Maldonado, A., Catena, A., Cándido, A., & García, I. (1999). The belief revision model: Asymmetrical effects of noncontingency on human covariation learning. *Animal Learning & Behavior*, *27*(2), 168–180. https://doi.org/10.3758/BF03199673

Maldonado, A., Jiménez, G., Herrera, A., Perales, J. C., & Catena, A. (2006). Inattentional blindness for negative relationships in human causal learning. *Quarterly Journal of Experimental Psychology*, *59*(3), 457–470. https://doi.org/10.1080/02724980443000854

Marsh, J. K., & Ahn, W. (2006). Order effects in contingency learning: The role of task complexity. *Memory & Cognition*, *34*(3), 568–576. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2677761/

Marshall, P. H., & Werder, P. R. (1972). The effects of the elimination of rehearsal on primacy and recency. *Journal of Verbal Learning and Verbal Behavior*, *11*(5), 649–653. https://doi.org/10.1016/S0022-5371(72)80049-5

Matías Avellaneda, Joaquín Menéndez, Federico Sánchez, Mateo Santillán, Sebastián Idesis, Victoria Papagna, & Alberto Iorio. (2016). *Equivalence Class Formation is Influenced by Stimulus Contingency*. https://link.springer.com/article/10.1007%2Fs40732-016-0187-y

Matute, H. (1995). Human Reactions to Uncontrollable Outcomes: Further Evidence for Superstitions Rather Than Helplessness. *The Quarterly Journal of Experimental Psychology Section B*, *48*(2), 142–157. https://doi.org/10.1080/14640749508401444

Matute, H. (1996). Illusion of Control: Detecting Response-Outcome Independence in Analytic but Not in Naturalistic Conditions. *Psychological Science*, *7*(5), 289–293. https://doi.org/10.1111/j.1467-9280.1996.tb00376.x

Matute, H., Blanco, F., Yarritu, I., Díaz-Lago, M., Vadillo, M. A., & Barberia, I. (2015). Illusions of causality: How they bias our everyday thinking and how they could be reduced. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00888

Matute, H., Steegen, S., & Vadillo, M. A. (2014). Outcome probability modulates anticipatory behavior to signals that are equally reliable. *Adaptive Behavior*, *22*(3), 207–216. https://doi.org/10.1177/1059712314527005

Matute, H., Vadillo, M. A., Vegas, S., & Blanco, F. (2007a). Illusion of Control in Internet Users and College Students. *CyberPsychology & Behavior*, *10*(2), 176–181. https://doi.org/10.1089/cpb.2006.9971

Matute, H., Vadillo, M. A., Vegas, S., & Blanco, F. (2007b). Illusion of Control in Internet Users and College Students. *CyberPsychology & Behavior*, *10*(2), 176–181. https://doi.org/10.1089/cpb.2006.9971

Matute, H., Yarritu, I., & Vadillo, M. A. (2011a). Illusions of causality at the heart of pseudoscience. *British Journal of Psychology*, *102*(3), 392–405. https://doi.org/10.1348/000712610X532210

Matute, H., Yarritu, I., & Vadillo, M. A. (2011b). Illusions of causality at the heart of pseudoscience. *British Journal of Psychology (London, England: 1953)*, *102*(3), 392–405. https://doi.org/10.1348/000712610X532210

McCullough, M. E., Emmons, R. A., & Tsang, J.-A. (2002). The grateful disposition: A conceptual and empirical topography. *Journal of Personality and Social Psychology*, *82*(1), 112. https://doi.org/10.1037/0022-3514.82.1.112

McFall, M. E., & Wollersheim, J. P. (1979). Obsessive-compulsive neurosis: A cognitive-behavioral formulation and approach to treatment. *Cognitive Therapy and Research*, *3*(4), 333–348. https://doi.org/10.1007/BF01184447

McLaren, S., & Crowe, S. F. (2003). The contribution of perceived control of stressful life events and thought suppression to the symptoms of obsessive–compulsive disorder in both non-clinical and clinical samples. *Journal of Anxiety Disorders*, *17*(4), 389–403. https://doi.org/10.1016/S0887-6185(02)00224-4

Meiser, T. (2003). Effects of processing strategy on episodic memory and contingency learning in group stereotype formation. *Social Cognition*, *21*(2), 121–156. https://doi.org/10.1521/soco.21.2.121.21318

Meiser, T., & Hewstone, M. (2004). Cognitive Processes in Stereotype Formation: The Role of Correct Contingency Learning for Biased Group Judgments. *Journal of Personality and Social Psychology*, *87*(5), 599–614. https://doi.org/10.1037/0022-3514.87.5.599

Michotte, A. (1963). *The perception of causality* (pp. xxii, 424). Basic Books.

Miller, D. T. (1978). What constitutes a self-serving attributional bias? A reply to Bradley. *Journal of Personality and Social Psychology*, *36*, 1221–1223. https://doi.org/10.1037/0022-3514.36.11.1221

Miller, D. T., & Ross, M. (1975). Self-serving biases in the attribution of causality: Fact or fiction? *Psychological Bulletin*, *82*, 213–225. https://doi.org/10.1037/h0076486

Miller, R. R., Barnet, R. C., & Grahame, N. J. (1995). Assessment of the Rescorla-Wagner model. *Psychological Bulletin*, *117*, 363–386. https://doi.org/10.1037/0033-2909.117.3.363

Miller, W. R., & Seligman, M. E. (1975). Depression and learned helplessness in man. *Journal of Abnormal Psychology*, *84*, 228–238. https://doi.org/10.1037/h0076720

Mineka, S., Watson, D., & Clark, L. A. (1998). Comorbidity of Anxiety and Unipolar Mood Disorders. *Annual Review of Psychology*, *49*(1), 377–412. https://doi.org/10.1146/annurev.psych.49.1.377

Moore, F. R., Filippou, D., & Perrett, D. I. (2011). Intelligence and attractiveness in the face: Beyond the attractiveness halo effect. *Journal of Evolutionary Psychology*, *9*(3), 205–217. https://doi.org/10.1556/jep.9.2011.3.2

Moore, M. T., & Fresco, D. M. (2012). Depressive realism: A meta-analytic review. *Clinical Psychology Review*, *32*(6), 496–509. https://doi.org/10.1016/j.cpr.2012.05.004

Moreno-Fernández, M. M., Blanco, F., & Matute, H. (2017). Causal illusions in children when the outcome is frequent. *PLoS ONE*, *12*(9). https://doi.org/10.1371/journal.pone.0184707

Moulding, R., & Kyrios, M. (2006). Anxiety disorders and control related beliefs: The exemplar of Obsessive-Compulsive Disorder (OCD). *Clinical Psychology Review*, *26*(5), 573–583. https://doi.org/10.1016/j.cpr.2006.01.009

Msetfi, R. M., Kornbrot, D. E., Matute, H., & Murphy, R. A. (2015). The relationship between mood state and perceived control in contingency learning: Effects of individualist and collectivist values. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.01430

Msetfi, R. M., Murphy, R. A., & Simpson, J. (2007). Depressive realism and the effect of intertrial interval on judgements of zero, positive, and negative contingencies. *Quarterly Journal of Experimental Psychology*, *60*(3), 461–481. https://doi.org/10.1080/17470210601002595

Msetfi, R. M., Murphy, R. A., Simpson, J., & Kornbrot, D. E. (2005). Depressive realism and outcome density bias in contingency judgments: The effect of the context and intertrial interval. *Journal of Experimental Psychology. General*, *134*(1), 10–22. https://doi.org/10.1037/0096-3445.134.1.10

Msetfi, R. M., Wade, C., & Murphy, R. A. (2013a). Context and Time in Causal Learning: Contingency and Mood Dependent Effects. *PLoS ONE*, *8*(5), e64063. https://doi.org/10.1371/journal.pone.0064063

Msetfi, R. M., Wade, C., & Murphy, R. A. (2013b). Context and Time in Causal Learning: Contingency and Mood Dependent Effects. *PLoS ONE*, *8*(5), e64063. https://doi.org/10.1371/journal.pone.0064063

Murphy, R. A., Schmeer, S., Vallée-Tourangeau, F., Mondragón, E., & Hilton, D. (2011). Making the illusory correlation effect appear and then disappear: The effects of increased learning. *Quarterly Journal of Experimental Psychology*, *64*(1), 24–40. https://doi.org/10.1080/17470218.2010.493615

Musca, S. C., Vadillo, M. A., Blanco, F., & Matute, H. (2010). The role of cue information in the outcome-density effect: Evidence from neural network simulations and a causal learning experiment. *Connection Science*, *22*(2), 177–192. https://doi.org/10.1080/09540091003623797

Mussweiler, T., Englich, B., & Strack, F. (2012). Anchoring Effect. In *Cognitive illusions* (pp. 195–212). Psychology Press.

Mutter, S. A., Strain, L. M., & Plumlee, L. F. (2007). The role of age and prior beliefs in contingency judgment. *Memory & Cognition*, *35*(5), 875–884. https://doi.org/10.3758/BF03193462

Neunaber, D. J., & Wasserman, E. A. (1986). The effects of unidirectional versus bidirectional rating procedures on college students' judgments of response-outcome contingency. *Learning and Motivation*, *17*(2), 162–179. https://doi.org/10.1016/0023-9690(86)90008-1

Nickerson, R. S. (1998). Confirmation Bias: A Ubiquitous Phenomenon in Many Guises. *Review of General Psychology*, *2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

Nisbett, R. E., & Wilson, T. D. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology*, *35*(4), 250–256. https://doi.org/10.1037/0022-3514.35.4.250

Novovic, Z., Kovac, A., Djuric, V., & Biro, M. (2012). Positive and negative affect in illusion of control. *Psihologija*, *45*(4), 395–407. https://doi.org/10.2298/PSI1204395N

Owen, A. M. (2011). The Monte Carlo fallacy. *The Medical Journal of Australia*, *195*(7), 421. https://doi.org/10.5694/mja11.10937

Perales, J. C., Catena, A., Maldonado, A., & Cándido, A. (2007). The role of mechanism and covariation information in causal belief updating. *Cognition*, *105*(3), 704–714. https://doi.org/10.1016/j.cognition.2006.12.003

Perales, J. C., Catena, A., Shanks, D. R., & González, J. A. (2005). Dissociation between judgments and outcome-expectancy measures in covariation learning: A signal detection theory approach. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *31*(5), 1105–1120. https://doi.org/10.1037/0278-7393.31.5.1105

Perales, J., & Catena, A. (2006). Human causal induction: A glimpse at the whole picture. *European Journal of Cognitive Psychology*, *18*(2), 277–320. https://doi.org/10.1080/09541440540000167

Peterson, C. (1980). Recognition of noncontingency. *Journal of Personality and Social Psychology*, *38*, 727–734. https://doi.org/10.1037/0022-3514.38.5.727

Pine, D. S., Cohen, P., Gurley, D., Brook, J., & Ma, Y. (1998). The Risk for Early-Adulthood Anxiety and Depressive Disorders in Adolescents With Anxiety and Depressive Disorders. *Archives of General Psychiatry*, *55*(1), 56–64. https://doi.org/10.1001/archpsyc.55.1.56

Pineño, O., & Miller, R. R. (2007). Comparing associative, statistical, and inferential reasoning accounts of human contingency learning. *Quarterly Journal of Experimental Psychology*, *60*(3), 310–329. https://doi.org/10.1080/17470210601000680

Plante, T. G., & Boccaccini, M. T. (1997). The Santa Clara Strength of Religious Faith Questionnaire. *Pastoral Psychology*, *45*(5), 375–387. https://doi.org/10.1007/BF02230993

Plous, S. (1993). *The Psychology of Judgment and Decision Making*. McGraw-Hill Education.

Poutvaara, P., Jordahl, H., & Berggren, N. (2009). Faces of politicians: Babyfacedness predicts inferred competence but not electoral success. *Journal of Experimental Social Psychology*, *45*(5), 1132–1135. https://doi.org/10.1016/j.jesp.2009.06.007

Prolific. (2020). [Computer software]. Prolific. https://www.prolific.co

Rakow, T., Cheung, N. Y., & Restelli, C. (2020). Losing my loss aversion: The effects of current and past environment on the relative sensitivity to losses and gains. *Psychonomic Bulletin & Review*, *27*(6), 1333–1340. https://doi.org/10.3758/s13423-020-01775-y

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical Conditioning II: Current Research and Theory: Vol. Vol. 2*.

Riley, W. T., Treiber, F. A., & Woods, M. G. (1989). Anger and hostility in depression. *The Journal of Nervous and Mental Disease*, *177*(11), 668–674. https://doi.org/10.1097/00005053-198911000-00002

Rodríguez-Ferreiro, J., & Barberia, I. (2017). The moral foundations of illusory correlation. *PLoS ONE*, *12*(10). https://doi.org/10.1371/journal.pone.0185758

Rosenberg, M. (2006). *Rosenberg Self-Esteem Scale (RSE)*.

Rundus, D. (1980). Maintenance rehearsal and long-term recency. *Memory & Cognition*, *8*(3), 226–230. https://doi.org/10.3758/BF03197610

Safer, M. A., Bonanno, G. A., & Field, N. P. (2001). ‘It was never that bad’: Biased recall of grief and long-term adjustment to the death of a spouse. *Memory*, *9*(3), 195–203. https://doi.org/10.1080/09658210143000065

Sahu, A., Gupta, P., & Chatterjee, B. (2014). Depression is More Than Just Sadness: A Case of Excessive Anger and Its Management in Depression. *Indian Journal of Psychological Medicine*, *36*(1), 77–79. https://doi.org/10.4103/0253-7176.127259

Salkovskis, P. M., Forrester, E., & Richards, C. (1998). Cognitive-behavioural approach to understanding obsessional thinking. *The British Journal of Psychiatry. Supplement*, *35*, 53–63.

Sanbonmatsu, D. M., Shavitt, S., Sherman, S. J., & Roskos-Ewoldsen, D. R. (1987). Illusory correlation in the perception of performance by self or a salient other. *Journal of Experimental Social Psychology*, *23*(6), 518–543. https://doi.org/10.1016/0022-1031(87)90019-9

Sawa, K. (2009). Predictive behavior and causal learning in animals and humans1. *Japanese Psychological Research*, *51*(3), 222–233. https://doi.org/10.1111/j.1468-5884.2009.00396.x

Schaller, M., & O’Brien, M. (1992). ‘Intuitive Analysis of Covariance’ and Group Stereotype Formation. *Personality and Social Psychology Bulletin*, *18*(6), 776–785. https://doi.org/10.1177/0146167292186014

Schwarz, N., Sanna, L. J., Skurnik, I., & Yoon, C. (2007). Metacognitive Experiences and the Intricacies of Setting People Straight: Implications for Debiasing and Public Information Campaigns. In *Advances in Experimental Social Psychology* (Vol. 39, pp. 127–161). Academic Press. https://doi.org/10.1016/S0065-2601(06)39003-X

Serafini, K., Malin-Mayor, B., Nich, C., Hunkele, K., & Carroll, K. M. (2016). Psychometric properties of the Positive and Negative Affect Schedule (PANAS) in a heterogeneous sample of substance users. *The American Journal of Drug and Alcohol Abuse*, *42*(2), 203–212. https://doi.org/10.3109/00952990.2015.1133632

Shanks, D. R. (1993). Human instrumental learning: A critical review of data and theory. *British Journal of Psychology*, *84*(3), 319–354. https://doi.org/10.1111/j.2044-8295.1993.tb02486.x

Shanks, D. R. (2007). Associationism and cognition: Human contingency learning at 25. *Quarterly Journal of Experimental Psychology*, *60*(3), 291–309. https://doi.org/10.1080/17470210601000581

Shanks, D. R. (2010). Learning: From Association to Cognition. *Annual Review of Psychology*, *61*(1), 273–301. https://doi.org/10.1146/annurev.psych.093008.100519

Shanks, D. R., & Dickinson, A. (1988). Associative Accounts of Causality Judgment. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 21, pp. 229–261). Academic Press. https://doi.org/10.1016/S0079-7421(08)60030-4

Shanks, D. R., & Dickinson, A. (1991). Instrumental judgment and performance under variations in action-outcome contingency and contiguity. *Memory & Cognition*, *19*(4), 353–360. https://doi.org/10.3758/BF03197139

Shanks, D. R., Holyoak, K., & Medin, D. L. (1996). *Causal learning* (pp. xii, 435). Academic Press.

Shanks, D. R., & Lopez, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, *24*(4), 511–522. https://doi.org/10.3758/BF03200939

Sharma, P., Murthy, P., & Bharath, M. M. S. (2012). Chemistry, Metabolism, and Toxicology of Cannabis: Clinical Implications. *Iranian Journal of Psychiatry*, *7*(4), 149–156. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3570572/

Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1191–1210. https://doi.org/10.1037/0278-7393.18.6.1191

Sigall, H., & Ostrove, N. (1975). Beautiful but dangerous: Effects of offender attractiveness and nature of the crime on juridic judgment. *Journal of Personality and Social Psychology*, *31*(3), 410–414. https://doi.org/10.1037/h0076472

Skinner, B. F. (1948). 'Superstition' in the pigeon. *Journal of Experimental Psychology*, *38*(2), 168–172. https://doi.org/10.1037/h0055873

Smith, E. E., & Medin, D. L. (1981). *Categories and Concepts:* Harvard University Press.

Sobel, D. M., & Kirkham, N. Z. (2006). Blickets and babies: The development of causal reasoning in toddlers and infants. *Developmental Psychology*, *42*, 1103–1115. https://doi.org/10.1037/0012-1649.42.6.1103

Sowislo, J. F., & Orth, U. (2013). Does low self-esteem predict depression and anxiety? A meta-analysis of longitudinal studies. *Psychological Bulletin*, *139*, 213–240. https://doi.org/10.1037/a0028931

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. https://doi.org/10.1016/j.cogpsych.2005.10.003

Streufert, S., & Streufert, S. C. (1969). Effects of conceptual structure, failure, and success on attribution of causality and interpersonal attitudes. *Journal of Personality and Social Psychology*, *11*, 138–147. https://doi.org/10.1037/h0027047

Suetens, S., Galbo-Jørgensen, C. B., & Tyran, J.-R. (2016). Predicting Lotto Numbers: A Natural Experiment on the Gambler's Fallacy and the Hot-Hand Fallacy. *Journal of the European Economic Association*, *14*(3), 584–607. https://doi.org/10.1111/jeea.12147

Surawski, M. K., & Ossoff, E. P. (2006). The effects of physical and vocal attractiveness on impression formation of politicians. *Current Psychology*, *25*(1), 15–27. https://doi.org/10.1007/s12144-006-1013-5

Suslow, T., Bodenschatz, C. M., Kersting, A., Quirin, M., & Günther, V. (2019). Implicit affectivity in clinically depressed patients during acute illness and recovery. *BMC Psychiatry*, *19*(1), 376. https://doi.org/10.1186/s12888-019-2365-3

Switzer, F. S., & Sniezek, J. A. (1991). Judgment processes in motivation: Anchoring and adjustment effects on judgment and behavior. *Organizational Behavior and Human Decision Processes*, *49*(2), 208–229. https://doi.org/10.1016/0749-5978(91)90049-Y

Talamas, S. N., Mavor, K. I., & Perrett, D. I. (2016). The influence of intelligence on the endorsement of the intelligence–attractiveness halo. *Personality and Individual Differences*, *95*, 162–167. https://doi.org/10.1016/j.paid.2016.02.053

The Digital, Culture, Media and Sport Committee. (2018). *Disinformation and 'fake news': Government Response to the Committee's Fifth Report—Digital, Culture, Media and Sport Committee—House of Commons*. https://publications.parliament.uk/pa/cm201719/cmselect/cmcumeds/1630/163002.htm

Thorndike, E. L. (1920). A constant error in psychological ratings. *Journal of Applied Psychology*, *4*(1), 25–29. https://doi.org/10.1037/h0071663

Turkey earthquake: The false images shared online. (2023, February 8). *BBC News*. https://www.bbc.com/news/64557407

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, *76*(2), 105–110. https://doi.org/10.1037/h0031322

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, *185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

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

Vadillo, M. A., De Houwer, J., De Schryver, M., Ortega-Castro, N., & Matute, H. (2013). Evidence for an illusion of causality when using the Implicit Association Test to measure learning. *Learning and Motivation*, *44*(4), 303–311. https://doi.org/10.1016/j.lmot.2013.04.001

Vadillo, M. A., Musca, S. C., Blanco, F., & Matute, H. (2011a). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin & Review*, *18*(1), 110–115. https://doi.org/10.3758/s13423-010-0032-2

Vadillo, M. A., Musca, S. C., Blanco, F., & Matute, H. (2011b). Contrasting cue-density effects in causal and prediction judgments. *Psychonomic Bulletin & Review*, *18*(1), 110–115. https://doi.org/10.3758/s13423-010-0032-2

Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action–outcome contingency judgements. *Quarterly Journal of Experimental Psychology: Section B*, *58*(2), 177–192. https://doi.org/10.1080/02724990444000104

van Praag, H. M. (2004). Can stress cause depression? *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *28*(5), 891–907. https://doi.org/10.1016/j.pnpbp.2004.05.031

Van Tilburg, W. A. P., & Igou, E. R. (2014). Moving Onwards: An Action Continuation Strategy in Finding the Way. *Journal of Behavioral Decision Making*, *27*(5), 408–418. https://doi.org/10.1002/bdm.1817

van Tilburg, W. A. P., & Igou, E. R. (2014). The impact of middle names: Middle name initials enhance evaluations of intellectual performance. *European Journal of Social Psychology*, *44*(4), 400–411. https://doi.org/10.1002/ejsp.2026

Vance, L., Raciti, M. M., & Lawley, M. (2016). Beyond brand exposure: Measuring the sponsorship halo effect. *Measuring Business Excellence*, *20*(3), 1–14. https://doi.org/10.1108/MBE-07-2015-0037

Venkatesh, S., Moulds, M. L., & Mitchell, C. J. (2018). Testing for Depressive Realism in a Clinically Depressed Sample. *Behaviour Change*, *35*(2), 108–122. https://doi.org/10.1017/bec.2018.12

Verhulst, B., Lodge, M., & Lavine, H. (2010). The Attractiveness Halo: Why Some Candidates are Perceived More Favorably than Others. *Journal of Nonverbal Behavior*, *34*(2), 111–117. https://doi.org/10.1007/s10919-009-0084-z

Vigen, T. (2015). *Spurious Correlations* (Gift edition). Hachette Books.

Vos, T., Abajobir, A. A., Abate, K. H., Abbafati, C., Abbas, K. M., Abd-Allah, F., Abdulkader, R. S., Abdulle, A. M., Abebo, T. A., Abera, S. F., Aboyans, V., Abu-Raddad, L. J., Ackerman, I. N., Adamu, A. A., Adetokunboh, O., Afarideh, M., Afshin, A., Agarwal, S. K., Aggarwal, R., … Murray, C. J. L. (2017). Global, regional, and national incidence, prevalence, and years lived with disability for 328 diseases and injuries for 195 countries, 1990–2016: A systematic analysis for the Global Burden of Disease Study 2016. *The Lancet*, *390*(10100), 1211–1259. https://doi.org/10.1016/S0140-6736(17)32154-2

Wardenaar, K. J., van Veen, T., Giltay, E. J., de Beurs, E., Penninx, B. W. J. H., & Zitman, F. G. (2010). Development and validation of a 30-item short adaptation of the Mood and Anxiety Symptoms Questionnaire (MASQ). *Psychiatry Research*, *179*(1), 101–106. https://doi.org/10.1016/j.psychres.2009.03.005

Wasserman, E. A. (1990). Detecting Response-Outcome Relations: Toward an Understanding of the Causal Texture of the Environment. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 26, pp. 27–82). Academic Press. https://doi.org/10.1016/S0079-7421(08)60051-1

Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *16*(3), 509–521. https://doi.org/10.1037//0278-7393.16.3.509

Wasserman, E. A., Kao, S.-F., Hamme, L. J. V., Katagiri, M., & Young, M. E. (1996). Causation and Association. In D. R. Shanks, D. L. Medin, & K. J. Holyoak (Eds.), *Psychology of Learning and Motivation* (Vol. 34, pp. 207–264). Academic Press. https://doi.org/10.1016/S0079-7421(08)60562-9

Watson, D. (2005). Rethinking the mood and anxiety disorders: A quantitative hierarchical model for DSM-V. *Journal of Abnormal Psychology*, *114*(4), 522–536. https://doi.org/10.1037/0021-843X.114.4.522

Watson, D., & Clark, L. A. (1994). *The PANAS-X: Manual for the Positive and Negative Affect Schedule - Expanded Form* [dataset]. University of Iowa. https://doi.org/10.17077/48vt-m4t2

Weigl, M., Mecklinger, A., & Rosburg, T. (2018). Illusory correlations despite equated category frequencies: A test of the information loss account. *Consciousness and Cognition*, *63*, 11–28. https://doi.org/10.1016/j.concog.2018.06.002

Weiner, B., Frieze, I., Kukla, A., Reed, L., Rest, S., & Rosenbaum, R. (1971). *Perceiving the Causes of Success and Failure*.

Werker, J. F., Cohen, L. B., Lloyd, V. L., Casasola, M., & Stager, C. L. (1998). Acquisition of word–object associations by 14-month-old infants. *Developmental Psychology*, *34*(6), 1289–1309. https://doi.org/10.1037/0012-1649.34.6.1289

Wheatley, D. (1997). Stress, Anxiety and Depression. *Stress Medicine*, *13*(3), 173–177. https://doi.org/10.1002/(SICI)1099-1700(199707)13:3<173::AID-SMI739>3.0.CO;2-6

White, P. A. (1995). Use of prior beliefs in the assignment of causal roles: Causal powers versus regularity-based accounts. *Memory & Cognition*, *23*(2), 243–254. Scopus. https://doi.org/10.3758/BF03197225

Williams, D. A. (1995). Do Rats and People Judge Negative Contingencies in the Same Way? *Current Directions in Psychological Science*, *4*(6), 207–211. https://doi.org/10.1111/1467-8721.ep10772650

Wittchen, H. U., Jacobi, F., Rehm, J., Gustavsson, A., Svensson, M., Jönsson, B., Olesen, J., Allgulander, C., Alonso, J., Faravelli, C., Fratiglioni, L., Jennum, P., Lieb, R., Maercker, A., van Os, J., Preisig, M., Salvador-Carulla, L., Simon, R., & Steinhausen, H.-C. (2011). The size and burden of mental disorders and other disorders of the brain in Europe 2010. *European*

Neuropsychopharmacology, *21*(9), 655–679. https://doi.org/10.1016/j.euroneuro.2011.07.018

Wortman, C. B., Costanzo, P. R., & Witt, T. R. (1973). Effect of anticipated performance on the attributions of causality to self and others. *Journal of Personality and Social Psychology*, *27*(3), 372–381. https://doi.org/10.1037/h0034949

Yarritu, I., & Matute, H. (2015). Previous knowledge can induce an illusion of causality through actively biasing behavior. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.00389

Yarritu, I., Matute, H., & Luque, D. (2015). The dark side of cognitive illusions: When an illusory belief interferes with the acquisition of evidence-based knowledge. *British Journal of Psychology*, *106*(4), 597–608. https://doi.org/10.1111/bjop.12119

Yarritu, I., Matute, H., & Vadillo, M. A. (2014). Illusion of Control. *Experimental Psychology*, *61*(1), 38–47. https://doi.org/10.1027/1618-3169/a000225

Yates, J. F., & Curley, S. P. (1986). Contingency judgment: Primacy effects and attention decrement. *Acta Psychologica*, *62*(3), 293–302. https://doi.org/10.1016/0001-6918(86)90092-2

Appendices

Appendix A

The following images show the cover stories used in Experiment 6 along with the belief questions asked in the experiment.

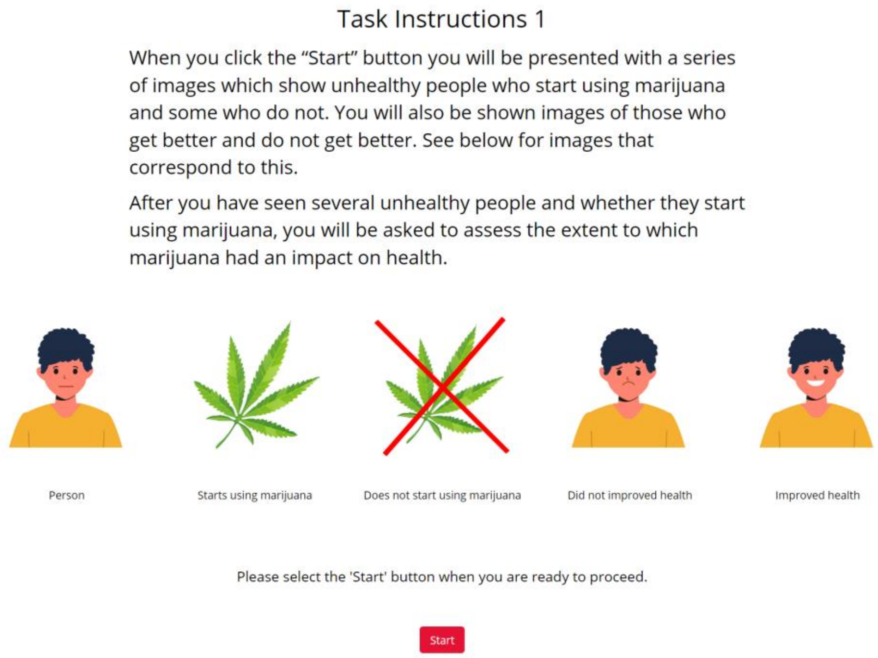

Note. Attribution to Upklyak - Freepik.com. (2022). Set of cannabis production and equipment icons. Freepik.com. (https://www.freepik.com/free-vector/set-of-cannabis-production-and-equipmenticons_16793902.htm). CC BY 2.0.

Note. Attribution to Felicities - Freepik.com. (2022). Young man character creation design for Animation cartoon flat design. Freepik.com. (https://www.freepik.com/free-vector/young-man-character-creation-design-for-animationcartoon-flat-design_14475313.htm). CC BY 2.0.

On the scale below, please indicate to what extent you think that using marijuana had an impact on health?

-100 | 100

It worsened health in all cases | It improved health in all cases

Next

Task Instructions 1

When you click the "Start" button you will be presented with a series of images which show people, some who start using social media and some who do not. You will also be shown images of how this affects their sociability. See below for images that correspond to this.

After you have seen several people and whether they start using a new social media platform, you will be asked to assess the extent to which the use of social media had an impact on sociability.

Normal | Starts using a new social media platform | Does not start usign new a social media platform | Did not increased sociability | Increase sociability

Please select the 'Start' button when you are ready to proceed.

Start

Note. Attribution to Felicities - Freepik.com. (2022). Young man character creation design for Animation cartoon flat design. Freepik.com. (https://www.freepik.com/free-vector/young-man-character-creation-design-for-animationcartoon-flat-design_14475313.htm). CC BY 2.0.

Note. Attribution to Vectorjuice - Freepik.com. (2022). Artificial intelligence in social media abstract concept illustration. Freepik.com. (https://www.freepik.com/free-vector/artificial-intelligence-in-social-media-abstract-concept-illustration_12291060.htm). CC BY 2.0.

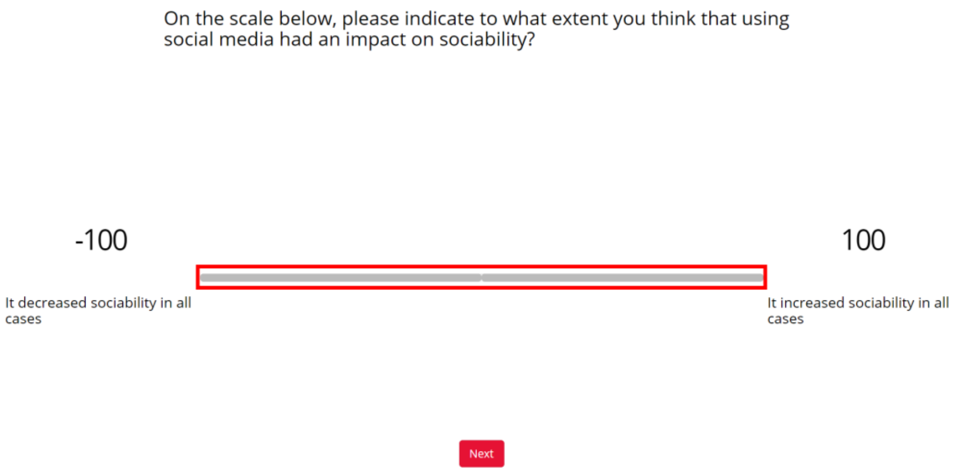

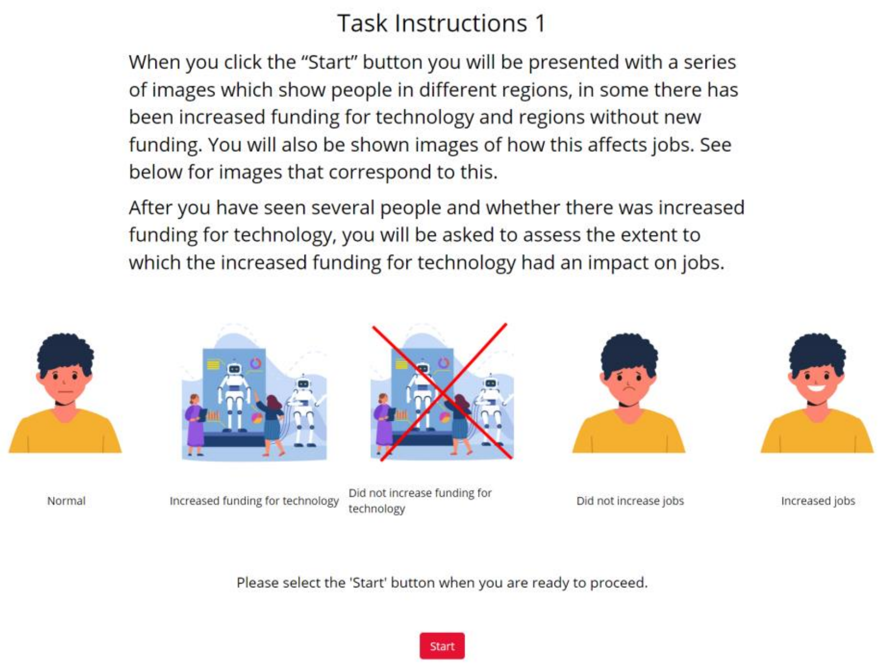

Note. Attribution to Pch.vector - Freepik.com. (2022). Cartoon tech team people creating robots in lab. Machine or hardware engineering, scientist with computer, new invention flat vector illustration. Technology, development, science concept for banner. Freepik.com. (https://www.freepik.com/free-vector/cartoon-tech-team-people-creating-robots-in-lab-machine-or-hardware-engineering-scientist-with-computer-new-invention-flat-vector-illustration-technology-development-science-concept-for-banner_24644156.htm). CC BY 2.0.

Note. Attribution to Felicities - Freepik.com. (2022). Young man character creation design for Animation cartoon flat design. Freepik.com. (https://www.freepik.com/free-vector/young-man-character-creation-design-for-animationcartoon-flat-design_14475313.htm). CC BY 2.0.

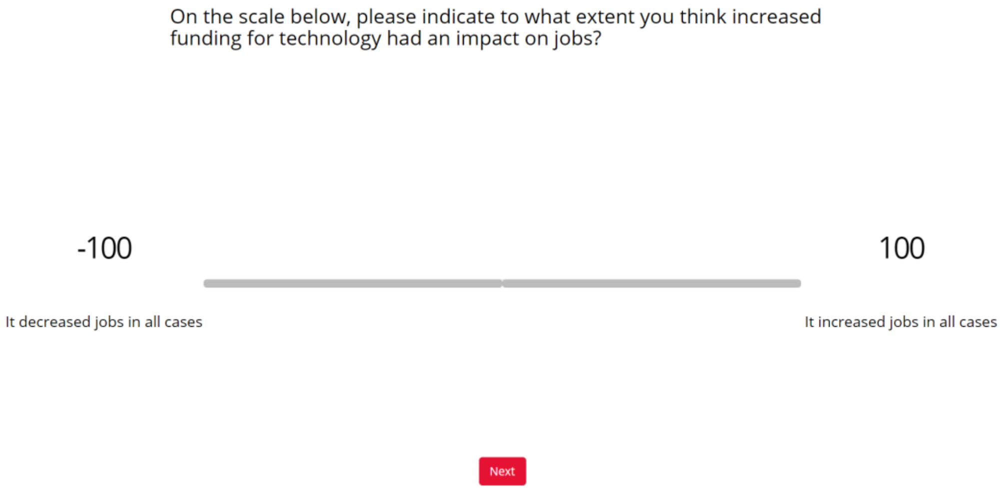

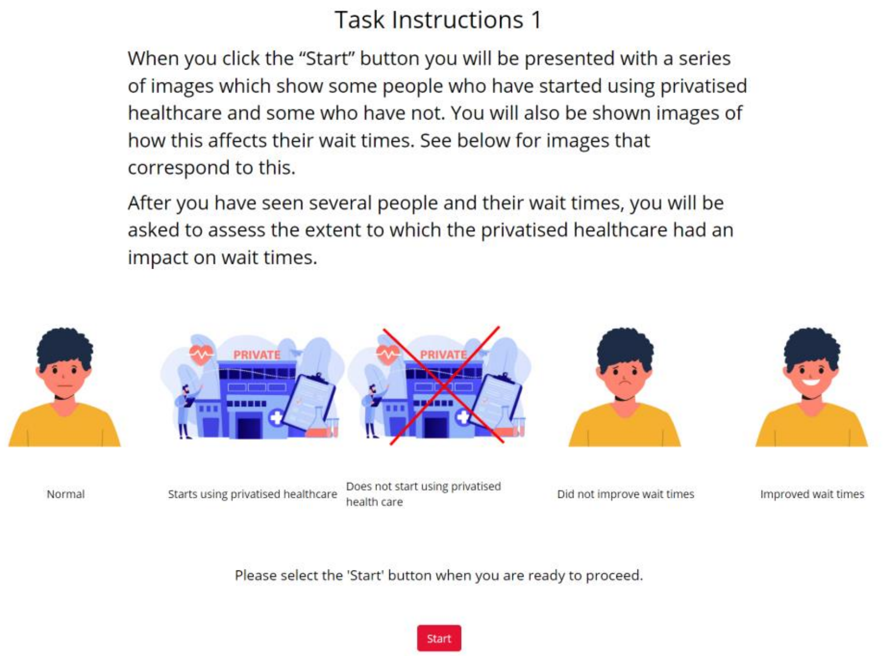

Note. Attribution to Vectorjuice - Freepik.com. (2022). Doctor pointing at private healthcare center with medical services. Private healthcare, private medical services, health care center concept. Pinkish coral bluevector vector isolated illustration. Freepik.com. (https://www.freepik.com/free-vector/doctor-pointing-at-private-healthcare-center-with-medical-services-private-healthcare-private-medical-services-health-care-center-concept-pinkish-coral-bluevector-vector-isolated-illustration_11666946.htm). CC BY 2.0.

Note. Attribution to Felicities - Freepik.com. (2022). Young man character creation design for Animation cartoon flat design. Freepik.com. (https://www.freepik.com/free-vector/young-man-character-creation-design-for-animationcartoon-flat-design_14475313.htm). CC BY 2.0.

On the scale below, please indicate to what extent you think that privatised healthcare had an impact on wait times?

-100 100

It worsened wait times in all cases

It improved wait times in all cases

Next

You will be asked a series of questions about different topics. In these questions please click on the statement that you agree with the most.

Marijuana use has...

Many health benefits	Some health benefits	Limited health benefits	Neither benefits nor harms	Limited health harms	Some health harms	Many health harms

Privatising health care can...wait times

Often improve	Sometimes improve	Occasionally improve	Neither improve or worsen	Occasionally worsen	Sometimes worsen	Often worsen

Social media use can... sociability

Often increase	Sometimes increase	Occasionally increase	Neither increase nor decrease	Occasionally decrease	Sometimes decrease	Often decrease

New technology can ... jobs

Often increase	Sometimes increase	Occasionally increase	Neither increase nor decrease	Occasionally decrease	Sometimes decrease	Often decrease

Next