UNIVERSIDAD DE DEUSTO

FACULTAD DE PSICOLOGÍA Y EDUCACIÓN

DPTO. DE FUNDAMENTOS Y MÉTODOS DE LA PSICOLOGÍA

PhD in Experimental Psychology

In Search of Rationality in Human Causal Learning

-

Evidence of non-normative strategies when combining causes

Thesis by NEREA ORTEGA-CASTRO

Supervisor: Dr. MIGUEL A. VADILLO

Co-supervisor: Dr. ANDY G. BAKER

Bilbao, August of 2014

UNIVERSIDAD DE DEUSTO

FACULTAD DE PSICOLOGÍA Y EDUCACIÓN

DPTO. DE FUNDAMENTOS Y MÉTODOS DE LA PSICOLOGÍA

PhD in Experimental Psychology

In Search of Rationality in Human Causal Learning

-

Evidence of non-normative strategies when combining causes

Author Supervisor Co-supervisor

Bilbao, August of 2014

A mis padres…

'I admit that I feel inclined to protest about certain exaggerations [...] of the irrationality of man'

Karl Popper, 1963 (p. 480)

This thesis was supported by a predoctoral fellowship from the Basque Government (Ref: BFI09.102, 2010 – 2013). The research team that I have joined, led by Dr. Helena Matute, was funded during the performing of this thesis by Dirección General de Investigación Científica y Técnica del Ministerio de la Ciencia y la Investigación (Ref: PSI2011-26965), and by Departamento de Educación, Universidades e Investigación del Gobierno Vasco (Ref: IT363-10).

Agradecimientos

Lo primero de todo, quiero agradecer a mi director, Miguel Ángel Vadillo, por permitirme el honor de ser tu primera doctoranda. Por la confianza y cercanía mostrada durante todos estos años. Por enseñarme qué es ser un buen investigador, y sobre todo por tu amistad. La sensación de pérdida que supuso tu partida a la UCL lo resume todo. Mil veces gracias.

Andy Baker, for giving me the opportunity to work in his lab and learn from him. For all the inspiring conversations and, above all, for becoming my co-supervisor in the midway of my thesis, giving his time and knowledge to a person who, in principle, was going to be there only for 2 months. Knowing you have greatly enriched me, both at research and personal level. Thank you very much.

Muchas gracias a todos los que han costeado esta tesis con sus impuestos.

Estoy también muy agradecida a...

Helena Matute, gracias por darme esta oportunidad y por haber apostado por mí con ese primer contrato. Sin tu ayuda, nada de lo que está escrito aquí podría haberse conseguido. Te estoy muy agradecida.

La gente del equipo de Málaga, mencionar en especial a Pedro, Paco, David, Joaquín y Amanda. Gracias por la cercanía, la amabilidad, y por estar siempre dispuestos a ayudar con cualquier cosa. Eso no tiene precio.

The team from Oxford University, among all the great people I met there I want to make special mention of Robin A. Murphy and Nicola Byrom. It is to them I owe the longing and not the place. Thanks for your warm welcome.

Mis compañeros de laboratorio. Qué decir, habéis sido como mi familia durante todo este recorrido. Al menos os he visto más que a nadie ;) Recuerdo vuestro apoyo el día de la entrevista con el Gobierno Vasco, las cañas que nos tomamos en el Dubliners, la sensación de unión y de formar parte de un equipo. Sois los mejores compañeros que una persona podría desear. Ion Yarritu, compañero de fatigas, no puedo imaginarme estando en un laboratorio donde no estés tú. Carmelo Pérez Cubillas, la persona más generosa que conozco, siempre dispuesto a ayudar. Estar contigo hace mejor la vida, a pesar de tu percusión continua. Espero y deseo que el destino vuelva a reunirnos. Itsaso Barberia, gracias por hacerme sentir como en casa a 6000 kilómetros de casa. Eres un ejemplo a seguir, en todos los sentidos. Fernando Blanco, por todas las conversaciones inspiradoras mantenidas (excepto las bayesianas), y las que aún quedan. Txipi, por dotar no sólo de sentido tecnológico, sino humano, cada conversación. Unai, por demostrarme que siempre se puede ser aún más friki. Sara, por entendernos con una sola mirada incluso sin hablar el mismo idioma. Marcos, aunque haya sido breve, ha sido muy bueno. Y al resto de la gente que ha ido pasando por Labpsico, gracias por todo.

Mis amigos, Alberto, Pablo, Juanjo, Vanesa, Enrique, Txema, Zigor. Gracias por tantos años aguantando mis “chapas” sin quejaros ni una sola vez.

Tiene mucho mérito y lo aprecio profundamente. Habéis sido mi válvula de escape, a quien recurría cada vez que necesitaba desconectar.

Mi familia. Gracias en primer lugar a mis padres, Rosa Mari y Andoni, por haber sido siempre quienes más han confiado en mi valía. Por apoyarme, instruirme, y estar siempre ahí. Por haberme convertido en una persona de provecho, aunque haya costado un poco. Este logro es totalmente vuestro. Agradecer también a mi hermano Aitor, por aguantar mis ensayos de comunicaciones orales con estoicismo, entereza y una sonrisa, siempre dispuesto.

Y por supuesto, agradecer a mi compañero, Jon, por muchas razones, pero sobre todo por amor.

¡Gracias a todos!

Abstract

Causal learning models make different assumptions about how people should combine the influence of different potential causes presented in combination. Based on the linear integration rule, some models propose that the causal impact of a compound should equal the linear sum of each of the causes presented in isolation. Other models such as the Power PC theory are based on a different integration rule, the noisy-OR, suggesting that the rational way of computing the causal impact of a compound involves correcting the sum of the causes by subtracting the overlap between them. The present experiments tested which integration rule people use. Four different cover stories were used to ensure that the participants understood the independence of the causes. The experiments used different sets of probabilities and several formats for presenting information. The results of most experiments do not confirm the predictions of the noisy-OR integration rule. Only one experiment (of ten) supports the predictions of the noisy-OR rule. In spite of having mixed evidence, people do not appear to spontaneously use this rule. We discuss the implications of our results and alternative explanations for our pattern of data, including inhibitory mechanisms and an averaging heuristic.

Resumen

Los modelos de aprendizaje causal parten de diferentes supuestos acerca de cómo la gente debería combinar la influencia de diferentes causas potenciales cuando se presentan en compuesto. Algunos modelos parten de la regla de integración lineal, proponiendo que el impacto causal de un compuesto debe ser igual a la suma lineal de cada una de las causas presentadas en solitario. Otros modelos como Power PC están basados en una regla de integración diferente, noisy-OR, que sugiere que la manera racional de computar el impacto causal de un compuesto requiere corregir la suma lineal de las causas restando la superposición entre ambas. Los experimentos de esta tesis pusieron a prueba qué regla de integración usa la gente. Se utilizaron cuatro escenarios causales para asegurar que los participasen percibiesen las causas como independientes. Los experimentos constaban de diferentes sets de probabilidades y de diversos formatos a la hora de presentar la información. Los resultados de la mayoría de los experimentos no confirman las predicciones de la regla de integración noisy-OR. Sólo un experimento (entre diez) apoya dichas predicciones. A pesar de tener evidencia contradictoria, la gente no parece usar de manera espontánea esta regla de integración. Se discuten las implicaciones de nuestros resultados y explicaciones alternativas a nuestro patrón de datos, incluyendo mecanismos de carácter inhibitorio y el posible uso de un heurístico de promedio.

INDEX

Chapter 1. Introduction

'How dare we speak of the laws of chance? Is not chance the antithesis of all law?'

Joseph Bertrand

What do we mean when we talk about causal knowledge? I will try to explain what the term "causal knowledge" means, and how we extract knowledge about each present and absent cause of a final result. Frequently, the best way to check how people learn about causes, consequences and their interactions is to design an experiment in which all the information available is fully controlled. As the reader will soon discover, all the experiments that will be described in this thesis have been, to a greater or lesser extent, influenced by passion for science fiction. Therefore, I will use one of the top novels from this genre to explain what causal knowledge is and how in certain situations having a precise causal knowledge can be a matter of life or death.

Hyperion (Simmons, 1989) is the first novel in the *Hyperion Cantos* tetralogy. The author, following a narrative structure similar to *The Canterbury*

Tales by Geoffrey Chaucer and Boccaccio's *Decameron*, presents a future where humanity has colonized dozens of worlds in the galaxy and dominates the teleportation and space travel at supraluminal velocities.

In the distant world called Hyperion, beyond the Hegemony of Man network, await the Time Tombs, artifacts sent from the far future which are about to open, and hide a terrible secret related to an impossible creature that only communicates through death. The Shrike, or Lord of Pain, is considered by some as a deity and by others as an avatar of the impending human atonement. Also, external creatures known as 'exters', descendants of the old Earth, have shown the same interest in this strange world. For unknown reasons, all powerful groups in the known universe want to take control of the planet at the moment that the Time Tombs will be opened. On the eve of Armageddon, and against the backdrop of a possible war between the Hegemony, exter swarms, and artificial intelligences from the 'TechnoCore', seven pilgrims undertake the route to Hyperion to resurrect an ancient religious ritual, and here is where the story really begins.

They all carry impossible hopes and terrible secrets. A diplomat, a priest, a soldier, a poet, a teacher, a detective, and a navigator entwine their lives and their destinies in this journey in search of the Shrike and the Time Tombs. The Shrike granted one desire to only one of the pilgrims. The rest shall surely die, in terribly painful ways. Their personal stories compose a kaleidoscopic and evocative vision of the complex society in which they live and which, perhaps, can save.

How do the pilgrims manage to survive? It is their first trip to Hyperion, they have no prior knowledge of its dangers. They do not know what things are dangerous, and what beneficial, or if suffering one of the causes prevents them from suffering others. They do not know if there is any way to offset the effects of certain things. They know nothing. *They have no causal knowledge, yet.* But they need it to succeed. They are aware at all times that survival depends directly on the knowledge that they are able to extract about the environment. The more precise and accurate the knowledge, the more likely they are to remain alive. There are a number of factors that can be potentially dangerous, the pilgrims need to learn as much as possible about all the elements they are likely to be exposed to, and use this knowledge to develop a strategy that allows them to maximize their chances of survival. This is the causal knowledge, and it is important because it ensures adaptation to the environment and, moreover, for survival when life or death depends on knowing what causes and consequences are connected. I will not describe more about the end of the book, and who among pilgrims was able to survive, if any did, I refer the reader to read and enjoy this work by Simmons.

Causal reasoning, reasoning about causes and consequences, represents one of the most basic but important cognitive processes that underpin all high-order activities (Jonassen & Ionas, 2008). The ability to detect the contingencies (or relations) between the events in the environment is central to most types of behavior, including learning, causal judgment, categorization, problem solving, and hypothesis testing (Crocker, 1981). This capacity is critical for planning, acting and reasoning (Spellman, 1997; Buehner & Cheng, 2005), and allows us to develop adaptive behavior, through the prediction and control of the events in our lives (Tolman & Brunswik, 1935). Without this ability, our potential to adapt would be severely compromised. How could any cognitive system learn causal relations from simple associations of events? How do people do it daily? There is, obviously nothing new in these questions, they have been the focus of philosophical accounts since Aristotle, and of many great modern philosophers including Descartes, Hume (1739), and Kant (1781).

Causal learning and probabilistic reasoning are essential to predict how a system will behave. Often we infer causal relations on the basis of probabilistic data, but not every correlation indicates a direct causal relation. Conceptions of probability have been around for thousands of years (Stigler, 1990; Franklin, 2001; Hald, 2003; Hacking, 2006), but probability theory did not arise as a branch of mathematics until the mid-seventeenth century. Luca Pacioli (1494) authored the first printed work on probability, but the mathematical theory of probability has its roots in attempts to analyze games of chance by Gerolamo Cardano. His gambling led him to formulate elementary probability rules, making him one of the founders of the field (Ore, 1953). In 17th century French society, gambling was popular and fashionable, and not restricted by law. As the games became more complicated there was a need for mathematical methods for computing chances. A well-known gambler, Antoine Gombaud, the chevalier De Mere, consulted Blaise Pascal in Paris about some games of chance, and Pascal began to correspond with his friend Pierre Fermat about these problems.

The correspondence between Pascal and Fermat is the origin of the mathematical study of probability (Devlin, 2008). The method they developed is now called the classical approach to computing probabilities. Through the 18^{th} century, the application of probability theory expanded from games of chance to scientific problems. The best exemplar of applied probability theory was Laplace's (1812) book. Many workers have contributed to the theory since Laplace's time, among the most important are Chebyshev, Markov, von Mises, and Kolmogorov. Andrey Kolmogorov (1933) developed the first rigurous approach to probability, where he built up probability theory from fundamental axioms in a way comparable with Euclid's treatment of geometry (for a review, see Heath, 1956).

Probabilistic reasoning involves estimating the probability of the occurrence of an event based on some knowledge. It is critical to understand the relationships between the significant events in the environment in order to exert the higher possible degree of control over them, and inevitably, chance and random phenomena permeate our lives and our environment (Bennet, 1998). Nonetheless we may not be well suited to accurate do this type of reasoning. In the words of Persi Diaconis (1989): 'Our brains are just not wired to do probability problems very well'. Moreover, many recent psychological studies have provided evidence that people have trouble with probability (e.g., Lewens, 2007; Nickerson, 2008; Taynes, 2003). It is not just the general population that has difficulty with numerical tasks, studies have shown that even highly educated laypersons and health professionals have an inadequate understanding of probabilities and other chance-related concepts (Estrada, Barnes, Collins & Byrd, 1999; Lipkus, Sansa, & Rimer, 2001; Nelson, Reyna, Fagerlin, Lipkus, & Peters, 2008; Reyna, Lloyd, & Whalen, 2001; Sheridan & Pignone, 2002). There is no *a priory* reason to expect correct, normative solutions for dealing with numbers, frequencies and probabilities in causal learning. Perhaps there is no domain in which such a normative solution exists. As outlined below, many theories describe causal reasoning as a particular type of probabilistic reasoning. Or, at least, propose that probabilistic reasoning is part of the processes necessary for causal reasoning. Therefore, the major debates about the rationality of probabilistic thinking are relevant to the subject of this thesis.

To give a proper explanation of how humans deal with probabilistic causal learning, we must first take into account a number of issues. First, how do we measure causal learning in the laboratory? In most modern studies of human contingency learning, participants receive information about situations in which certain cues (the potential causes in our Hyperion example) and outcomes (the probability of dying or living) are either present or absent, and are asked to evaluate to what extent the presence of a given cause is related to the presence of the outcome (Allan, 1980; Dickinson, Shanks, & Evenden, 1984; Baker, Berbrier, & Vallée-Tourangeau, 1989; De Houwer & Beckers, 2002, Treisman, 1998; Buehner, Cheng, & Clifford, 2003). At the procedural level, these studies of contingency learning are very similar to those of Pavlovian conditioning: Stimuli (cues and outcomes) are paired in a certain way, and the resulting changes in the responses to the stimuli (contingency judgments) are assessed. It is therefore not surprising that many of the causal learning models that will be explained later have their roots in the same models that have been used to explain all kinds of associative learning, including classical and instrumental conditioning in animals (Rescorla and Wagner, 1972; Dickinson, 1980; Mackintosh, 1974).

Chapter 2. Causal Integration Rules

The present work focuses on how people combine the influence of several causes of an outcome (or effect). Although at first glance it might seem that acquiring causal knowledge and using this causal knowledge are different processes, the fact is that they are closely related. Theories that attempt to explain how learning occurs also contain explicit or implicit assumptions not only about how we learn but also how we combine our knowledge about different causative factors. To compute the causal strength of a possible candidate cause, most causal induction models subtract the probability that the effect occurs when the cause is absent from the probability of the effect when the cause is present (Rescorla & Wagner, 1972; Cheng & Novick, 1990; Cheng, 1997). The models differ in how this subtraction should be performed, depending on different assumptions about how the causal impact of multiple causes should be combined. When we know that A is a potential cause, and B is another potential cause, then we can ask what would be the combined effect of A and B based on the different learning models at our disposal. The simplest assumption about how several causes presented together determine the probability of the effect is that the impact of each cause should be computed by simple linear addition. But there are alternative and arguably more normative assumptions.

This notion can be illustrated with an example that will be used in several of my experiments. Imagine that scientists have discovered a number of chemical substances that can change one's eye colour to pink. Each substance has a different causal force, that is, a different probability of achieving the effect (changing the eye colour). For commercial reasons, drug companies sell combinations of these substances. We are informed that the substance A has a 40% probability of achieving the effect, and that the substance B also has a probability of 40% of achieving the effect. The product that is marketed has a combination of substances A and B. How likely is the product to change the eye color to pink? Following the simple linear addition rule in this scenario, if we know that cause A produces an outcome 40% of the time and B produces the same outcome 40% of the time, when both are present (and assuming that no other effective cause is present) the probability of the effect should be equal to 80%. This is the law for combining mutually exclusive probabilities (Dieks, González, Hartmann, Stöltzner, & Weber, 2012). Following this reasoning, isolating the role of a single cue from the effect of a larger collection of alternative causes is based on a simple linear subtraction. As I will show later, this assumption is made in many formal models of causal learning (Allan, 1980; Cheng & Novick, 1992; Jenkins & Ward, 1965; Rescorla & Wagner, 1972). The linear integration rule is a very simple strategy to combine the effect of several causes on an outcome. But this does not necessarily mean that it is the correct strategy. In other words, it does not mean that it is *normative,* but it can be seen as normative. By using a simple coin tossing game we can show how this solution can be non-normative.

Imagine a game in which two coins (i.e., candidate causes, Coin 1 and Coin 2) are tossed. To win, you need at least one of the tosses (or both) come up heads. Thus, a head on either coin means that you win. If both coins are tossed together, either or both of them might come up heads and hence meet the conditions for generating the desired outcome of at least one head. If the coins are unbiased the probability of getting heads tossing Coin 1 would be 50% and the probability of getting heads tossing Coin 2 would also be 50%. Using the linear summation rule, the probability of the outcome (getting heads) would be 50% + 50%= 100%. This illustrates the problem. We have all thrown coins many times, enough times to know that tossing two coins at the same time does not guarantee that we will get at least one head. The linear integration rule leads to an incorrect result; this is easily appreciated by studying the sequence of possible outcomes depicted below:

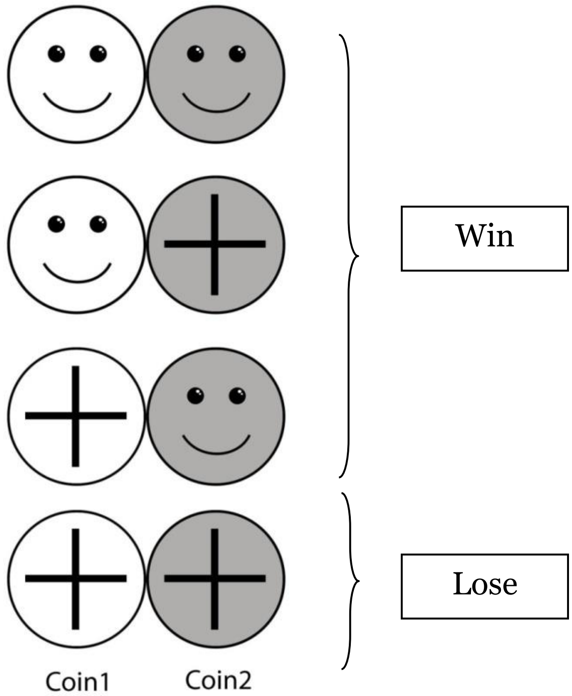

Figure 1. Four permutations when both coins are tossed. Smilies denote heads, and crosses denote tails.

Figure 1 shows that when both coins are tossed, there are only four possible permutations of heads and tails: both coins are heads; the first is heads and the second tails; the first is tails and the second heads; or both are tails. In the first three cases the winning condition is achieved because each includes at least one head but in the final case the outcome does not happen – you lose. If only one of the coins is tossed, we win in 1/2 of the tosses (50% of the time). But if Coin 2 is tossed along with Coin 1 then we win on 3⁄4 of the tosses (75% of the time). If we follow the logic of the linear integration rule assuming 75% wins, the contribution of tossing Coin 2 to winning is the linear difference between the probability of winning if Coin 1 is tossed (50%) and the probability of winning when both are tossed (75%). Thus, by the linear integration rule, we would have to deduce that the power of Coin 2 is 25%. However we know that actual likelihood of an outcome given the tossing of Coin 2 is 50%. This is one of the most important problems of the linear integration rule assuming independent probabilities: If a coin is tossed first its effectiveness is 50%; if it is tossed second it is 25%. When extrapolating the results of a causal inference to a novel context, the linear integration rule sometimes leads to problems of coherence, as some authors have pointed out (Cheng, Novick, Liljeholm, & Ford, 2007; Liljeholm & Cheng, 2007).

Patricia Cheng (1997) suggested that the combined causal power of several causes should be computed by means of a different integration rule, the noisy-OR rule integration rule. This integration rule is based on the law of independent probabilities. It is clear from our coin tossing example that when the second coin is tossed, at least part of its effectiveness is masked because sometime both it and the first coin meet the conditions for an outcome. An Euler diagram illustrating this is shown in Figure 2. It depicts the assumptions made by the noisy-OR integration rule. The plus signs (+) denotes the occurrence of the effect, and the minus signs (-) the absence of the effect. When the two causes, A and B, are present, some instances of the effect are due to A, some to B, and some others to the joint effect of A and B, thus, fall into the intersection between them.

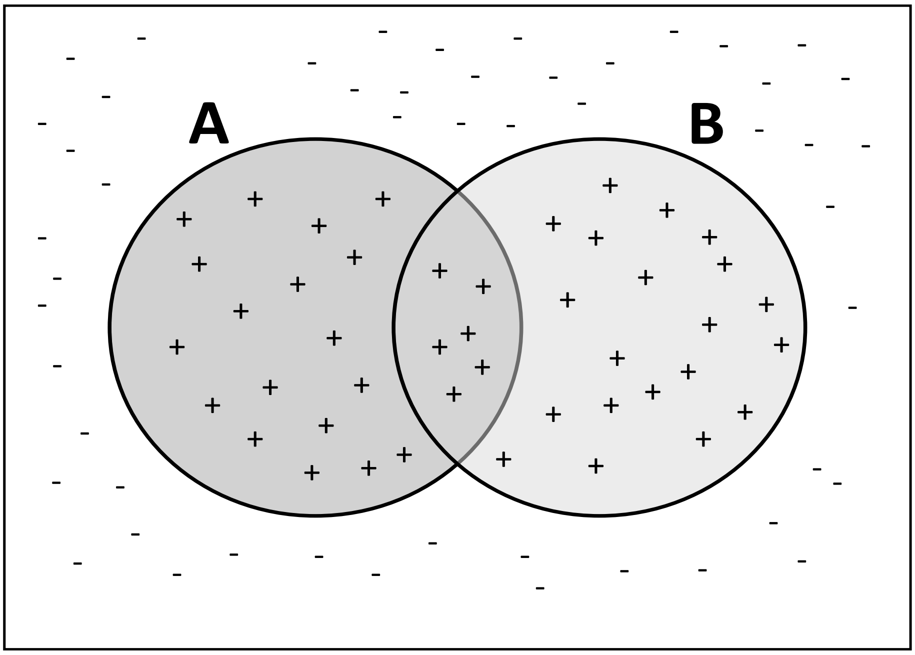

Figure 2. Euler diagram illustrating the assumptions of the noisy-OR rule. Plus signs denote the presence of the outcome, whereas minus signs denote no outcome. In the intersection, some instances of the effect are due to the joint effect of A and B.

According to the noisy-OR rule, the probability of the effect if two potential causes A and B are present (and no other effective cause is present) can be computed as follows:

$$p\ (e|A\&B) = qA + qB - (qA \cdot qB) \qquad \text{Equation 1}$$

Where qA and qB are the causal influences of A and B to produce the effect e. The strength or causal influence of a potential binary cause is the proportion of times it produces the outcome, in a hypothetical context in which no other effective causes of the same outcome exist (Cheng, 1997).

Consider the product from the pink eyes scenario that combined A and B substances, each with a 40% of probability of causing pink eyes. If we use the noisy-OR integration rule, we correct the linear sum of the two independent probabilities (80%) by subtracting the overlap between them (by applying Formula 1), generating a likelihood of producing pink eyes of 64% [i.e., 40% + 40% - 16].

This noisy-OR rule is incorporated in several causal learning models that I will discuss below. It is particularly clear in the Power PC model, but it is used in associative models too (Danks, Griffiths, & Tenenbaum, 2003), and it also appears in some form in Bayesian causal induction models (c.f., Griffiths & Tenenbaum, 2005, 2009). According to the noisy-OR rule the probability of an effect falls between the probability of the effect given the stronger single cause and the linear sum of the probabilities. Returning to the pink eyes example, if we want to know the probability of the effect given A (40%) and B (40%), the noisy-OR rule predicts that it should be above 40% and below 80%, thus, the compound of the two causes is more effective than the stronger cause alone but less than their sum. But, do people spontaneously use this rule for independent binary causes? Or, do they default to the simpler linear summation rule? The "rational" solution, assuming independent probabilities, is the noisy-OR strategy, but people do not always follow rationality when making decisions.

Chapter 3. Models of Causal Learning

'Concepts, like individuals, have their histories, and are just as incapable of withstanding the ravages of time as are individuals'

Kierkegaard

The nature of the psychological process underlying human causal learning and judgment has been the subject of some controversy (Holyoak & Cheng, 2011; Mitchell, De Houwer & Lovibond, 2009; Shanks, 2010). This debate concerns the relative viability of three major positions in this field: Rule-based models, associative models, and Bayesian models. Rule-based models represent organisms as "intuitive statisticians" who can extract contingency information by applying a rule to integrate probabilities or frequencies of events (Peterson & Beach, 1967). By contrast, associative models postulate that contingency learning is, in fact, the result of associations formed between all contiguously presented events. Finally, Bayesian approaches to cognition take into account beliefs and expectations in the form of a prior probability distribution over potential hypotheses. Our previous beliefs are transformed into posterior beliefs in light of new data. Each of them will be described in detail.

3.1. Rule-based models

Edwards (1954) postulated that humans behave as if they calculated probabilities and acted as 'intuitive statisticians'. This concept is at the core of rule-based models. According to rule-based accounts, people encode representations of event frequencies during learning, and then combine this information according to some arithmetical rule to arrive at a final contingency judgment. But how does the person fulfill this role as an "intuitive statistician" and then use this information to make a judgment? In many cases, the events are binary (something happens or does not happen, like getting heads on a coin tossing game). With binary causes and effects, the statistical relation between the cause and the effect can be cast in a traditional 2 x 2 contingency table (see Table 1) where cell *a* contains the frequency of co-occurrence of the presumed cause (C) and the outcome or effect (E), cell *b* contains the frequency of the occurrence of the cause without the effect (~E), cell *c* the frequency of the outcomes when the cause is absent (~C), and cell *d* contains the frequency of the joint absence of the cause and the outcome. According to rule-based models, it is assumed that people store some mental representation of this contingency table and to make their judgments they apply some arithmetic rule to this stored information.

		Effect	
		Present (E)	Absent ($\sim E$)
Potential Cause	Present (C)	$\boldsymbol{a}$	$\boldsymbol{b}$
	Absent ($\sim C$)	$\boldsymbol{c}$	$\boldsymbol{d}$

Table 1. 2 x 2 contingency table.

3. 1. 1. ΔP and Probabilistic Contrast Model

Although they are different in some aspects, in this section ΔP and Probabilistic Contrast Model (PCM hereinafter) models will be covered together as they share many factors like the integration rule they use, and some theoretical problems. According to the ΔP model, the relation between such binary events mentioned in Table 1 is quantified using the ΔP coefficient (Allan, 1980; Cheng & Novick, 1992; Jenkins & Ward, 1965), which is defined as the difference betweem the probability of the outcome given the presumed cause, p(e|c), and the probability of the outcome in the absence of the presumed cause, p(e|~c). This is formally represented in an equation that relates the probabilities to frequencies:

$$\Delta P = p(e|c) - p(e|{\sim}c) = \frac{a}{(a+b)} - \frac{c}{(c+d)} \qquad \text{Equation 2}$$

Intuitively, ΔP refers to the difference a potential cause makes in the probability of the occurrence of the effect (Perales & Shanks, 2008). When p(e|c) exceeds p(e|~c) the contingency is positive, and we would assume that c is a generative cause of the effect (e), as long as c is not itself correlated with another event which is the true cause. When p(e|~c) exceeds p(e|c) the contingency is negative and we would assume that c prevents or reduces the likelihood of the effect. Finally, when p(e|c) and p(e|~c) are equal there is a zero contingency, and we would assume that the cause has no influence over the outcome.

On other hand, the PCM (Cheng & Novick, 1990, 1991, 1992) proposes that people start out with some pre-existing conceptions of possible causal factors. Then, contrasts (or covariation like the one defined by the ΔP rule) are computed for 'focal sets' that are restricted to events in which each of these factors are systematically present or absent to identify which is the specific influence of the target cause over the outcome (Cheng & Novick, 1992). Specifically, Cheng and Novick (1990, 1992) proposed in their PCM that ΔP index should be calculated holding constant all relevant factors except the potential cause being valued. This is analagous to the scientific processs of holding other factors constant while studying the effect of a single factor. To differentiate the general ΔP index, this new index is called conditional ΔP.

A very large number of studies have shown that associative judgments are highly sensitive to variations in ΔP (Allan, 1993; Peterson, 1980, Shanks, 1985, 1987, 1995; Shanks & Dickinson, 1987; Ward & Jenkins, 1965; Wasserman, Elek, Chatlosh & Baker, 1993), and that the correspondence between ΔP and judgments is often remarkably close. However, although it is true that causal judgments tend to correlate with ΔP, this correlation is far from perfect (Perales & Shanks, 2007). There are a number of studies showing that judgments systematically vary across conditions in which ΔP is held constant (Allan, Siegel, & Tangen, 2005; Blanco, Matute & Vadillo, 2013; Blanco, Matute, & Vadillo, 2011; Buehner, Cheng, & Clifford, 2003; Lober & Shanks, 2000; López, Almaraz, Fernández, & Shanks, 1999; Matute, 1995, 1996; Matute, Yarritu, & Vadillo, 2011; Hannah & Beneteau, 2009; Yarritu, Matute, & Vadillo, 2014; Wasserman et al, 1993); so it is apparent that ΔP does not provide a complete description human causal reasoning. As White (2009) pointed out, one problem is that the ΔP rule only indicates the degree of empirical association between a posible cause and an outcome, and not the strength of the cause or the likelihood that it does cause the outcome.

The most important thing to keep in mind regarding these models is that they use a linear integration rule to combine the influence of several potential causes, which is not only related with being rule-based models, as these models can employ either a linear or a noisy-OR rule. As this thesis is focused on testing which integration rule people use when combining multiple causes, this work is relevant to the understanding of causal learning models. The experiments that comprise the thesis should show which integration rules people use, supporting some models but not others.

3. 1. 2. Power PC model

Although, as mentioned above, people's causal and contingency judgments often map onto ΔP (Wasserman, Elek, Chatlosh, & Baker, 1993), Patricia Cheng (1997) has argued that ΔP is an inappropriate measure of contingency because it states that a pair of events can be known to be unrelated in conditions in which any reasonable person would conclude that there is insufficient evidence to decide whether or not they are related (Wu & Cheng, 1999). For this reason, she presented a new model of causal induction called the power PC theory, intended to be both a normative and descriptive account of causal induction. Although the theory is in some aspects very complex, at its core two very simple equations specify the degree of causal power between a target cause and an effect as a function of the probability of the effect in the presence and the absence of the potential cause. The starting point of this account is the existence of some normative problems that appear not to be solved properly by either the ΔP rule or the PCM, and as mentioned above, by some empirical findings that are difficult to reconcile with the predictions of these models.

Causal power and causal structures are theoretical entities that need to be estimated, combining causes to infer the causal structure behind them, and assess their causal strength (Cheng, 1997; Waldmann, Cheng, Hagmayer, & Blaisdell, 2008; Waldmann & Holyoak, 1992). Starting from a rational analysis of situations involving causal judgments, Cheng (1997) showed that deriving a normative computation of the strength of a causal relation is possible (assuming compliance with certain restrictions of the model). For example, if one suffers spontaneous irritation (e or effect) from time to time, it is reasonable to assume there is something in the environment (a or background causes) responsible for that. If the introduction of a new factor such as a new food (c or cause) makes the irritation more probable, then we could reasonably argue that the joint probability of the effect and the cause, p(e|c), reflects the additive effects of the candidate factor and the background causes. If we now assume that a and c are independente and non-interactive, we can segregate the influence of c from the a*c compound by applying probability calculus. Causal strength computed following this way is called *generative causal power* or q:

$$q = \frac{\Delta p}{1 - p(e|\sim c)} \qquad \text{Equation 3}$$

Thus, causal power is directly related with contingency and the base rate of the effect in the absence of the target cause. Therefore, the power PC theory suggests that ΔP is a conservative estimator of causal power. For generative causes, the higher p(e|~c) is, the more conservative ΔP is. In the case of preventive causes, causal power is computed as p:

$$p = \frac{-\,\Delta p}{p(e|{\sim}c)} \qquad \text{Equation 4}$$

For preventive causes, the lower p(e|~c), the more conservative ΔP is. In the case where the effect does never occur, the preventive power cannot be computed. It follows from Equation 3 that the generative causal power of the target cause to be judged can not be estimated if the effect occurs in every case in which the target factor is absent. Similarly, it follows from equation 4 that the preventive causal power of the target cause to be judged can not be estimated if the effect does not occur in any case in which the target factor is absent.

Whether or not human causal intuitions actually conform to the power PC theory's predictions for intermediate probability base rates is a matter of some dispute (Allan, 2003; Barberia, Baetu, Sansa & Baker, 2014; Collins & Shanks, 2006; Griffiths & Tenenbaum, 2005; Perales & Shanks, 2007, 2008). Some studies have demonstrated that causal judgments are well calibrated to causal power (Buehner, Cheng, & Clifford, 2003), but clear violations of the rule have also been reported (Lober & Shanks, 2000). Buehner and his colleagues argued that these violations of the power PC theory can be attributable either to inherent ambiguities in the typical wording of causal judgment questions, or to memory biases. The typical wording regarding causal judgments leaves open whether causal influence should be judged in the same context as learning, or in a different (conterfactual) context in which background causes of the effect are not present. Participants may truly base their causal estimates on power, but misrepresent the conditional probabilities $p(e|c)$ and $p(e|\sim c)$ on which the power calculation is based. If that were the case, any bias in power-based causal judgments could be due to bias in perceiving the conditional probabilities.

More important for the present work is the problem raised by Luhman and Ahn (2005) and White (2009). The authors pointed out that this theory seeks to compute a context-free estimate of causal strength for individual causal links, and to achieve that it requires many assumptions. Alternative causes (especially the unobserved ones) must work in very specific ways. They must be generative and may not interact with the target cause or target causes. It is not unreasonable to start from the assumption that the world is a messy place, disorganized and uncontrolled and idealized controlled situations such as these ones are few even in well-controlled experiments. Such required conditions are so restrictive that causal power could not be computed in the real world. These difficulties imply that an accurate computation of causal power as predicted by the theory requires a huge amount of accurate knowledge, much of which reasoners are unlikely to possess (Cheng & Novick, 2005).

It is important to mention again that this rule-based model assumes that the noisy-OR integration rule is used to assess the causal impact of a compound of possible causes, whereas ΔP and PCM models used the linear integration rule. Apart from their different integration rules, the three models are formally very similar.

3.2. Associative models

From the associative perspective, the processes involved in causal learning are analogous to (if not exactly the same as) those underlying simple conditioning in animal learning (Rescorla & Wagner, 1972). Following this account, people are assumed to form psychological associations between pairs of events during the learning phase, and then to base their judgments directly on the strengths of these associations. Generally when two events occur together, the relationship between them is strengthened. Conversely, when the events occur separately, the relationship between them usually weakens. The models usually use the notion of prediction error (a linear operator) in which the change in associative strength that occurs on each trial is a function of the difference between the pre-existing strength of an association and the maximum strength possible. The aim of this chapter is to discuss the main models of associative learning that exist today, and analyze in more detail the Rescorla-Wagner (1972) model.

3.2.1. Rescorla-Wagner Model

The Rescorla and Wagner (1972) model has possibly been the most influential model in the psychology of learning since its appearance, 40 years ago. It conceives learning as an adjustment between the expectations that the participant (or animal) has, and what actually happens. Learning occurs when there is surprise. Surprise understood here as the discrepancy between what the participant expects and what he or she actually experiences (Kamin, 1969). The greater this difference is, the larger the increment in learning (strengthening of the association), and vice versa. The second principle postulates that learning about a stimulus also depends on all the stimuli that are present in the situation, and not just on the presence of the stimulus of interest (target). What individuals learn on each new conditioning trial adds information to what they already knew before the trial began. The power of the relationship between the conditioned stimulus (CS) and the unconditioned stimulus is known as associative strength. After repeated association with an unconditioned stimulus, the CS comes to elicit the response previously generated by the unconditioned stimulus itself. The unconditioned stimulus (US) naturally, unconditionally and automatically triggers a response. For example, if I smell one of my favorite foods I may immediately feel very hungry, that is the US. But if I paired a neutral stimulus (a tone) along with the smell of the food several times, the tone alone would eventually trigger the response.

According to the Rescorla-Wagner (1972) model, changes in the associative strength are described by Equation 5. According to this formulation of the delta rule, the associative strength (the strenght of the association between the mental representation of the cause and the mental representation of the outcome) of Cause *i* increases or decreases on trial *t* as follows:

$$\Delta Vit = \alpha\beta(\lambda rft - \Sigma V t - 1) \quad \text{Equation 5}$$

Where ΔVit is the increment or decrement to Vi on trial t, α and β are learning rate parameters with values from zero to one, and λrft refers to the strength of the effect on trial t, which takes a value of one when the effect occurs and a value of zero when the effect does not occur. Finally, $\Sigma Vt - 1$ denotes the sum of the associative strengths of all potential causes that occurred on trial t. The change in associative strength is proportional to the difference between the expected status of the effect and the true status of the effect. Essentially, the Rescorla-Wagner (RW, 1972) model describes the popular delta rule that has been subsequently been used in connectionist approaches to contingency learning and judgment (Chapman, 1991; Gluck & Bower, 1988; Shanks, 1991; Shanks & Dickinson, 1987; Wasserman, Elek, Chatlosh, & Baker, 1993).

This model was originally developed to explain classical conditioning phenomena, most particularly those involving the interaction of simultaneously presented stimuli, that could not been explained with previous learning theories. Thus, the Rescorla-Wagner model provided an explanation for blocking (Kamin, 1968), relative validity (Wagner, Logan, Haberlandt, & Price, 1968; Wasserman, 1974), overshadowing (Pavlov, 1927), and conditioned inhibition (Pavlov, 1927). It also become apparent that many of the phenomena which at first sight were outside the model, could also be explained using alternative models that have assumptions in common with Rescorla-Wagner model.

It is important to note that at assymptote, assuming the presence of a constant context and equal betas for reinforcement and non reinforcement, the RW (1972) model computes ΔP (Chapman and Robbins, 1990; Danks, 2003; Tangen & Allan, 2003). Wasserman, Chatlosh, Elek & Baker (1993) generalized it to the more realistic case of unequal betas.

3. 2. 2. Other associative models

There are a number of effects that have important implications for theories of learning, and need to be explained before the rest of the models. Kamin's (1969) blocking effect refers to failures of learning and/or at least the expression of classically conditioned responses when the CS is presented as part of a compound that includes another CS that has previously been used to establish the conditioned response. Lubow and Moore's (1959) latent inhibition effect refers to the observation that a familiar stimulus takes longer to acquire associations (act as a CS) than a new stimulus, possibly because a lack of attention or because this stimulus has been deemed irrelevant. Another phenomenon that challenges learning theories is retrospective reevaluation (Dickinson & Burke, 1996; Wasserman & Berglan, 1998), where the established response to a stimulus is modified by later experience with other stimuli. Finally, the overshadowing effect demonstrates that when subjects are trained with a compound (composed of elements that may or may not differ in salience), there is less conditioning to a weak CS if it is combined with a more salient CS during conditioning trials.

There are other models of associative learning that, although not directly tested in this thesis, should be noted. The basic assumption of the Mackintosh (1975) model is that more attention is paid to stimuli that are better predictors of their outcomes or consequences, that is, relevant stimuli. The model is an attempt to respond to difficulties that the Rescorla-Wagner (1972) model had in accounting for phenomena like latent inhibition. It combines the assumptions of selective attention theories (Zeaman & House, 1963) with associative processes, emphasizing the role of the reinforcement in associative learning. This model is based on two fundamental assumptions. First, the associability of stimuli changes through the experience. Thus, the model explains latent inhibition (Lubow, 1973) assuming that the learning rate parameter for the CS (α) decreases with repeated presentations. Second, theories of selective attention explain overshadowing and blocking phenomena as a result of the competition between stimuli for control of attention. The attention the subject lends to stimuli (or its associability), changes with the experience that the subject has of such stimuli. Changes in stimulus associability are directly related to its predictiveness compared to the predictive power of other stimuli.

By contrast, Pearce and Hall (1980) start from the supposition that the participant pays attention to stimuli that are not good predictors of their consequences (or there is not enough information to know if they are), to determine their relevance. Similar to Mackintosh (1975), this model focuses on the processing of the CS and emphasizes the role of attention in learning, but it has two unique features that differentiate it from the Mackintosh's model. It proposes a new principle to determine the changes that occur in α and that lead to decreases in α when the CS is a good predictor of the US, and it also completely abandons the idea that over the course of conditioning the effectiveness of the US changes.

Most relevant to the experiments in this thesis is Pearce's (1987) configurational model. This model postulates that the presentation of a compound of two or more stimuli results in the formation of a unitary representational entity that is distinct from its elements. It is this entity that creates a unique association with the reinforcer. A compound stimulus is considered to be a single entity that is different from the individual elements that made it up. Returning to the pink eyes example, this means that a combination of products, A and B for example, is processed as a unique configuration, named X, and not as individual elements A and B. Each configuration of stimuli is encoded independently, and participants learn about complete configurations, not about the stimuli that compose them. Because each configuration is represented independently, this implies that the associative strength of a compound does not necessarily have any relation to the associative strength of its elements (although it does in some versions, its similarity arises from the common elements in A and the AB compound). In other words, there may be no summation. And if there is summation, it does not have to obey either the linear or the noisy-OR rule.

Van Hamme & Wasserman (1994) proposed a modified version of the Rescorla-Wagner model (1972), by assuming that the parameter that defines the intensity or associability of the CS (α) has two different values depending on whether the stimulus is present or absent in the trial. If the CS is present, the parameter (α) will be positive, but if the CS is absent, α will be negative, although usually smaller in absolute value than when the CS is present. This means that stimuli indirectly activated through within-compound-associations have a negative learning parameter, thus phenomena of retrospective reevaluation can be now explained.

To sum up, it is important to note that all the associative models discussed so far (except the Pearce model) work under the assumptions of the linear integration rule, as they are variations based on the Rescorla-Wagner (1972) model, which used the linear integration rule, and the rest of the models have kept that idea. As mentioned above, in Pearce's model a compound has its own representation, different from the representation of its elements. It is worth mentioning that Rescorla (1970) made a similar assumption claiming that a compound includes a configural element, i.e., A and B equal AXB, with X being the common configural element. When a novel compound is presented for the first time and has not been learned about it, that compound response is summated from what was previously learned about the elements that compose it. This generalization depends on how much the compound resembles the elements. Due to this rule of generalization, the configural view predicts that the strength of a given compound will approximate the average strength of each of its components in isolation. That is the main reason to explore the contribution of the different models of causal learning, and its relevance to the study of the way in which different causes should be added when presented in combination.

3.3. Bayesian Models

Although I will not focus on Bayesian approaches throughout this thesis, our results have some important implications for the development of these models. This increasingly influential approach to causal induction proposes that subjects solve causal reasoning problems in a way consistent with the Bayesian formalism, originally developed in the field of computer science and statistics (Pearl, 2000; Spirtes, Glymour, & Scheines, 2001; Woodward, 2003; Sloman, 2005; for an interesting review see Glymour, 2003). Bayesian models offer a formal framework for representing and reasoning about causal systems using causal models, a form of graphical representation of both deterministic and probabilistic causal systems (Hagmayer & Sloman, 2009).

Probabilistic models aim to explain human cognition by appealing to the principles of statistics and probability theory, which dictates how an agent should act in conditions that involve some degree of uncertainty. A Bayesian model is a probabilistic graphical model that represents a set of random variables and their conditional dependencies via a directed acyclic graph. For example, a Bayesian network could represent the probabilistic relationships between causes and consequences, such as smoking and lung cancer. Given symptoms, the network can be used to compute the probabilities of the presence of various diseases.

Different versions of the Bayesian approach propose different learning algorithms for inferring the causal structure underlying a given set of covariation data. Bottom-up models focus on providing algorithms for inferring relationships based on statistical data in the absence of any other cues (Spirtes et al., 2001; Gopnik, Glymour, Sobel, Schulz, Kushnir, & Danks, 2004). By contrast, the "causal model" approach emphasizes the role of top-down domain-general assumptions that constrain and inform the process of causal induction (Waldmann, 1996; Waldmann & Hagmayer, 2001). Finally, the "theory-based" approach focuses on the influence of domain-specific prior knowledge (Tenenbaum & Griffiths, 2003; Tenenbaum, Griffiths, & Kemp, 2006).

Most importantly, Bayesian models can use the noisy-OR integration rule or the linear one, exactly as associative and rule-based models do. Which integration rule is used is a prior issue in the formulation of the models. Therefore, any evidence about which integration rule is more prevalent or intuitive for people has direct implications for the development of this type of models.

Chapter 4. Combination of causes in described situations

The main objective of this thesis is to verify what type of integration rule was used by participants in a series of experiments that combined different cover stories, sets of probabilities, and the way information was presented. In the introductory chapters the two main integration rules have been presented, the linear summation and the noisy-OR rule. It was also mentioned that learning models differ in the integration rules they incorporate.

Previous studies have compared the predictions of a learning model using a given integration rule against the predictions of other model using other integration rule. For example, studies in which the Power PC model's predictions, which use the noisy-OR integration rule, are compared with the ΔP model's predictions, using the linear integration rule. The results have been mixed. On one hand, there are some studies indicating that the noisy-OR rule does not fit the data well (Allan, 2003; Perales & Shanks, 2007). On the other hand, other studies seem to show evidence that people's behavior conforms to the noisy-OR standard (Buehner et al., 2003; Griffiths & Tenenbaum, 2005; Holyoak & Cheng, 2011), but can also support alternative explanations (Lober & Shanks, 2000).

In most of these experiments, participants were trained with a potential cause in a noisy context involving the presence of other potential effective causes of the same outcome. Each experiment starts with a learning phase composed of a series of experimental trials where the participant is learning the exposed information. In each of the trials, the subject is presented with information about the potential causes and the effect, following the model of the 2x2 contingency table explained in Causal Learning Models section (Table 1). Specifically, a trial can have a present or absent potential causes, and the presence of the effect or not. In causal learning experiments every cell (a, b, c and d) is manipulated in order to expose participants to different experimental situations where causal learning models make different predictions. After the learning phase, the participants are asked to evaluate the influence of the target cause over the final outcome. The wording of the question varied slightly depending on the experiments, but usually was something like: *To what extent does X produce Y?* An important factor to consider when analyzing the responses of the participants is that they may have used a variety of strategies based on the available information, and we cannot conclude that the use of one completely excludes the representation of others. Even if participants' answers match the linear integration rule in a particular experiment, this does not demonstrate *per se* that they cannot represent power in terms of the noisy-OR rule. Any of these possible answers is a correct description of the action of a cause in the training situation and can be directly derived from a representation of power.

In the experiments presented here a more direct test of whether participants spontaneously use the noisy-OR integration rule is used. Previous experiments asked participants to infer the individual influence of a target cause from situations in which multiple causes simultaneously influence the outcome. Quite surprisingly, none of the experiments mentioned so far involving information about several candidate causes asked participants to predict the probability of the combined effect of some given potential causes. If this is done the noisy-OR rule makes very specific predictions. When people combine the influence of several causes, they should logically fall in between the value of the strongest individual cause and the linear sum of the probabilities of both causes. Although most studies explore the contrast between two different learning models (this issue is addressed in the Discussion), none of them were designed specifically to test if the participants used the linear or noisy-OR integration rule. The experiments of this thesis were not designed to compare causal models with different predictions, but to investigate the integration rules that are the basis of each of them.

In the present series of experiments, participants were provided with percentage information about the individual influence of different potential causes and asked to predict what their combined effect over the outcome would be when acting simultaneously. For example, imagine that we have information regarding the effect of two potential factors. Each of these causes produces the effect half of the times it is present. What would be the probability of the effect be if both are presented simultaneously? Similar to our coin tossing game, each of the causes would produce the outcome half of the times, and therefore the probability of the outcome (getting heads) will equal .75. Participants' inferences might be normative and approximate this prediction (following the noisy-OR rule). Alternatively, people might use other simpler heuristics when faced with this situation. If both causes are summed, participants might overestimate the probability of the outcome.

Importantly, the noisy-OR rule is the normative description of the influence of several causes over the same outcome only if the participants assume that each of the causes *independently* affects the outcome, that is, if there are no interactions between causes. Models that use this rule, such as Cheng's model (1997), argue that independence is a default assumption. In other words, people face reality assuming independence between causes, and this belief is only abandoned if strong evidence against it is observed. Contrary to this approach, it is possible that participants do not always assume independence, or assume independence not attending to the cover story. The main reason science fiction scenarios are used in all of our experiments is to encourage participants to assume independence between causes and ensure, as far as possible, that there is no prior knowledge that could contaminate the results.

The first experiments have in common the same experimental structure, with only the set of specific probabilities varying across experiments. In general, participants received a written questionnaire in which information about a number of chemical substances with Greek names was given. These substances had the capacity to produce pink eyes. This effect was chosen to ensure that the base rate of the effect in the absence of the candidate cues is zero. Each substance had different causative force, expressed as a percentage. The participants were asked to imagine that they wanted to have pink eyes. To do this, they must decide between two combinations of substances. The probabilities associated with the substances were designed so that we can discriminate between linear and noisy-OR integration rules.

4.1. Experiment 1A

4. 1. 1. Method

Participants, procedure and design

Twenty-six Psychology degree students from the University of Deusto volunteered to participate in Experiment 1A. All the materials were presented in booklet (see Appendix 1 in page 153 for the specific instructions for this experiment). Participants were asked to imagine that a number of chemical substances (Alpha, Beta, Delta, Gamma and Omega) had been discovered to have the ability to change the color of a person's eyes to pink. The assignment of substance names, and therefore their correspondent probabilities over the outcome (see Table 2), was partially counterbalanced following a latin-square design. The order in which the substances appeared on the sheet of paper was also counterbalanced. The set of probabilities for this experiment was as follows:

Cue	*Cue in graphs*	*Probability*
Alpha	A	40%
Beta	B	40%
Gamma	C	80%
Delta	D	64%
Omega	E	0%

Table 2. Set of probabilities for Experiment 1A.

Upon receiving this information, participants were told that different companies were selling different products containing different pairs of chemical substances. They were asked to choose which product they would buy if they wished to have pink eyes. Specifically, participants were asked whether they would prefer a product containing substances A and B together versus one containing substances C and E, and whether they would prefer a product containing substances A and B together to another compound containing D and E. This consisted on a simple forced choice, they had to choose one or the other. The order of these two questions and the order in which both compounds were presented in each question (AB first vs. CE/DE first) was counterbalanced across participants.

Would the participants prefer a compound including A and B to one including C and E? These preference questions were designed to directly test the linear integration rule versus the noisy-OR integration rule, being the preference score the proportion of participants choosing one compound or another. If participants use a simple non-normative linear summation strategy to make this decision, they should be indifferent in the first preference question, because, from the Equation in Figure 1, the linear addition of A and B (40% + 40% = 80%) and the linear addition of C and E (80% + 0% = 80%) are exactly the same. However, if participants use a noisy-OR integration rule they should prefer the compound CE, whose ability to produce the effect is 80% + 0% - (80% · 0%) = 80%, over the AB compound, whose ability to produce the effect is 40% + 40% – (40% · 40%) = 64%. In the second preference question the participants were asked to decide between the compound AB and the compound DE (64% + 0%). In this case, if the participants use a linear integration rule, they are expected to choose AB. However, if they use a noisy-OR integration rule they are expected to be indifferent to both compounds. The pattern of predictions regarding the two integration rules for this experiment is depicted below. The calculated probabilities predicted by the integration rules for all the experiments can be found on page 105 (Table 9).

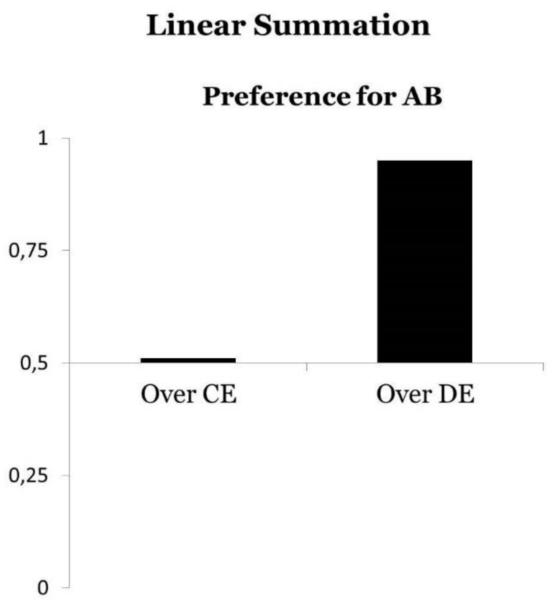

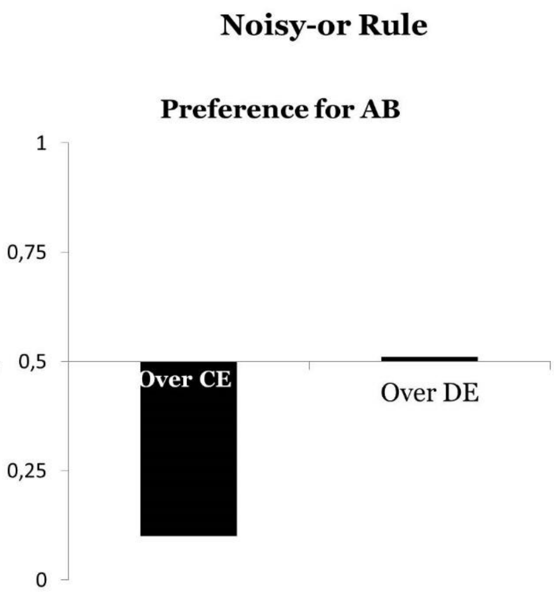

Figure 4. Pattern of preference results as predicted by the two integration rules for Experiment 1A, linear summation in the left versus noisy-OR integration rule in the right.

4.1.2. Results

The choice pattern for this experiment is shown in Figure 5 and does not fit the prediction of the noisy-OR integration rule well. By contrast, it corresponds perfectly with the predictions of the linear summation rule, showing a preference for AB compound over DE (binomial test, p = .009). Although there appears to be a tendency to prefer the AB compound over CE, this trend is not statistically significant (binomial test, p= .327).

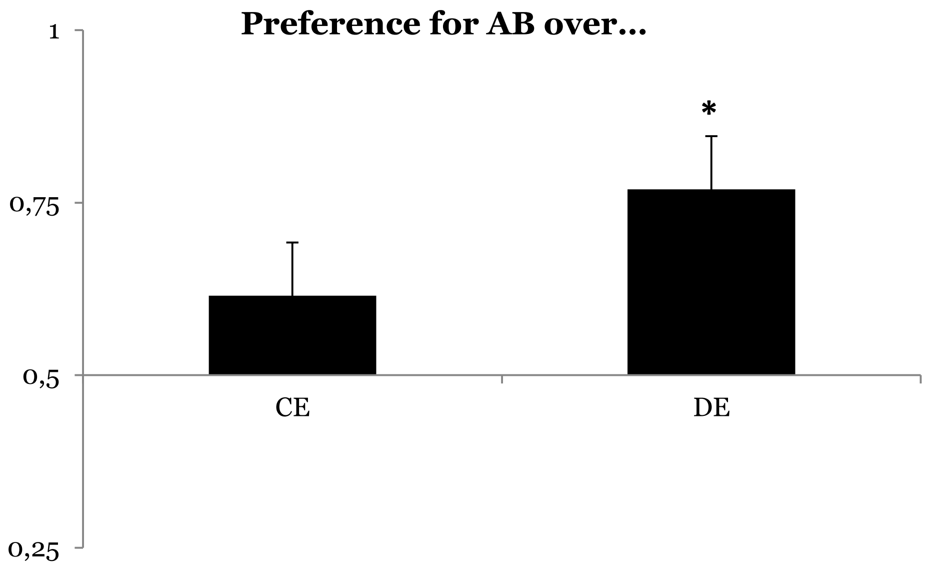

Figure 5. Preference for AB over CE (left), and preference for AB over DE (right) for Experiment 1A. The axis in 0.5 indicates indifference. Scores close to 1 indicate preference for AB compound, while ratings close to 0 indicate preference for CE/DE compounds.

4.2. Experiment 1B

In Experiment 1A a significant tendency to respond according to the linear integration rule was observed. In the first preference question (AB vs CE) no significant differences were found, participants responded equally to both compounds, but in the second question (AB vs DE), there was a significant tendency to prefer the AB compound over DE, as the linear integration rule predicts. The results contradict the predictions of the noisy-OR integration rule.

In Experiment 1B, the procedure and design were the same as the previous one, except that different values were assigned to the chemicals (see Table 3 below).

Cue	***Cue in graphs***	***Probability***
Alpha	A	60%
Beta	B	65%
Gamma	C	97%
Delta	D	86%
Omega	E	0%

Table 3. Set of probabilities on experiment 1B.

The new set of probabilities was designed so that the linear sum of some of the cues was higher than 100%. If a participant merely makes a spontaneous linear summation, thinking that it is a correct answer, it is possible that when calculating the sum but getting an outcome higher than 100% he or she might realize that this result is not normative, and decide to use a different strategy. Therefore these cues could promote the processing using the noisy-OR rule. It will be noted that the probabilities used in this experiment are not all regular multiples of ten or simple binary expansions (e.g., 64). This was done to discourage participants from falling back on a strategy of simple practice arithmetic calculations.

The probabilities assigned to each cue were constructed in a way to discriminate between the two integration strategies, so that if the participants used the linear rule integration, they would prefer the AB compound in both preference questions. By contrast, the noisy-OR rule predict a preference for the CE over the AB compound; and indifference between AB and DE (see Figure 6). For all the specific calculated values, see Table 9 on page 105. Finally, to show that that the results do not rely on substances A and B having equal probabilities, substance A had a value of 60, and substance B 65.

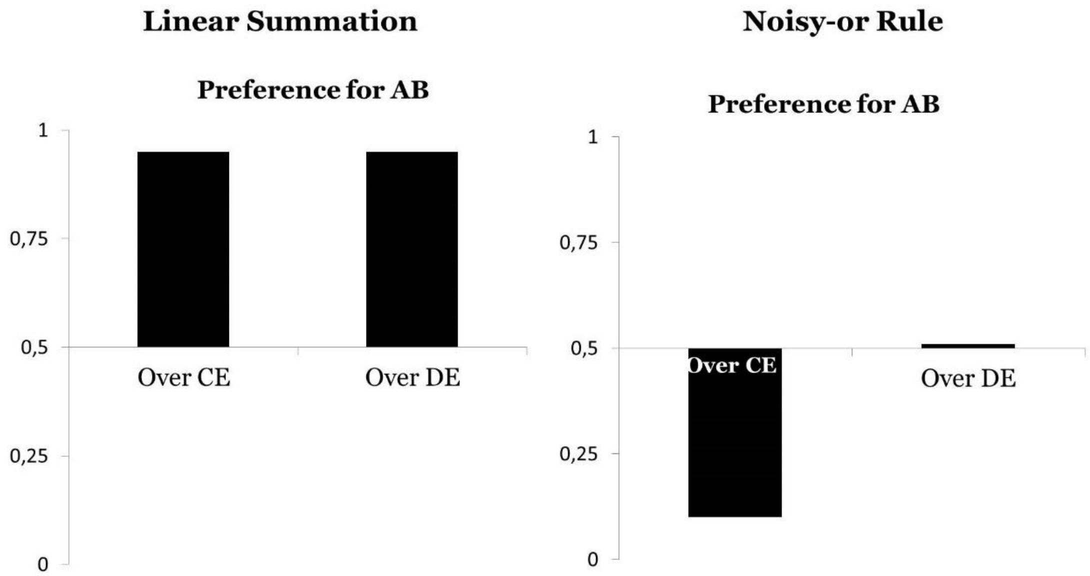

Figure 6. Pattern of preference results as predicted by the two integration rules for Experiment 1B.

4.2.1. Method

Participants, procedure and design

The experiment was completed by a sample of 43 students, 28 of them students of the Psychology MSc, and 15 from the Psychology Bachelor's degree Programmes, from Deusto University. The design and procedure were the same as in the previous study, except the probability set (shown in Table 3). The instructions (see Appendix 1) were also the same. The order of the questions and the order in which items appear on each question were counterbalanced.

4.2.2. Results

The linear summation rule predicts preference for AB over CE and DE, whereas the noisy-OR rule predicts preference for CE over AB, and indifference between AB and DE. As can be seen in Figure 7, there were significant tendencies to prefer the AB compound over the CE compound, as predicted by the linear integration rule (binomial test, *p*=.014). In the second preference question, the results again fit perfectly with the predictions of the linear integration rule, participants preferred AB compound over DE (binomial test, *p*= .032).

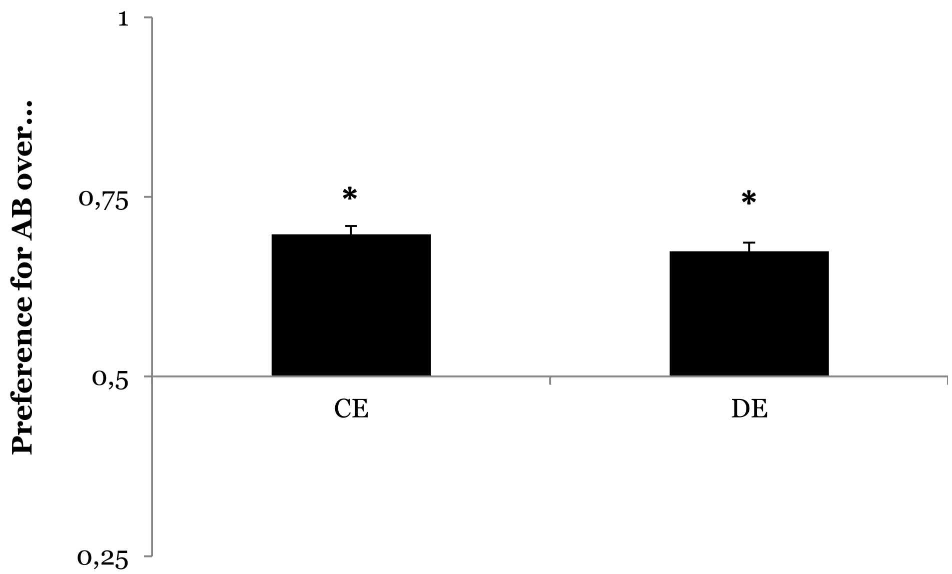

Figure 7. Preference for AB over CE (left bar) and over DE (right bar). The axis in 0.5 indicates indifference at the group level.

This experiment raises an obvious problem for the noisy-OR rule: the results are not consistent with what was expected from this integration rule. One factor that could be on the basis of this result is cue E which is 0%. Participants could perceive this cue not just as ineffective or without causal power, as intended, but as have the ability to inhibit the effect of other potential and effective causes which are paired with it in a compound. In this way, participants may prefer compounds with two ate least partly effective cues (AB compound). Taking this into account, in the following experiment the set of probabilities were again modified to see if the pattern of preference is maintained.

4.3. Experiment 1C

In the previous experiment a significant tendency to prefer the compound AB over the alternatives, CE and DE, was found. This result is predicted by the linear integration rule. As mentioned in Experiment 1B, one factor that may have promoted this result is the cause with 0% effectiveness. It is possible that participants believe that this cause is not just ineffective, but inhibits the causal power of the other causes paired with in a compound.

Experiment 1c is designed to investigate that doubt. In this experiment there is no cause with 0% causal power (see Table 4). Thus, all the candidate causes have at least some generative causal power (above 0%). As in the previous experiment, there were no causes with identical probabilities. The patter of preference based on the different strategies is here again the same as in the Experiment 1A (see Figure 4, and see Table 9 for the calculated probabilities made by several integration rules).

Cue	*Cue in graphs*	*Probability*
Alpha	A	60%
Beta	B	65%
Gamma	C	95%
Delta	D	80%
Omega	E	30%

Table 4. Set of probabilities for experiment 1C.

4.3.1. Method

Participants, procedure and design

The experiment was completed by a sample of 57 students from the Psychology degree, all coming from Deusto University. The design, instructions and procedure were the same as in the previous study, except the probability set (Table 4). As in the previous experiments, the order of the questions and the order in which items appear on each question were also counterbalanced.

4.3.2. Results

The linear summation rule predicts indifference between AB and CE, and preference for AB over DE. By contrast, noisy-OR rule predicts preference for CE over AB, and indifference between AB and DE. As shown in Figure 8, the results replicate those of the previous experiments, despite the change in the probabilities associated to each cause. A significant tendency to prefer AB compound over CE was found (binomial test, p= .008). This result is not predicted by either theory. In the second preference question, participants also seemed to prefer AB compound over DE (binomial test, p= .000).

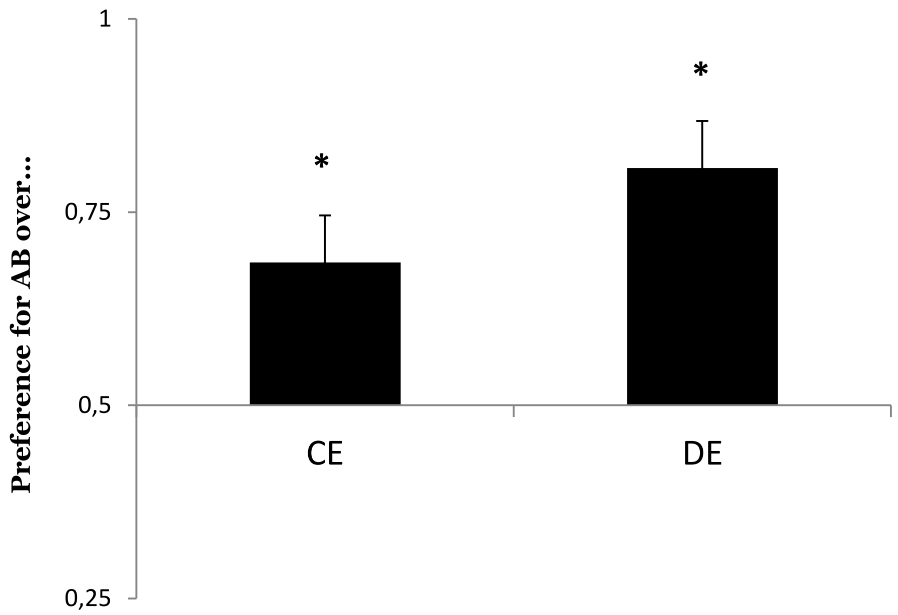

Figure 8. Preference for AB over CE (left bar) and over DE (right bar). The axis in 0.5 indicates indifference.

Removing the non-effective (0%) cause did not change the pattern of results obtained in Experiments 1A and 1B. It seems that the preference for the compound AB remains the same despite the change in the probability set. Therefore, the results of those experiments cannot be explained assuming that a 0% effective cause was perceived as a preventive cause.

4.4. Experiment 1D

This experiment introduces a novel deterministic cause into the set of probabilities to see if the participants' preferences remain the same. It should be easier for participants to reason according to the noisy-OR strategy if a cause is effective 100% of the time, because no combination of powers that are less than 100% can be more effective than a combination of powers that includes a 100% cue. In this case, it should be clearer that 100% combined with anything is 100%. Conversely, when two values below 100% are combined, the final result is less than 100%. This intuition is consistent with the logic of the noisy-OR, but not with the linear logic. Therefore, this should encourage noisy-OR style reasoning.

According to some studies, the deterministic probabilities like 100% or 0% are interpreted in a different way to other probabilities (Li & Chapman, 2009; Kahneman & Tversky, 1979). In the case of the 100%, the cue is seen as a very salient reference point. The authors propose that this salient point receives special attention above its numerical weight. In the present experiment all instructions and procedure are the same as in the previous experiments, but the set of probabilities was modified, to include new potential cause with 100% causal power (see Table 5). We hypothesize that, following the reasoning of Li and Chapman's study (2009), participants would pay more attention to the 100% effective cause, resulting in a reversal of the pattern of results obtained in this experimental series.

Cue	*Cue in graphs*	*Probability*
Alpha	A	85%
Beta	B	40%
Gamma	C	100%
Delta	D	88%
Omega	E	25%

Table 5. Set of probabilities on experiment 1D.

The predictions of the two integration rules are the same in this experiment as in Experiments 1A and 1C, see Figure 4.

4.4.1. Method

Participants, procedure and design

Forty five students from different schools at the University of Deusto volunteered for Experiment 1D. The design, instructions and procedures were kept the same as in previous experiments but the probabilities shown in Table 5 were used. The preference questions and the order of the elements within the questions were counterbalanced.

4.4.2. Results

In this experiment, linear integration rule predicts indifference between AB and CE, and preference for AB when in contrast with DE. Noisy-OR rule predicts preference for CE over AB, and indifference between AB and DE compounds. As can be seen, the results fit with the predictions of the linear integration rule. Although at first glance it seems to be a tendency to prefer CE over AB, this tendency was non-significant (binomial test, p= .072), so in this question participants seems to be indifferent between the two combinations. In the second question, participants significantly preferred the AB compound over DE (binomial test, p= 0.16).

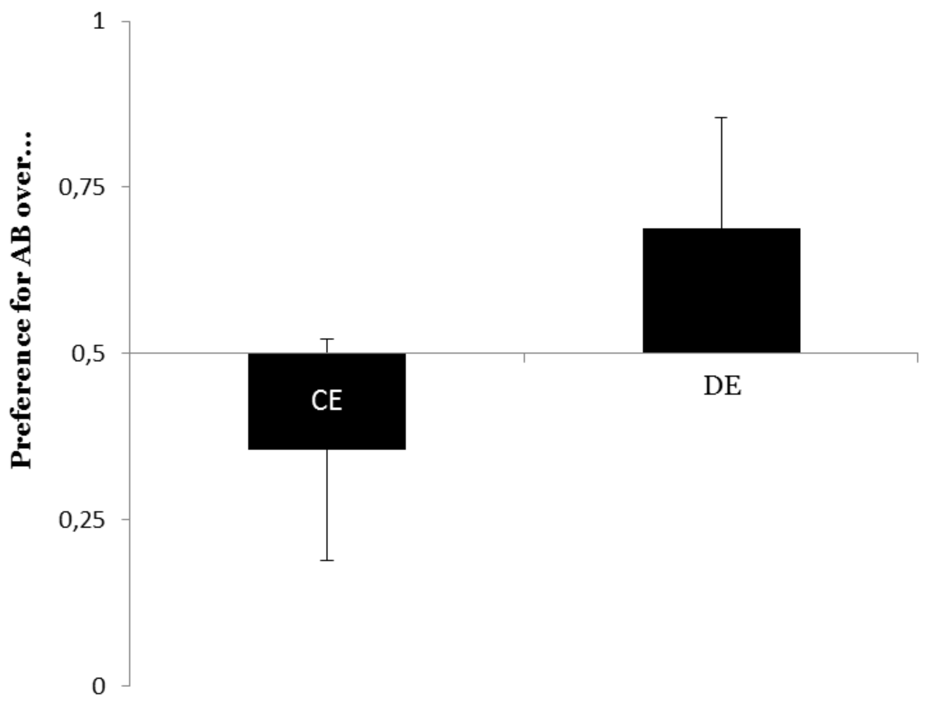

Figure 9. Preference for AB over CE (left bar) and over DE (right bar) for Experiment 1D. The axis in 0.5 indicates indifference.

One could argue that with more participants in the sample, the tendency to prefer CE to AB would become significant. If that were the case this would possibly indicate that, regardless of the integration rule participants are following, when they are indifferent between two compounds they can adopt the other integration rule to “maximize” their choice of winning and avoid uncertainty or indifference. So, if a given participant follows the linear integration rule (and it should be mention that only about 70% of the people made that choice), he or she might be very certain of the AB vs. DE decision (preferring the former). But in AB vs. CE, the linear integration rule predicts indifference, so in order to make a choice, the participant could follow the predictions of the noisy-OR rule and choose the CE compound. And vice versa, if a participant follows the noisy-OR integration rule, he/she will be sure of choosing CE over AB, but in the second preference question, in which noisy-OR predicts indifference, the participant can change his/her strategy in order to make a decision, preferring DE although it is not predicted by the noisy-OR rule. In any case, it can be argued that there was little evidence for this prediction in the previous experiments. So this idea that when a strategy invites indifference in a given preference question then participants used the other does not accurately explain the previous results.

Chapter 5. Summation in adolescents

All previous experiments were conducted with university students, and mostly with Psychology students. These students are trained in Statistics and Probability, and it is possible they may be using that knowledge to perform the experiments. It is therefore possible that the previous results are not representative of the performance of the general population, untrained in Probability Theory. In this case, the aim of these experiments was to check if the results of the previous experiments also hold for with 13 year old students. This chapter includes replications of Experiments 1A and 1B using early adolescent children.

5.1. Experiment 2A

5.1.1. Method

Participants, Procedure, and Design

Thirty-eight secondary school students participated in Experiment 2A. All the participants were 13 years old, 24 were male and 14 female. The procedural details and materials were exactly the same as in Experiments 1A (see Table 2).

5.1.2. Results

As in Experiment 1A, the linear summation rule predicts indifference between AB and CE, but a preference for AB over DE. The noisy-OR rule predicts preference for CE over AB, and indifference between AB and DE. As can be seen in Figure 10, the results of Experiment 2A were very similar to those of Experiment 1A, which was based on the same probability set. Although the graph seems to suggest otherwise, the preference for AB over CE did not reach statistical significance (binomial test, p = .143). However, participants preferred the compound AB to the compound DE (binomial test, p= .014*)* so, again, these results are at odds with the predictions of the noisy-OR rule, and fit the linear integration rule.

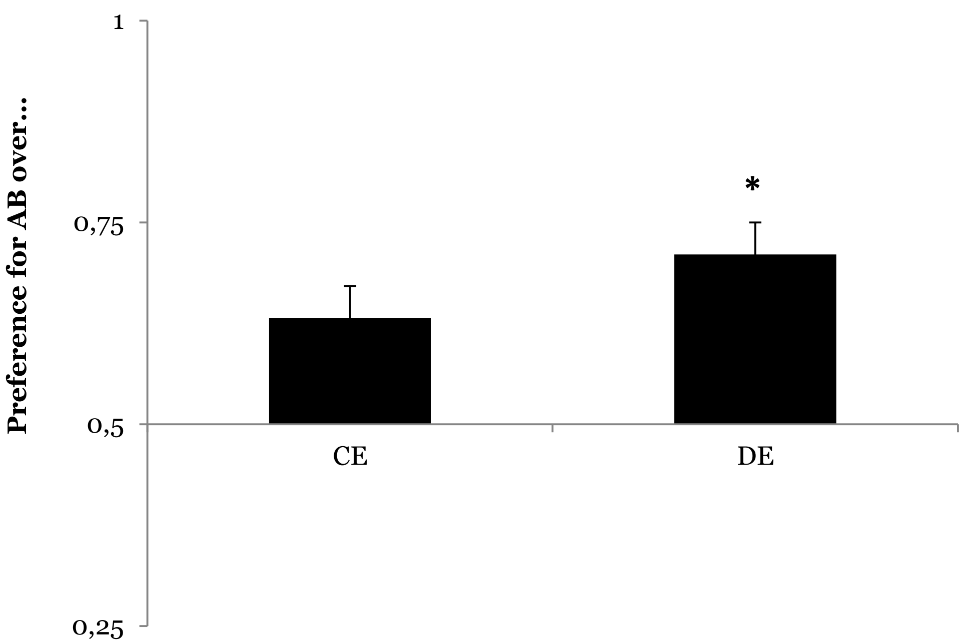

Figure 10. Preferences for experiment 2A with secondary school students.

5.2. Experiment 2B

5.2.1. Method

Participants, Procedure, and Design

Twenty-eight secondary school students participated in Experiment 2B. All were aged 13, and were 15 male and 13 female. The set of probabilities used is depicted in Table 3. The instructions, procedure and materials were exactly the same as in the experiment 1B.

5.2.2. Results

In this experiment, the linear rule predicts preference for AB compound over CE and DE, whereas the noisy-OR rule predicts preference for CE compound over AB, and indifference between AB and DE. The results of experiment 2B are shown in Figure 11. Participants preferred the compound AB to the compound CE (binomial test, p= .013), and also preferred compound AB over DE (binomial test, p= .004). This demonstrates that the results of the previous set of experiments generalize to other populations, and not rely on specific attributions of Psychology students.

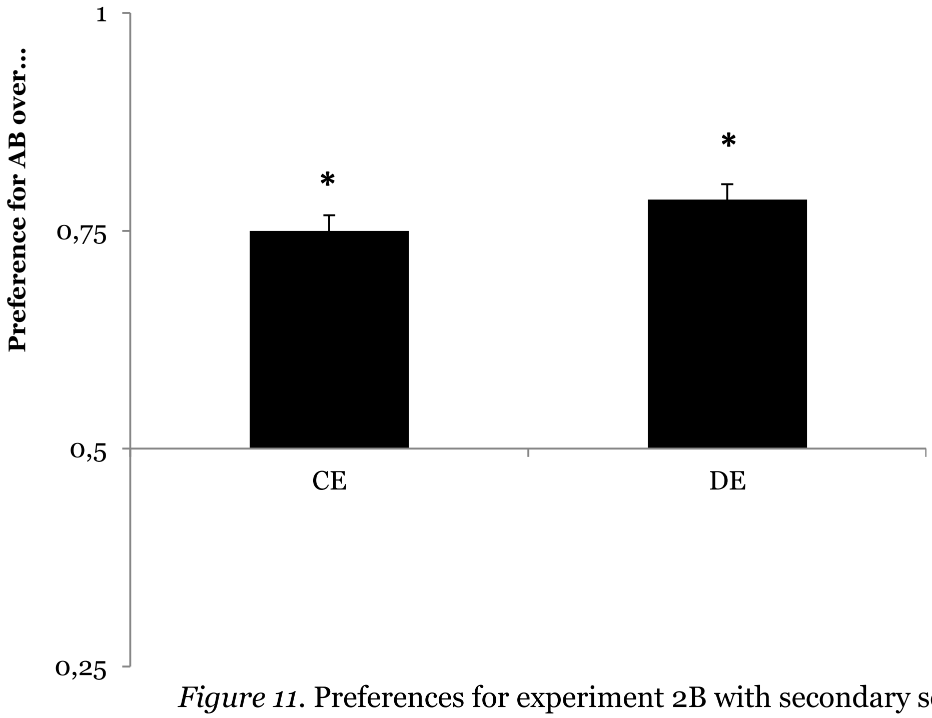

Figure 11. Preferences for experiment 2B with secondary school students.

The pattern of preference choice in AB vs. DE is at odds with the predictions of the noisy-OR integration rule again, but exactly as predicted by the linear integration rule. The preference for AB over CE fits again with the predictions of the linear integration rule.

Chapter 6. Summation under alternative scenarios and information formats

It has become apparent that the way information is provided to participants can have a large effect on the patterns of behavior observed in decision making (Barron & Erev, 2003; Hertwig, Barron, Weber, & Erev, 2004) and judgment domains (Fiedler, Brinkmann, Betsch, & Wild, 2000). The latest studies in this field (Bergert & Nosofsky, 2007; Bröder & Schiffer, 2003, 2006; Bröder & Newell, 2008; Platzer & Bröder, 2012) indicate that it is not the presentation format *per* se, but the *accessibility* of the information as determined by the presentation format which most influences which strategies are employed. Up to this point the information in my experiments was provided to the participants in form of probabilities. However, people do not always correctly use information presented as probabilities or percentages (Sloman, Over, Slovak, & Stibel, 2003).

Gigerenzer and Hofrage (1995) made a strong claim for an evolutionary approach to ecological rationality, proposing that humans evolved cognitive algorithms for making statistical inferences that are more prone to reason correctly with frequencies rather than probabilities. This argument relies on how the organism acquires causal information about the environment, by natural sampling of event frequencies. Previous research has shown that, from primates and lower animals to neural networks, all systems are highly sensitive to changes in the frequency of events (Brunswick, 1939; Gallistel, 1990; Real, 1991; Shanks, 1991), which seem to be encoded easily and automatically (Hasher & Zacks, 1984; Hintzman, 1976). By contrast while people directly observe frequencies they must compute probablities from experience. Indeed, as mentioned in the Introduction, probabilities were late inclusions into mathematics with the development of mathematical probability theory (Gigerenzer, Switjink, Porter, Daston, Beatty, & Krliger, 1989).

To explore the impact of this factor, Experiment 3A included information presented in a frequency format (e.g., participants were told that when cause A was present the effect occurred in 160 out of 200 occasions, instead of telling them that produced the effect 80% of the time). In Experiment 3B pie charts were used to present the probabilities in a graphical manner. Experiment 3C used graphical banners to present the numerical information. Finally, Experiment 3D used a cover story involving a coin tossing game with percentages.

Additionally, it is important to ensure that the previous results would generalize to other scenarios. Therefore three new cover stories were included in this series of experiments, one for Experiments 3A and 3B, a different one for Experiment 3C, and finally, Experiment 3D used a cover story based on a coin tossing game. Testing the generality of the previous results is not only important for methodological reasons, but it can also rule out alternative explanations. The results of our previous experiments are inconsistent with the predictions of the noisy-OR integration rule, but they can be accommodated assuming that the participants perceived an interaction between the elements, so that the effect of the combination might have a different outcome than the sum of its elements. Although according to the Power PC theory, that follows the noisy-OR integration rule, people naturally assume by default that causes do not interact with each other and this basic premise would only be rejected when there is strong disconfirming evidence. Otherwise, the problem of causal induction would become computationally intractable. Nonetheless, it is possible to assume that prior knowledge of our participants about how medicines and chemical substances work in general has led them to doubt that the various chemicals used in previous experiments do not interact with each other – so there actions that may be not independent. Daily experience shows that often, the effect of a drug may interact with other drugs. If the participants have assumed that something like this might be happening, then previous findings may not be valid tests of the models based on the noisy-OR integration rule.

6.1. Experiment 3A

A new cover story inspired by a short science-fiction story that we expected most participants to be unfamiliar with (Chiang, 2002) was used. In this cover story, the participants were asked to imagine that scientists have discovered a treatment that makes people insensitive to physical beauty. Some people would demand this experimental treatment to encourage them to be more sensitive to other people's inner beauty than to their physical beauty (see Appendix 2 in page 155 for further information). As in the previous pink-eyes experiments, participants were given information about the effectiveness of five different proteins that could produce this effect. In this experiment the probabilistic information was given as frequencies (e.g., "120 out of 200" people injected with protein Alfa stop perceiving physical beauty). Next, participants were asked to report which compounds they would choose, AB versus CE and AB versus DE, if they wished to stop being influenced by the physical appearance of other people.

6.1.1. Method

Participants, Procedure, and Design

Fifty-three undergraduate psychology students from University of Deusto volunteered to participate in Experiment 3A, 43 female and 10 male. The set of frequencies and materials used this experiment was the same as the one used in Experiment 1C (see Table 4).

Cue	***Cue in graphs***	***Probability***
Alpha	A	60%
Beta	B	65%
Gamma	C	95%
Delta	D	80%
Omega	E	30%

Table 6. Set of probabilities on experiment 3A.

6.1.2. Results

The choice pattern depicted in Figure 12 is similar to that of previous experiments. Contrary to the predictions of the noisy-OR integration rule, the participants preferred the compound AB to the compound DE (binomial test, $p < .001$). In this experiment the participants also preferred the compound AB to the compound CE (binomial test, $p= .005$), although this choice pattern is not predicted by either the noisy-OR or the linear integration rule.

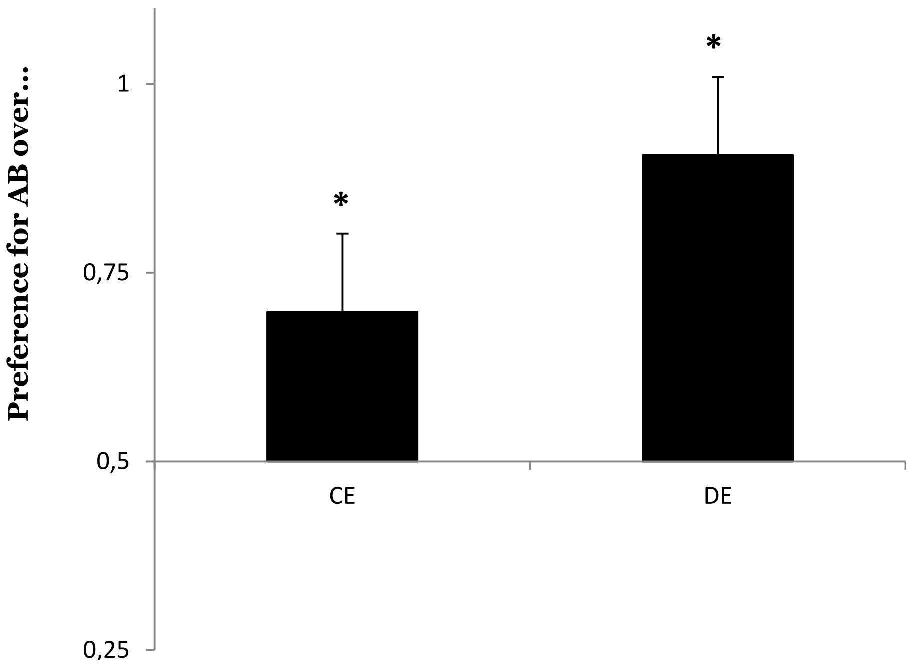

Figure 12. Preferences for experiment 3A with frequency format.

These results strengthen the idea that the data observed in previous experiments is not attributable to either peculiarities of the cover story or to the use of probabilities instead of frequencies.

6.2. Experiment 3B

There is an additional way to create probability versions that reveals the information of the experiment: to include a graphical representation that shows the probabilistic information. This technique was used in previous experiments in causal reasoning domain (Cosmides & Tooby, 1996; Sloman et. al, 2003; Yamagishi, 2003). Despite the widespread use of graphical displays and statistical graphs, the knowledge of how people read and interpret graphs is relatively meager. Hollands and Spence (1998) indicated that the bar graph most commonly used is not an effective display for judging proportion. Divided bars, reference bars, and pie charts generated generally better performance across their experiments. We included pie charts in this experiment because this display produced the smallest absolute error values. Regardless of some critics (Cleveland, 1985; Macdonald-Ross, 1977; Tufte, 1983), the accumulated evidence indicates that the pie chart is an effective format for judgments of proportion and yields accuracies as good as or better than other graph types for that task (Eells, 1926; Hollands & Spence, 1992; Simkin & Hastie, 1987; Spence, 1990; Spence & Lewandowsky, 1991).

In Experiment 3B the information was given by means of pie charts (see Figure 13). A pie chart is a circular chart divided into sectors, illustrating numerical proportion. In a pie chart, the arc length of each sector is proportional to the quantity it represents. In this experiment only two numerical proportions are depicted in each chart, corresponding to the probability of the effect and the complementary probability of no-effect for each of the protein treatments. The scenario used in this experiment was the protein treatment to stop perceiving physical beauty that was used in Experiment 3A. After seeing this information, the participants were asked to say which compounds they would choose to stop perceiving physical beauty.

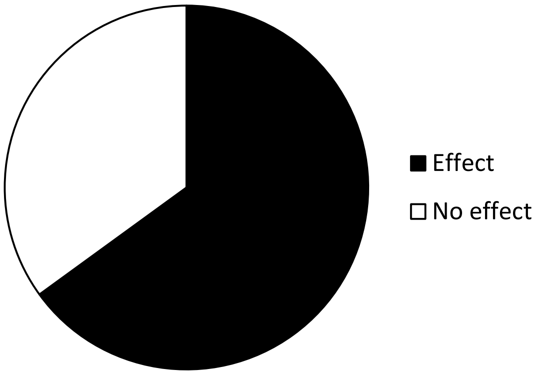

Figure 13. Pie chart used in experiment 3B to express the numerical information. The white part represents no effect of the given protein, and the black one denotes effect (stop perceiving the physical beauty in this scenario).

6.2.1. Method

Participants, Procedure, and Design

34 undergraduate students from different programmes volunteered for the experiment, 32 female and 2 male. The set of probabilities and materials are the same as in the previous Experiment 3A, and are the same as those used in the experiment 1C (see Table 4).

6.2.2. Results

The pattern of results, depicted in Figure 14, is exactly the one predicted by the linear integration rule. In contrast to the predictions of the noisy-OR rule, the participants preferred the compound AB to the compound DE (binomial test, p= .003). They showed no particular preference of either compound in the second preference question (binomial test, p = .392).

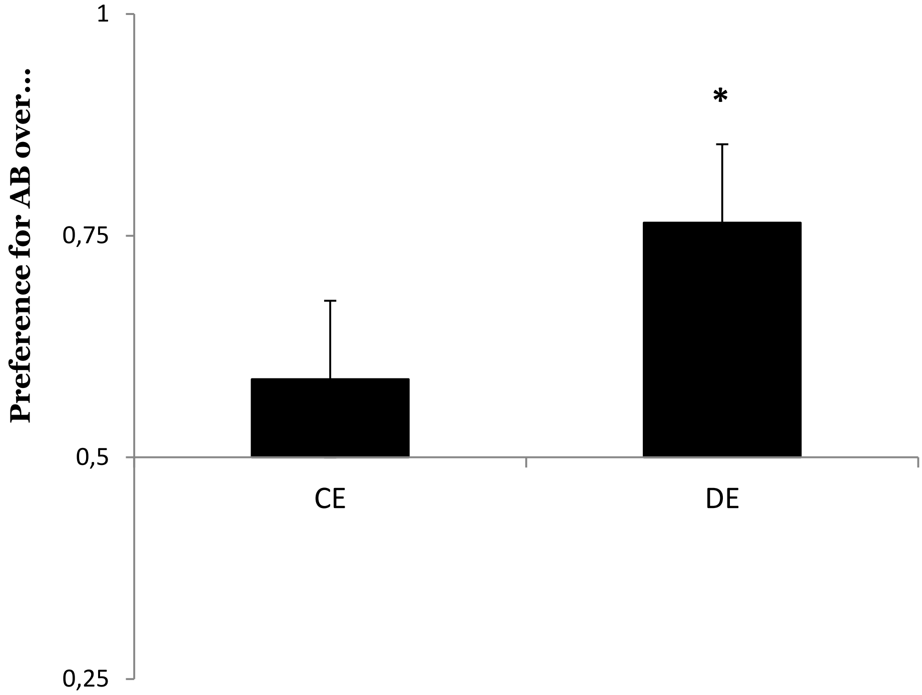

Figure 14. Results for experiment 3B with pie charts.

What can be concluded from this experiment is that using a graphical format to depict the probabilistic information did not change the pattern of behavior observed so far in this experimental series. Participants consistently differ from the predictions of the noisy-OR integration rule, preferring the compound AB over the compound DE.

6.3. Experiment 3C

In this experiment a new cover story derived from the Hyperion stories, described in the Introduction, was used. The two previous cover stories involved proteins, chemical substances and treatments, so it is possible that participants still perceiving interaction between the causes or use their previous knowledge about chemicals in order to derive their predictions. In this study, four new features were included: a new cover story, a new format for presenting information (graphical banners), preference questions involving a compound versus one candidate cause alone, and new questions involving the numerical judgment of the combined effectiveness of compounds involved in the preference questions.

6.3.1. Method

Participants, Procedure, and Design

Thirty participants volunteered for experiment 3C. All of them were from different courses, mainly engineering and medicine from the University of the Basque Country.

A new cover story related to the science fiction scenario presented in the Introduction was included. Participants were presented with an alien world scenario in which they had to learn about several risk factors or dangers that were more or less likely to kill humans. There were five different dangers represented by different names and ways of dying. Some of them always or never kill the explorer, whereas others sometimes were fatal and sometimes not (see Appendix 4 in page 159 for further information). After seeing this information, displayed as graphical banners based on frequencies of 24 (see Figure 15), participants answered preference questions, involving the choice of a route that contained one or two of the previously trained dangers. In this experiment, if participants perceive A and B as a very effective (the causes kill them), then they should choose the other alternative. For this reason, we need to reverse the interpretation of the preference data to make the results more directly comparable with those of the previous experiments.

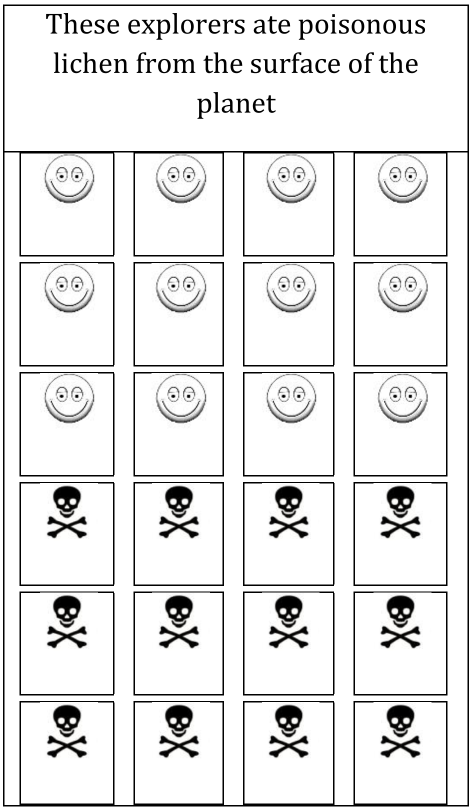

Figure 15. Graphical banners with information about the chances of living and dying out of 24 for experiment 3B with Hyperion scenario.

It is important to highlight that the results in the previous experimental series are quite consistent not only with the linear integration rule, but also with an averaging strategy that has not been mentioned thus far. According the averaging strategy, the causal impact of a compound would be equal to the average of the causal impact of its elements. Although the binary decision test used through the previous experiments was well designed for the purpose of contrasting linear and noisy-OR predictions, taking this issue into account we decided to introduce changes in the experimental design to test if the participants used some form of averaging strategy. In Experiment 3C we requested the same choice pattern data for the compounds but also we include the ratings of the probability of the effect given AB, CE and DE. Additionally, we asked participants to choose between a compound versus a cue alone, specifically between AB versus C and versus D. As depicted in Table 8, E element lacks any power to produce the effect (i.e., the probability of the effect in the presence of cue E is zero). Hence, removing it from any compound should have no impact on pattern choices: The participants' behavior should be the same when deciding between AB and CE and when to deciding between AB and C. Nevertheless, if removing this element has an impact on the pattern of choices made by participants, this might be of great interest to understand the strategy they used.

The linear and the noisy-OR rules predict no effect of adding the neutral compound E. If participants were linearly summing the independent probabilities, then they would still prefer the combination of causes over the cause alone. By contrast, if they were averaging the two elements within a compound, with this new experimental design (see Table 7), they would prefer the stronger cause alone over compounds of it and a weaker cause. Finally, probability ratings about the combination of causes were included in order to compare those estimates with the preference pattern.

Cue	*Probability*	*Number of outcomes*	*Preference Test*	*Predictions for compounds*
A	50%	12/24	AB vs CE	AB?
B	50%	12/24	AB vs C	CE?
C	75%	18/24	AB vs DE	DE?
D	100%	24/24	AB vs D	
E	0%	0/24		

Table 7. Set of probabilities and experimental design for experiment 3C, in Hyperion scenario using graphical banners.

6.3.2. Results

The preference responses for Experiment 3C are depicted in Figure 16. To facilitate comparison with the previous experiments in which the outcome was desired, the choices made by participants in this experiment (in which the outcome, death, was not desired) were reversed. Participants did not think that the combination of two causes with a 50% efficacy was more dangerous than the compound of a 75% effective cause and a 0% effective cause (AB vs. CE, binomial test, p=.585). Moreover, they thought the combination of the 100% effective cause and the 0% effective cause was more dangerous than the combination of the two 50% effective causes (AB vs. DE, binomial test, p<.001). The participants show indifference between the two 50% effective causes and the 75% cause (AB vs. C alone, binomial test, p=.099), and they significantly prefer the 100% effective cause over the two 50% effective causes (D over AB, binomial test, $p<.001$). This pattern of results is consistent with the noisy-OR integration rule. It is important to note that the pattern of choices was not influenced by the presence of the 0% cause, suggesting that in this experiment participants did not average the effect of several causes.

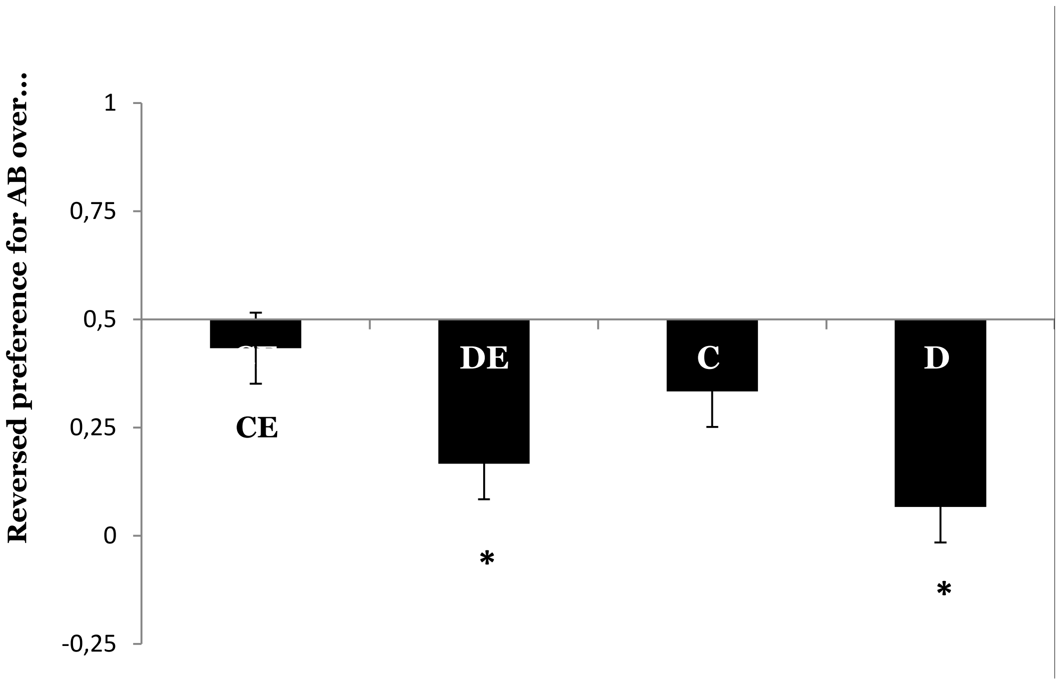

Figure 16. Choice results for Experiment 3C in Hyperion scenario using graphical banners.

As mentioned above, participants were also asked to rate the probability of the effect (dying) given several combinations of the causes A-E. Figure 17 shows the mean ratings of the probability of the effect given the combinations AB, CE, and DE. It is quite clear that these means are quite consistent with the noise-OR rule. Participants judged the combination of two 50% causes (AB) to generate a probability of around the expected 75% and likewise made similar judgments of the combination of a 0 and 75% probability cause (CE). Finally the judgments of the combination of the 100 % cause and the 0% (DE) cause were much higher.

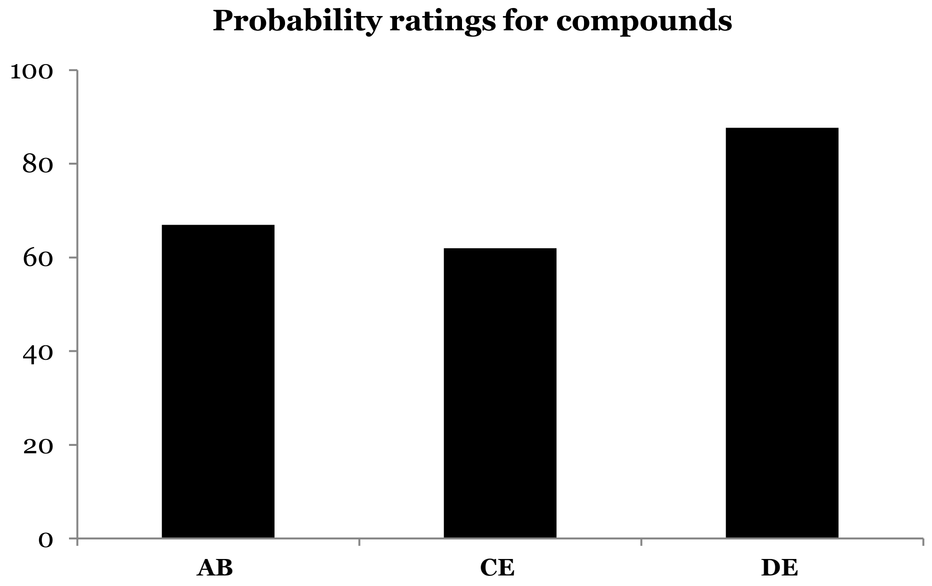

Figure 17. Probability ratings for the three compounds in Experiment 3C with Hyperion scenario using graphical banners.

Although these mean judgments are consistent with the noisy-OR rule they are not clear on the possible rules that each individual might use. For example a mean judgment of 75% for the AB compound might occur because everyone made judgments around 75% or because half of the participants summed the probabilities (50% +50% = 100%) while the other half chose one of the probabilities (here 50%). To shed some light on this we plot the number of people making judgments of individual frequencies or ranges of frequencies in the next three figures. Ranges are sometimes relevant because, while the noisy-or rule makes precise computational predictions (e.g. For AB = .5 + .5 - .5x.5 = .75), people may not make the exact computation but may reason that the result of two independent causes is somewhere between the larger of the two and their sum. So in this case a range between 50% and 100% might be relevant, although obviously less extreme values in the range are more convincing. Figure 18 which plots judgments of the AB compound with its expected value of 75% provides the most valuable insight into the strategies used by participants. It is important to note that, although the range depicted in the graph runs between 50 and 100, all the participants responded in this range gave 75% AB probability rating. As can be seen, most participants in this Experiment thought that the probability of the effect when two 50% effective causes were presented was 75%. This is consistent with the noisy-OR integration rule. Note, however, that one third of participants gave ratings equal to either 50 or 100, which are the values from the averaging and linear integration rules, respectively. Furthermore, 4 of the participants gave ratings that were outside the 50 to 100 range. Even in this experiment, the only one in the whole thesis that favors noisy-OR-based pattern of choices, only a small majority of participants adjusted their judgments predictions the noisy-OR integration rule.

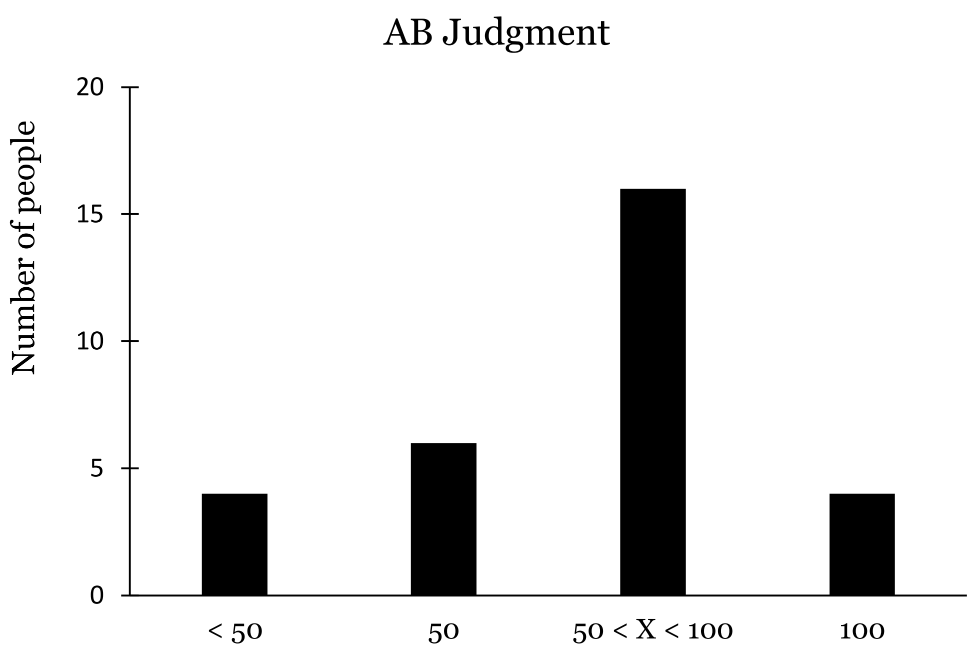

Figure 18. Judgment frequency for AB (50% + 50%) in Experiment 3C.

In Figure 19, probability judgments to compound CE are represented. It should be noted that the majority of people believed that the effect of the combination of a 75% effective cause plus a 0% effective cause is equal to the effect of the 75% effective cause in isolation. This is predicted by the linear integration rule and the noisy-OR rule but not the averaging heuristic. Although there are participants giving probability predictions below 75, it is unclear whether these participants were averaging the contribution of both causes because ratings were not always close to 37.5.

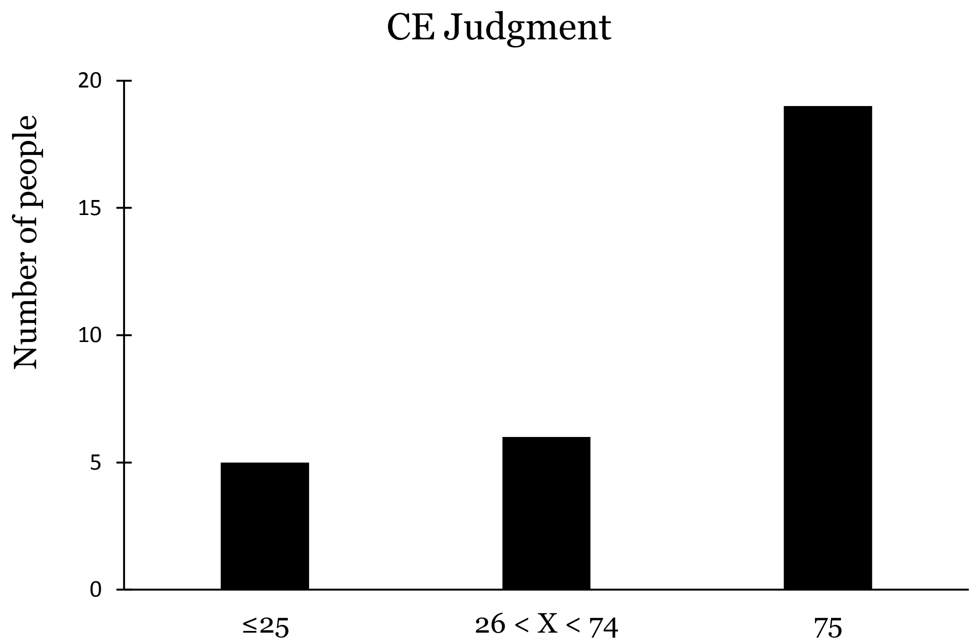

Figure 19. Judgment frequency for CE (75% + 0%) in Experiment 3C.

Finally, Figure 20 shows the probability predictions for compound DE. The majority of people correctly chose 100% for this compound. But again there is a nontrivial amount of people who did not do this.

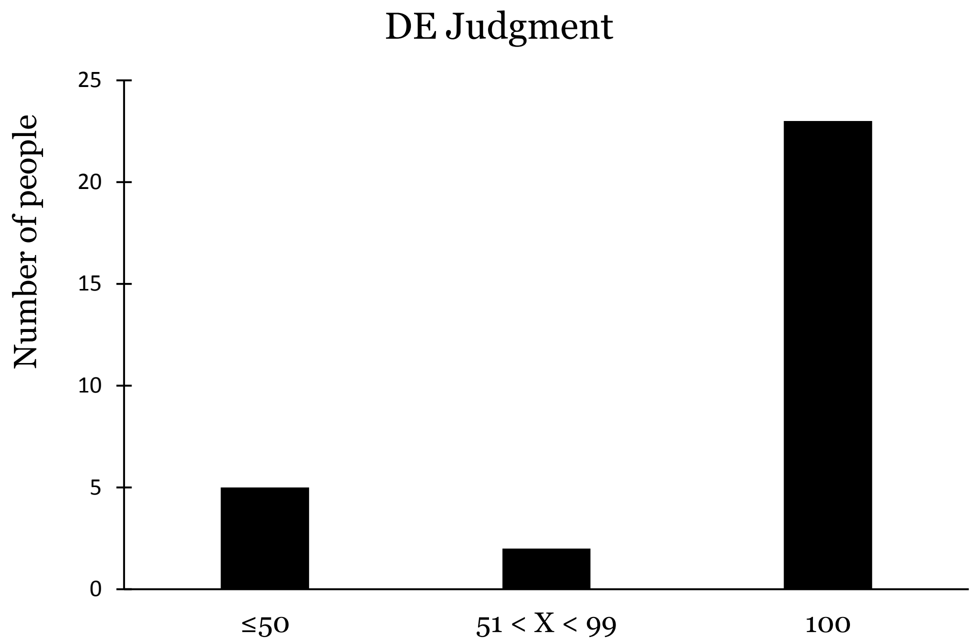

Figure 20. Judgment frequency for DE (100% + 0%) in Experiment 3C.

6.4. Experiment 3D

It is always possible that in each of our previous experiments participants assumed the causes were not independent and this precluded the noisy-OR strategy so in Experiment 3D we simply asked people to make judgments in the classic independent binary probability scenario – coin tossing. Participants were asked to imagine that they were playing a coin tossing game with three different biased coins. In this game, they would win if they obtained at least one head. They were told that would win regardless of whether they got one or two heads in the two coin tosses. Two heads are not better than one here. The three coins turned up heads 40%, 80%, and 64% of the times, respectively (see Appendix 3 in page 157 for further information). In the first test question participants were asked to choose whether they would prefer to toss the coin with the 40% chance of winning twice or to toss the coin with the 80% chance of winning once. In the second test, they had to choose whether they preferred tossing the 40% coin twice or tossing the 64% once. The noisy-OR rule predicts a probability of (.4 + .4 - .4X.4 = .64). So participants should prefer the single toss of the 80% coin and be indifferent with the 64% coin. After making these decisions, participants were asked to rate the probability of winning if they tossed the 40% coin twice.

6.4.1. Method

Participants, procedure and design

Forty-nine participants volunteered for experiment 3D. All of them were from different programs, mainly Psychology, from McGill University. See Table 8 below for the probabilistic information and experimental design for Experiment 3D.

Cues	*Probabilities*	*Preference Test*	*Predictions for compounds*
A	40%	A twice vs B	A twice?
B	64%	A twice vs C	
C	80%		

Table 8. Set of probabilities and experimental design for experiment 3D with coin tossing game scenario.

6.4.2. Results

The pattern of choices made by participants in Experiment 4D (Figure 21) was at first sight totally consistent with the predictions of the noisy-OR rule: They preferred to toss a single coin with 80% chance of winning over a double toss with a coin with 40% chance of winning. They were indifferent to the choice between tossing the 40% coin twice or the 64% coin once, which is also consistent with the noisy-OR rule.

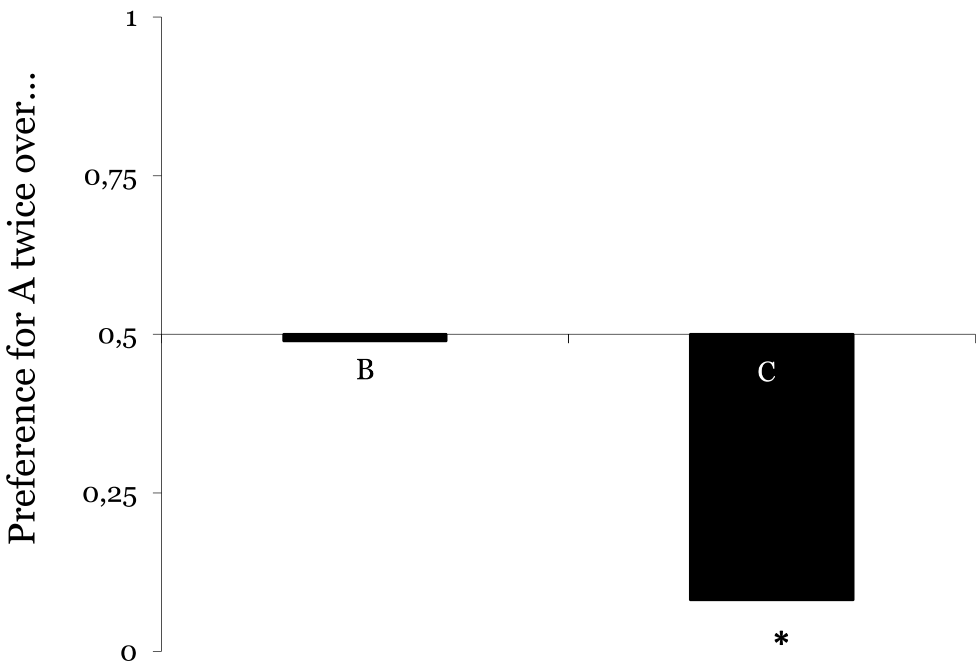

Figure 21. Preferences responses for 4D experiment with coin tossing game, depicting the choice pattern on 40% + 40% versus 64%, and 80%.

Figure 22 depicts the frequencies of the numerical ratings of the probability of winning tossing the 40%-winning coin twice. Their numerical ratings of the probability of winning if they tossed the 40% coin twice are represented in Figure 22. These ratings are in stark contrast with the pattern of decisions depicted in the previous graph. If participants were actually following the noisy-OR rule, one would expect most of ratings to fall in the range between 40 and 80. However, only 8 of the 49 participants gave responses within that very conservative range. Instead, 34 responded with either 40, which is consistent with either an averaging strategy or 'choose one or the other strategy' (as both predict 40%), or 80, which is the outcome if they are linearly summating both cues. According to the probability ratings in this a classical independent probabilities scenario, most participants are following either the averaging strategy or the linear integration rule and certainly not the noisy-OR rule even though the mean probability ratings are consistent with this. But, interestingly, when represented in a choice graph, these data indicate that the participants may be responding according to the predictions of the noisy-OR rule, which is not correct. These surprising results show that a pattern of choices that resemble a normative strategy might actually be the consequence of different, non-normative strategies used by different participants.

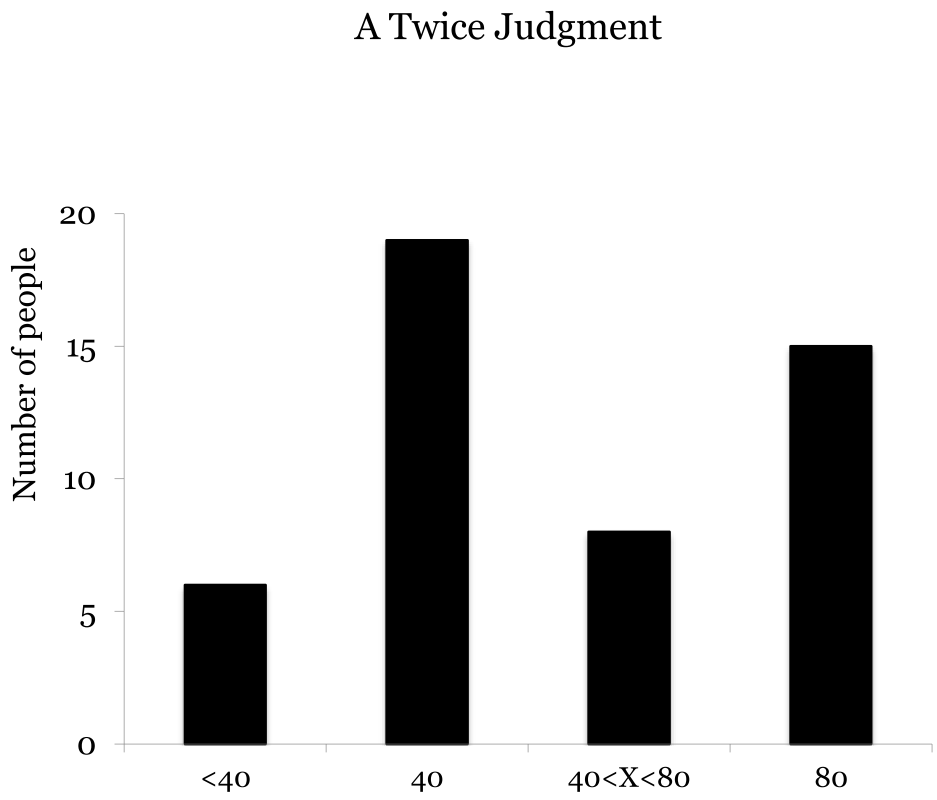

Figure 22. Probability ratings when coin 40% is double tossed.

People often have many strategies for solving a single problem. Averaging data generated by use of different strategies carries the same risks as averaging data generated by different individuals (Siegler, 1987). Just as data aggregated over people may not accurately reflect the behavior of any given person (Estes, 1956), data aggregated over different causal strategies may not accurately reflect the characteristics of any strategy.

Compound	*Cause 1*	*Cause 2*	*Linear*	*Noisy-OR*	*Average*
Experiments 1A, 2A, 3D					
AB	.40	.40	.80	.64	.40
CE	.80	.00	.80	.80	.40
DE	.64	.00	.64	.64	.32
Experiments 1B, 2B					
AB	.60	.65	1.25	.86	.625
CE	.97	.00	.97	.97	.485
DE	.86	.00	.86	.86	.43
Experiments 1C, 3A, 3B					
AB	.60	.65	1.25	.86	.625
CE	.95	.30	1.25	.965	.625
DE	.80	.30	1.10	.86	.55
Experiment 1D					
AB	.85	.40	1.25	.91	.625
CE	1.00	.25	1.25	1.00	.625
DE	.88	.25	1.05	1	.565
Experiment 3C					
AB	.50	.50	1.00	.75	.50
CE	.75	.00	.75	.75	.375
DE	1.00	0.00	1.00	1.00	.50

Table 9. Predictions made by several integration rules for all the experiments.

Chapter 7. General Discussion

'Not to be absolutely certain is, I think, one of the essential things in rationality.'

Bertrand Russell (1949)

7.1. Overall pattern of results

The results of the experiments in this thesis present an interesting picture. In Chapter 4, Experiment 1A explored the pattern of choices of participants with a given probability set using the pink-eyes scenario. The results perfectly fit the linear integration rule. In Experiment 1B, a new probability set was designed in a way that the sum of some cues is higher than 100%, to see if that change in the experimental design could promote a choice pattern consistent with the noisy-OR integration rule, but again the results were consistent with the linear predictions. Experiment 1C removed the 0% effective cause to avoid any particular effects of perceived inhibition associated with the non-effective cue. The results were as predicted by the linear integration rule. Finally, in Experiment 1D the introduction of a 100% effective cause made no difference to the pattern of results in this Chapter, and again participants responded according to the linear integration rule.

Chapter 5 explored causal reasoning in adolescents to see if the previous results could be due to some peculiarities of the university level participants, in the experiments in Chapter 4, who were mostly Psychology students trained in Probability. Experiment 2A replicated Experiment 1A with a sample of adolescents, and the same results favoring the linear strategy were obtained. Experiment 2B replicated Experiment 1B with another equivalent sample. The results where again consistent with the linear integration rule.

Chapter 6 introduced numerous changes in the cover stories and the probabilistic information formats. In Experiment 3A the information was presented as frequencies, and a new cover story was used to ensure the generality of the previous findings. The results were partially consistent with the linear integration rule, and did not fit the predictions of the noisy-OR strategy. In Experiment 3B pie charts were used to present the probabilistic information, using the same cover story as in Experiment 3A. The results were exactly as predicted by the linear integration rule. Experiment 3C included a number of changes: A new cover story, the Hyperion scenario, was used. Also, graphical banners based on frequencies of 24 were used as the information format. Finally, new questions involving the choice between a compound and a single cue alone were included to explore if the participants used some form of averaging strategy. Furthermore, probability ratings of the compounds were requested to provide more information regarding what strategies the participants used. The choice pattern in Experiment 3C was consistent with the noisy-OR integration rule, but the probability judgments did not always fit this rule. Experiment 3D employed a coin tossing game as a cover story, and the probabilistic information was given as percentages. In this experiment, the choice pattern was as predicted by the noisy-OR rule, but the probability judgments of the compounds suggested that the pattern of preferences consistent with noisy-OR strategy might actually be the consequence of different non-normative strategies (e.g., linear summation and averaging strategies) used by different participants. To sum up, this thesis explored a variety of probability sets, cover stories, and information formats to test if the participants behave according to the normative noisy-OR integration rule, or, by contrast, they employ non-normative strategies as the linear integration rule or an averaging strategy. Participants in the vast majority of the experiments throughout this thesis responded as predicted by the linear integration rule.

Our experimental work suffers from some problems. It is impossible to know what factor is the key one to modify the participants' strategies. The presentation format and the causal scenarios have not been independently manipulated, so, in a basis of these results, both the normative models and the heuristic approaches still have arguments supporting their views. One of the next steps to do would be to test these factors in isolation to see which of them is responsible for any changes in the behavior of participants.

Another important issue to be addressed is the lack of consistency between the choice and probability patterns. It seems to be pretty obvious that asking simultaneously preference choices and probability judgments can promote different rules, even in the same participant and with the same material. At this point, the factors that make people choose one rule over another remain unclear.

7.2. Formats of presentation

It has been widely claimed that people are much more competent in reasoning about frequencies than probabilities (Cosmides & Tooby, 1996; Gigerenzer & Hoffrage, 1995). Despite the extensive empirical evidence offered by these two studies, other authors have demonstrated that the use of frequencies rather than probabilities does not in itself make reasoning problems easier to solve when confounding factors were eliminated (Barbey & Sloman, 2007; Evans, Handley, Perham, Over, & Thompson, 2000; Griffin & Buehler, 1999; Neace, Michaud, Bolling, Deer, & Zecevic, 2008; Macchi, 2000; Mellers & McGraw, 1999; Reyna & Mills, 2007; Sloman et al., 2003; Waters, Weinstein, Colditz, & Emmans, 2006). This research shows that the participants' responses to such tasks are strongly influenced by subtle variations in the presentation of task information. Thus, there is little evidence to support the idea that frequencies per se, when not confounded with other factors, are more natural or easier to comprehend than percentages or other 'normalized' formats. Despite this mixed evidence, different presentation formats was an important factor to be tested within the context of our experiments.

The results obtained from different presentation formats did not significantly affect the results in the experiments throughout this thesis. A frequency format to depict the probabilistic information was used in Experiment 3A, and the results were in line with the linear integration rule. Similarly, Experiment 3B used pie charts to express the probability set, and the participants again behaved according to the linear summation rule. The last two experiments might be of more interest to this issue: In Experiment 3 the choice pattern changed from the linear to the noisy-OR integration rule, but not only was the presentation format changed (to graphical banners) but also a new cover story regarding the Hyperion world, with some potential dangers and an outcome of dying was used. Thus, we cannot conclude that the changes in choice patterns are due solely to the presentation format. It is important to mention that several of the experiments reported by Patricia Cheng and colleagues used this type of graphical banners, with different results to our experiments (Buehner, Cheng, & Clifford, 2003; Carroll & Cheng, 2010; Cheng & Buehner, 2012; Cheng, Novick, Liljeholm, & Ford, 2007; Liljeholm & Cheng, 2007, 2009; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008; Novick & Cheng, 2004).

Also, when analyzing the probability ratings given to the compounds, it is apparent that the participants used a mix of strategies, most of which were not consistent with the noisy-OR strategy rule. Finally, and in line with this previous finding, in Experiment 3D the information was presented as percentages, as in the four experiments in Chapter 4 which favored linear integration rule, but in this experiment the pattern of choices resembled a noisy-OR normative strategy. In this experiment the format did not change, the only novelty was a new cover story involving a coin tossing game. Hence, if anything, this finding could suggest that the cover stories employed in causal reasoning studies have more impact in the pattern of choices made by participants than the way the information is presented. It is important to mention that in this experiment although the choice pattern followed the noisy-OR rule's predictions, the probability ratings of the compounds demonstrated that most participants followed the non-normative linear and averaging strategies.

7. 3. Numeracy

Making good decisions in the real world requires some numerical ability. Although many causal judgments and decisions rely heavily on understanding basic mathematical concepts, little research has examined the role of numerical ability, or numeracy, in decision making (Peters, Vastfjall, Slovic, Mertz, Mazzocco, & Dickert, 2006; Reyna & Brainerd, 2007). Numeracy is defined as the ability to understand and use basic probability and numerical concepts to perform rudimentary mathematical operations, compare magnitudes, and comprehend ratio concepts (including fractions, proportions, percentages and probabilities). While it may appear that they are closely related, numeracy is not due to general intelligence (Peters et al., 2006; Reyna, Nelson, Han, & Dieckmann, 2009).

Research in numeracy suggests that people differ substantially in numerical ability (Lipkus, Samsa, & Rimer, 2001; Woloshin, Schwartz, Black, & Welch, 1999) and that many people are even 'innumerate' (Paulos, 1988). In today's increasingly technological world, when people have unprecedented access to numerically expressed information, low numeracy and most importantly innumeracy seems to be a critical obstacle to make good decisions in causal, financial and medical domains, among others. A number of studies (Fagerlin, Ubel, Smith, & Zikmund-Fisher, 2007; Keller, 2011; Lipkus & Peters, 2009) show that higher level of numeracy is related to better comprehension and integration of numerical information, usually leading to more informed and better decisions. More relevant to the present experiments, research on the effect of numeracy and presentation formats on risk perceptions and medical decision making suggests that less numerate participants were sensitive to the changes in presentation format, whereas highly numerate individuals were relatively unaffected (Dieckmann, Slovic, & Peters, 2009; Dickert, Kleber, Peters, & Slovic, 2011; Peters et al., 2006, 2011). This conclusion could be extrapolated to the results of Experiments 2A and 2B which studied adolescents who were presumably naïve to probability theory. We conducted replications of Experiments 1A and 1B, originally conducted with Psychology students. The results were exactly the same in the adolescents sample and Psychology students' sample. This finding can have two plausible explanations: One, in our experiments numeracy is not a crucial factor. Or two, if numeracy is a key factor, in Spanish educative system the level of numeracy obtained in adolescence is enough to obtain the same pattern of choices observed in University students.

It would have been interesting to measure numeracy across our experiments to check whether participants who scored low on numeracy were those that changed their strategies of decision making based on changes on different formats of presenting numerical information and cover stories, as found in Experiments 3C and 3D.

7.4. The independence assumption

Despite our efforts to prevent the participants from believing that there were interactions between causes, it is possible that we did not achieved this goal in the earlier experiments. If that were the case, our results could be consistent with the noisy-OR perspective: When the causes are clearly independent the responses are more in line with the noisy-OR rule predictions. It is interesting to note that we have two causal scenarios (the coin tossing game and that Hyperion scenario) that may more clearly promote the perception of the independence principle assumed by default in models like Cheng's (1997) Power PC theory. One could argue that the change in the choice pattern found in Experiments 3C and 3D could be due to the type of causal scenario employed in these two last experiments, with scenarios that emphasize the lack of interaction between the causes whereas in the pink eyes scenario chemical substances and drugs are more prone to lead the participants to assume interactions between them, because of previous knowledge about them. However, this assumption is hard to reconcile with the results of other studies favoring the noisy-OR rule that also involved causal stories regarding drugs and medical substances (Liljeholm & Cheng, 2007, 2009). Given the mixed evidence in the literature about causal learning models employing different integration strategies, one of the next steps might be to include a direct measure of whether or not participants believe there is an interaction between the potential causes.

7. 5. Implications for causal learning models

Considering the Bayesian perspective briefly covered in the Introduction, it should be mentioned that the experiments in this thesis were not designed to directly measure if participants' behavior is consistent with Bayesian models. But these models can explain our pattern of data attending to the flexibility in the strategies that our participants seem to have in all the experiments. The fact that participants seem to be sensitive to factors like the cover story or the presentation format can be understood as some type of previous knowledge resulting in an ability to choose the integration strategy that subjects deemed most appropriate in the situation where they are. Bayesian models reflect well this flexibility (Lu, Rojas, Beckers & Yuille, 2008). But this does not solve all the problems. The fact that Bayesian models are very powerful and flexible, incorporating this prior knowledge that has been proven so crucial, can make them less falsifiable than alternative models, given that they can represent one thing and its opposite. These models are usually unconstrained, given that they make few a priori predictions but a posteriori can account for nearly any result. It is not surprising that some authors (Lee, 2010) have suggested that Bayesian inference should be used now as a statistical method, not a model of mind.

In the earlier Causal Learning Models section, configural models have been mentioned but not covered in detail. These configural models can account for some of the experimental data obtained from the present experiments. Configural models (Kruschke, 1992; Pearce, 1987, 1994) propose that compound stimuli are processed not as separate entities, but as unique exemplars that form associations that are independent of those formed by their constitutive elements. When a novel compound is presented for the first time and the participant has not learned anything about it, what has been learned about the elements is generalized to the compound based on the similarity of the individual elements to the compound. As a consequence of this generalization, the configural view predicts that the strength of a given compound will approximate the average strength of each of its components in isolation, whereas the elemental approach predicts that the strength of the compound will be higher than the strength of either component presented in combination. In our experiments, the similarity between A and AB would be .50, because half of the elements of AB compound are present in A as well, and the same with B and AB. In many of our experimental probability sets, the strength of A is .40 and the strength of B is again .40, then the associative strength of the compound AB would be .40 [(.40 · .50) + (.40 · .50)]. The integration rule used by these configural models could potentially explain why many participants showed a pattern of preference responses consistent with an averaging strategy. This strong tendency to use the averaging strategy might have been developed by practice, because this strategy has been useful in the past.

A number of studies conducted over the past years have demonstrated that the way in which stimulus are processed is not fixed beforehand, and that a number of factors can heavily influence whether they are processed in an elemental or configural way. These factors include experimental instructions (Lipp, Cox, & Siddle, 2001; Mitchell & Lovibond, 2002), prior experience (Mehta & Williams, 2002; Williams & Bracker, 1999), stimulus properties (Kehoe, Horne, Horne, & Macrae, 1994; Lachnit, 1988; Myers, Vogel, Shin, & Wagner, 2001; Rescorla & Coldwell, 1995), stimulus organization (Glautier, 2002; Martin & Levey, 1991), and task demands (Lachnit & Kimmel, 1993; Lober & Lachnit, 2002, Shanks & Darby, 1998). The literature indicates that sometimes people process information as configurations, and sometimes as elements. In our experiments, on the probability ratings for compounds, we also showed that people sometimes average, and sometimes summate (using linear or noisy-OR summation). Thus, these results are consistent with the range of results that have been observed in the literature on configural versus elemental processing.

The results of our experiments are problematic for models of causal learning that rely exclusively on either the linear integration rule or the noisy-OR integration rule. A challenging possibility that has been proposed (Melcher, Shanks & Lachnit, 2008; Williams, 1995) is that stimulus processing involves both elemental and configural coding (Fanselow, 1999; Pearce & Bouton, 2001; Wagner, 2003), and that the predominance of one or another is a function not only of the task demands, but also of the individuals' learning history. Consistent with this assumption, there is some evidence that training with elemental discrimination disposes participants to use elemental coding in subsequent tasks, whereas training with configural discriminations predisposes participants to use configural coding in subsequent tasks (Williams & Braker, 1999; Williams et al., 1994). Because these assumptions are less parsimonious than a single rule it is imperative to be able to predict which rule will be used and when.

7.6. Previous experiments on summation

A small number of studies have reported evidence favoring summation and elemental processing in human causal learning (Van Osselaer, Janiszewski & Cunha, 2004; Collins & Shanks, 2006). Interestingly, these two studies that have found summation in human causal learning used a task in which the outcome during training was quantitative and submaximal. By contrast, Soto, Vogel, Castillo and Wagner (2009) reported evidence of summation in human causal learning in three experiments. They tried to find out whether the nature of the outcomes in the training phase (magnitude or categorical, maximal or submaximal) is a critical variable that can promote elemental-like results. The results suggest that summation is a relatively general finding in human causal learning, independent of the specific nature of the outcome used during training (binary or magnitude) and of the kind of scale used to measure causal ratings during test (magnitude or likelihood of the outcome), not as Collins and Shanks (2006) suggested.

The observance of summation in human causal learning is not peculiar to training and testing with different magnitude, but can be obtained with binary outcomes in training and estimations of likelihood of the consequence in testing. What appears to be important is that the training and testing conditions allow the participants to give submaximal causal ratings to the elemental stimuli, that is, to avoid producing such high ratings to the elemental stimuli alone that summation is precluded by a ceiling effect. These results are more easily explained by elemental theories of associative learning (Mackintosh, 1975; Rescorla & Wagner, 1972; Wagner, 1981), which suppose that the associative strength of a compound is the sum of the associative strengths acquired by each of its components. They would generally be considered to be at odds with configural learning theories (Estes, 1994; Pearce, 1987), which argue that the associative strength of a compound should be the same or lower than that acquired by its individual elements.

As can be seen along all the experimental work presented here, summation is unstable. Every single change in the experimental procedure has impact on the data. This is consistent with previous findings as well. For instance, van Osselaer et al. (2004) found that minor changes in the procedure when collecting judgments or presenting the information could result in the suppression of the summation. Later, Steven Glautier and colleagues (2010) found that when the similarity between two potential causes is increased by means of adding a common feature, summation is reduced, and can be even reversed. Interestingly, this reduction is logical starting from an averaging strategy: If participants average the causal power of each candidate cause, then the sum need not be higher than either of its parts. Most interestingly to the present experiments, Waldmann (2007) demonstrated that minor changes in the experimental procedure could influence whether participants combined causes in a basis of an averaging or an additive (linear or noisy-OR) integration rule. By contrast, when the experimental task involves subtracting causal powers instead of addition, participants' responses tended to be those predicted by an additive rule, even when the same instructions were used. This is consistent with the divergences we found between the pattern of preferences and the probability ratings in some of our experiments.

This conclusion is not very different from the one reached by previous experiments, participants are not fixed in employing a single causal induction strategy. Rather, qualitative differences in responding can be observed though subtle changes in the experimental environment such as judgment intervals (Collins & Shanks, 2002), magnitude of the effect (De Houwer, Beckers & Glautier, 2002), and the type of question that probes causal knowledge (Matute, Arcediano & Miller, 1996; Matute, Vegas & De Marez, 2002). Although there are also plenty of well-designed experiments suggesting that people are able to behave according to the Power PC assumptions, it is unclear that people use the same reasoning strategy in everyday-life situations. In everyday life, most problems of causal induction are full of confounding variables beyond the capacity of lay people to isolate. Making sure that the assumptions required for the inferential rule hold is difficult, as best.

To sum up, if we have to highlight just one finding from our experiments is that they clearly show that the notion that people's default assumption is the independence of the causes is incorrect in a variety of causal scenarios. And therefore, the claim of the noisy-OR rule as the default strategy is inconsistent with the results of the experiments in this thesis.

Chapter 8. References

'I have always imagined that Paradise will be a kind of library.'

Borges

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society, 15*, 147-129.

Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Review, 114*(3), 435-448.

Allan, L. G. (2003). Assessing power PC. *Learning & Behavior, 31*(2), 192-204.

Allan, L. G., Siegel, S., & Tangen, J. M. (2005). A signal detection analysis of contingency data. *Learning & Behavior*, *33*(2), 250-263.

Aydin, A., & Pearce, J. M. (1994). Prototype effects in categorization by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes, 20*(3), 264-277.

Aydin, A., & Pearce, J. M. (1995). Summation in autoshaping with short and long-duration stimuli. *Quarterly Journal of Experimental Psychology, 48*B, 215-234.

Aydin, A., & Pearce, J. M. (1997). Some determinants of response summation. *Animal learning & behavior, 25*, 108-121.

Baker, A. G., Berbrier, M., & Vallée-Tourangeau, F. (1989). Judgments of a 2 x 2 contingency table: Sequential processing and the learning curve. *Quarterly Journal of Experimental Psychology, 41B*, 65-97.

Baker, A. G., Mercier, P., Vallée-Tourangeau, F., Frank, R., & Pan, M. (1993). Selective associations and causality judgments: Presence of a strong causal factor may reduce judgments of a weaker one. *Journal of Experimental Psychology: Learning, Memory and Cognition, 19*(2), 414-432.

Barberia, I., Baetu, I., Sansa, J., & Baker, A. G. (2010). Choosing optimal causal backgrounds for causal discovery. *The Quarterly Journal of Experimental Psychology, 63*(12), 2413-2431.

Barberia, I., Baetu, I., Sansa, J., & Baker, A. G. (2014). When is a cause the "same"? Incoherent generalization across contexts. *Quarterly Journal of Experimental Psychology, 67,* 281-303.

Barbey, A. K., & Sloman, S. A. (2007). Base-rate respect: From ecological rationality to dual processes. *Behavior and Brain Sciences, 30*, 241-254.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making, 16(3)*, 215-233.

Bergert, F. B., & Nosofsky, R. M. (2007). A responsetime approach to comparing generalized rational and take-the-best models in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 107-129.

Bennet, D. J. (1998). *Randomness*. Cambridge, MA: Harvard University Press.

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,* 333-340.

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,* 1290-1304.

Bröder, A., & Newell, B. R. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making, 19*, 205-214.

Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General, 132*, 277-293.

Bröder, A., & Schiffer, S. (2006). Stimulus format and working memory in fast and frugal strategy selection. *Journal of Behavioral Decision Making, 19*, 361-380.

Brunswick, E. (1939). Probability as a determiner of rat behavior. *Journal of Experimental Psychology, 25*, 175-197.

Buehner, M. J., & Cheng, P. W. (2005). Causal learning. In K. J. Holyoak & R. Morrison (Eds.), *Handbook of thinking & reasoning* (pp. 143-168). Cambridge, England: Cambridge University Press.

Buehner, M. J., & Cheng, P. W. (1997). Causal induction: The Power PC Theory versus the Rescorla-Wagner theory. In M. G. Shafto, & P. Langley, *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (pp. 55-60). Mahwah, NJ: Erlbaum.

Buehner, M. J., & May, J. (2002). Knowledge mediates the timeframe of covariation assessment in human causal induction. *Thinking & Reasoning, 8*(4), 269-295.

Buehner, M. J., & McGregor, S. J. (2009). Contingency and contiguity trade-offs in causal induction. *International Journal of Comparative Psychology, 22*, 19-42.

Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 1119-1140.

Buehner, M. J., & Cheng, P. W. (2005). Causal learning. In K. J. Holyoak & R. G. Morrison (Eds.). *Cambridge handbook of thinking and reasoning* (pp. 143-168). Cambridge, England: Cambridge University Press.

Carroll, C. D., & Cheng, P. W. (2010). The induction of hidden causes: Causal mediation and violations of independent causal influence. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society,*eds. Ohlsson, S., & Catrambone, R. (pp. 913-918). Austin, Texas: Cognitive Science Society.

Chapman, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17*, 837-854.

Chapman, L. J., & Champan, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology, 74*(3), 271-280.

Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition, 18*, 537-545.

Cheng, P. W. (1993). Separating causal laws from causal facts: Pressing the limits os statistical relevance. In D. L. Medin, *The psychology of learning and motivation* (Vol. 30, pp. 215-264). New York: Academic Press.

Cheng, P. W. (1997). From covariation to causation: A causal power Theory. *Psychological Review, 104*(2), 367-405.

Cheng, P. W., & Buehner, M. J. (2012). Causal learning and inference. In Holyoak, K., and Morrison, R. (Eds.), *Oxford handbook of thinking and reasoning*. Oxford: Oxford University Press.

Cheng, P. W., & Holyoak, K. J. (1995). Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In J. -A. Meyer, & H. Roitblat, *Comparative approaches to cognition* (pp. 271-302). Cambridge, MA: MIT Press.

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58*, 545-567.

Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition, 40*, 83-120.

Cheng, P. W., & Novick, L. R. (1992). Covariation in natural causal induction. *Psychological Review, 99*(2), 365-382.

Cheng, P. W., & Novick, L. R. (2005). Constraints and nonconstraints in causal learning: Reply to White (2005) and to Luhmann and Ahn (2005). *Psychological Review, 112*(3), 694-707.

Cheng, P. W., Novick, L. R., Lilhejolm, M., & Ford, C. (2007). Explaining four psychological assymmetries in causal reasoning: Implications of causal assumptions for coherence. In M. O'Rourke, *Topics in Contemporary Philosophy (Vol. 4): Explanation and Causation*. Cambridge, MA: MIT Press.

Chiang, T. (2002). *Stories of your life and others. New York: Thor Books.*

Cleveland, W. S. (1985). *The elements of graphing data*. Monterey, CA: Wadsworth.

Collins, D. J., & Shanks , D. R. (2002). Momentary and integrative response strategies in causal judgment. *Memory & Cognition, 30*, 1138-1147.

Collins, D. J., & Shanks, D. R. (2006). Conformity to the power PC theory of causal induction depends on the type of probe question. *The Quarterly Journal of Experimental Psychology, 59*(2), 225-232.

Collins, D. J., & Shanks, D. R. (2006). Summation in causal learning: Elemental processing or configural generalization? *Quarterly Journal of Experimental Psychology, 59*, 1524-1534.

Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition, 58*, 7-73.

Crocker, J. (1981). Judgment of covariation by sociel perceivers. *Psychological Bulletin, 90*, 272-292.

Danks, D. (2003). Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology, 47,* 109-121.

Danks, D. (2005). The supposed cometition between theories of Human Causal inference. *Philosophical Psychology, 18*(2), 259-272.

Danks, D., Griffiths, T. L., & Tenenbaum, J.B. (2003). Dynamic causal learning. In S. Becker, S. Thrun, & K. Obermayer, eds. *Advances in Neural Information Processing Systems* (pp. 67-74).Cambridge, Mass: The MIT Press.

Darredeau, C., Baetu, I., Baker, A. G., & Murphy, R. A. (2009). Competition between multiple causes of a single outcome in causal reasoning. *Journal of Experimental Psychology: Animal Behavior Processes, 35*(1), 1-14.

De Houwer, J. (2002). Forward blocking depends on retrospective inferences about the presence of the blocked cue during the elemental phase. *Memory & Cognition, 30*, 24-33.

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, 55B*(4), 289-310.

De Houwer, J., & Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *The Quarterly Journal of Experimental Psychology, 55A*, 965-985.

Devlin, K. (2008). In *The Unfinished Game: Pascal, Fermat, and the Seventeenth-Century letter that made the world modern*. New York: Basic Books.

Dickert, S., Kleber, J., Peters, E., & Slovic, P. (2011). Numeracy as a precursor to pro-social behavior. The impact of numerac and presentation format on the cognitive mechanisms underlying donation decisions. *Judgment and Decision Making, 6(7)*, 638-650.

Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgments. *Quarterly Journal of Experimental Psychology, 49,* 60-80.

Dickinson, A., Shanks, D. R., & Evenden, J. L. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology, 36A*, 29-50.

Dieckmann, N. F., Slovic, P., & Peters, E. M. (2009). The use of normative evidence and explicit likelihood by decision makers varying in numeracy. *Risk Analysis, 29*, 1473-1488.

Diecks, D., González, W. J., Hartmann, S., Stöltzner, M., & Weber, M. (2012). *Probabilities, laws, and structures*. New York: Springer. ISBN 978-94-007-3029-8.

Edwards, W. (1954). The theory of decision making. *Psychological Bulletin, 51,* 380-417.

Eells, W. C. (1926). The relative merits of circles and bars for representing component parts. *Journal of the American Statistical Association, 21,* 119-132.

Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin, 53,* 134-140.

Estrada, C., Barnes, V., Collins, C., & Byrd, J. C. (1999). Health literacy and numeracy. *Journal of the American Medical Association, 282*(6)*,* 527.

Evans, J. B. T., Handley, S. J., Perham, N., Over, D. E., & Thompson, V. A. (2000). Frequency versus probability formats in statistical word problems. *Cognition, 77,* 197-213.

Fagerlin, A., Uber, P. A., Smith, D. M., & Zikmund-Fisher, B. J. (2007). Making numbers matter: Present and future research in risk communication. *American Journal of Health Behavior, 31*, 47-56.

Fanselow, M. S. (1999). Learning theory and neuropsychology: Configuring their disparate elements in the hippocampus. *Journal of Experimental Psychology: Animal Behavior Processes, 25*(3), 275-283.

Fiedler, K., Brinkmann, B., Betsch, T., & Wild, B. (2000). A sampling approach to biases in conditional probability judgments: Beyond base rate neglect and statistical format. *Journal of Experimental Psychology: General, 129*(3), 399-418.

Franklin, J. (2001). *The science of conjecture: Evidence and probability before Pascal.* Baltimore, MD: John Hopkins University Press.

Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: MIT Press.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review, 102*(4), 684-704.

Gigerenzer, G., Switjink, Z., Porter, T., Daston, L., Beatty, J., & Krliger, L. (1989). The empire of chance: How probability changed science and everyday life. Cambridge, England: Cambridge University Press.

Glautier, S. (2002). Spatial separation of target and competitor cue enhances blocking of human causality judgments. *Quarterly Journal of Experimental Psychology, 55*, 121-135.

Glautier, S., Redhead, E., Thorwart, A., & Lachnit, H. (2010). Reduced summation with common features in causal judgments. *Experimental Psychology, 57*(4), 252-259.

Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General, 117*(3), 227-247.

Glymour, C. (2003). Learning, prediction and Causal Bayes nets. *Trends in Cognitive Sciences, 7*(1), 43-48.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review, 111*(1), 1-31.

Griffin, D., & Buehler, R. (1999). Frequency, probability and prediction: Easy solutions to cognitive illusions? *Cognitive Psychology, 38*, 48-78.

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strenght in causal induction. *Cognitive Psychology, 51*, 334-384.

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review, 116*(4), 661-714.

Hacking, I. (2006). *The emergence of probability*. New York: Cambridge University Press, Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General, 138*(1), 22-38.

Hahn, U., & Warren, P. A. (2009). Perception of randomness; Why three heads are better than four. *Psychological Review, 116*(2), 454-461.

Hald, A. (2003). *A history of probability and statistics and their applications before 1750*. Hoboken, N. J.: Wiley.

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, 63,* 59-73.

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist*, 39, *1372-1388.*

Heath, T. L. (1956). *The thirteen books of Euclid's elements*. New York: Dover Publications.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science, 15*(8), 534-539.

Hintzman, D. L. (1976). Repetition and memory. En G. H. Bower. New York: Academic Press.

Hollands, J. G., & Spence, I. (1992). Judgments of change and proportion in graphical perception. *Human factors, 34,* 313-334.

Hollands, J. G., & Spence, I. (1998). Judging proportion with graphs: The summation model. *Applied cognitive psychology, 12,* 173-190.

Holyoak, K. J., & Cheng, P. W. (2011). Causal learning and inferece as a rational process: The new synthesis. *Annual Review of Psychology, 62*, 135-163.

Hume, D. (1978). A treatise of human nature (Original work published 1739). Oxford: Oxford University Press.

Hume, D. (2008). En *An enquiry concerning human understanding*. New York: Cosimo (Original work published 1748).

Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychologhical Monographs, 79*, (whole number X).

Jonassen, D. H., & Ionas, I. G. (2008). Designing effective supports for causal reasoning. *Education Technology Research and Development, 56*, 287-308.

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

Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.). *Miami symposium on the prediction of behaviour: aversive stimulation* (pp. 9-31). Miami: University of Miami Press.

Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (eds). *Punishment and aversive behavior,* 279-296. New York: Appleton - Century - Crofts.

Kant, I. (1781). *Critique of pure reason*. N. K. Smith. London: MacMillan.

Kehoe, E. J., Horne, A. J., Horne, P. S., & Macrae, M. (1994). Summation and configuration between and within sensory modalities in classical conditioning of the rabbit. *Animal Learning & Behavior, 22*, 19-26.

Keller, C. (2011). Using a familiar risk comparison within a risk ladder to improve risk understanding. *Risk Analysis, 31,* 1043-1054.

Kolmogorov, A. N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnun*. Berlin: Julius Sprinfer.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22-44.

Lachnit, H. (1988). Convergent validation of information processing constructs with Pavlovian methodology. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 143-152.

Lachnit, H., & Kimmel, H. D. (1993). Positive and negative patterning in human classical skin conductance response conditioning. *Animal learning and Behavior, 21*, 314-326.

Laplace, P. (1812). *Théorie analytique des probabilités*. Paris: Courcier.

Lee, M. D. (2010). Emerged and structured cognition in Bayesian models: Comment on Griffiths et al. and McClelland et al. *Trends in Cognitive Science, 14*(8), 345-346.

Lewens, T. (2007). *Risk. Philosophical perspectives*. New York: Routledge.

Li, M., & Chapman, G. B. (2009). "100% of anything looks good": The appeal of one hundred percent. *Psychonomic Bulletin & Review, 16*(1), 156-162.

Liljeholm, M., & Cheng, P. W. (2007). When is a cause the "same"? Coherent generalization across contexts. *Psychological Science, 18(11)*, 1014-1021.

Liljeholm, M., & Cheng, P. W. (2009). The influence of virtual sample sizes on confidence and causal-strength judgments. *Journal of Experimental Psychology: Learning, memory and cognition, 35*, 157-172.

Lipkus, I.M., & Peters, E. (2009). Understanding the role of numeracy in health: Proposed theoretical framework and practical insights. *Heatlh Education and Behavior, 36*, 1065-1081.

Lipkus, I.M., Samsa, G., & Rimer, B. K. (2001). General performance on a numeracy scale among hughly educated samples. *Medical decision making, 21(1)*, 37-44.

Lipp, O. V., Cox, D., & Siddle, D. A. T. (2001). Blink startle modulation during anticipation of pleasant and unpleasant stimuli. *Journal of Psychophysiology, 15,* 155-162.

Lober, K., & Lachnit, H. (2002). Configural learning in human Pavlovian conditioning: Acquisition of a biconditional discrimination. *Biological Psychology, 59,* 163-168.

Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review, 107*(1), 195-212.

López, F. J., Almaraz, J., Fernández, P., & Shanks, D. R. (1999). Adquisición progresiva de conocimiento sobre relaciones predictivas: Curvas de aprendizaje en juicios de contingencia. *Psicothema, 11,* 337-349.

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*, 672-694.

Lovibond, P. F., Been, S. -L., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory & Cognition, 31*(1), 133-142.

Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological Review, 115,* 955-984.

Lu, H., Rojas, R. R., Beckers, T., & Yuille, A. (2008). Sequential causal learning in humans and rats. In *Proceedings of the 29th Annual Conference of the Cognitive Science Society.*

Lubow, R. E. (1973). Latent inhibition. *Psychological Bulletin, 79*, 398-407.

Lubow, R. E.; Moore, A. U. (1959). Latent inhibition: The effect of non-reinforced preexposure to the conditional stimulus. *Journal of Comparative and Physiological Psychology, 52*, 415-419.

Luhmann, C. C., & Ahn, W. (2005). The meaning and computation of Causal Power: Comment on Cheng (1997) and Novick and Cheng (2004). *Psychological Review, 112*(3), 685-693.

Macchi, L. (2000). Partitive formulation of information in probabilistic problems: Beyond heuristics and frequency formats. *Organizational Behavior and Human Decision Processes, 82*, 217-236.

Macdonald-Ross, M. (1977). How numbers are shown: a review of research on the presentation of quantitative data in texts. *Audio-visual communication review, 25,* 359-407.

Mackintosh, N. J. (1974). *The psychology of animal learning*. New York: Academic Press.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychology Review, 82*, 276-298.

Martin, I., & Levey, A. B. (1991). Blocking observed in human eyelid conditioning. *Quarterly Journal of Experimental Psychology, 43,* 233-256.

Matute, H. (1995). Human reactions to uncontrollable outcomes: Further evidence for superstitions rather than helplessness. *Quarterly Journal of Experimental Psychology, 48B,* 142-157.

Matute, H. (1996). Illusion of control: Detecting response-outcome independence in analytic but not in naturalistic conditions. *Psychological Science, 7,* 289-293.

Matute, H., Arcediano, F., & Miller, R. R. (1996). Test question modulates cue competition between causes and between effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 22*, 182-196.

Matute, H., Vegas, S., & De Marez, P. –J. (2002). Flexible use of recent information in causal and predictive judgments. *Journal of Experimental Psychology: Learning, Memory and Cognition, 28,* 714-725.

Matute, H., Yarritu, I., & Vadillo, M. A. (2011). Illusions of causality at the heart of pseudoscience. *British Journal of Psychology, 102,* 392-405.

Mackintosh, N. J. (1975). A theory of attention: variations in the associability of stimuli with reinforcement. *Psychological Review, 82*, 276-298.

Mehta, R., & Williams, D. A. (2002). Elemental and configural processing of novel cues in deterministic and probabilistic taks. *Learning and Motivation, 33,* 456-484.

Melchers, K. G., Shanks, D. R., & Lachnit, H. (2008). Stimulus coding in human associative learning: Flexible representations of parts and wholes. *Behavioural Processes, 77,* 413-427.

Mellers, B. A., & McGraw, P. (1999). How to improve Bayesian reasoning: Comment on Gigerenzer and Hoffrage (1995). *Psychological Review, 106*, 417-424.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences, 32*(2), 183-198.

Mitchell, C. J., & Lovibond, P. F. (2002). Backwardand forward blocking in human electrodermal conditioning: Blocking requires an assumption of outcome additivity. *The Quarterly Journal of Experimental Psychology, 55,* 311-329.

Mitchell, C. J., Lovibond, P. F., & Condoleon, M. (2005). Evidence for deductive reasoning in blocking of causal judgments. *Learning and Motivation, 36*, 77-87.

Myers, K. M., Vogel, E. H., Shin, J., & Wagner, A. R. (2001). A comparison of the Rescorla-Wagner and Pearce models in a negative patterning ans a summation problem. *Animal Learning & Behavior, 29*(1), 36-45.

Neace, W. P., Michaud, S., Bolling, L., Deer, K., & Zecevic, L. (2008). Frequency formats, probability formats, or problem structure? A test of the nested-sets hypothesis in an extensional reasoning task. *Judgment and Decision Making, 3(2)*, 140-152.

Nelson, W., Reyna, V. F., Fagerlin, A., Lipkus, I., & Peters, E. (2008). Clinical implications of numeracy: Theory and practice. *Annals of Behavioral Medicine, 35*, 261-274.

Nickerson, R. S. (2008). *Cognition and chance: The psychology of probabilistic reasoning*. New Jersey: Taylor & Francis.

Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review, 111*, 455-485.

Ore, O. (1953). In *Cardano, the gambling scholar*. Publisher: Princeton University Press.

Pacioli, L. (1494). In *Summa de Arithmetica, geometria, proportioni et proportionalita*. Venecia.

Paulos, J. A. (1988). *Innumeracy: Mathematical illiteracy and its consequences*. New York: Hill and Wang.

Pavlov, I. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*. London: Oxford University Press.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review, 94*, 61-73.

Pearce, J. M. (1994a). Similarity and discrimination: A selective review and a connectionism model. *Psychological Review, 101*(4), 587-607.

Pearce, J. M. (1994b). Discrimination and categorization. In N. J. Mackintosh (Ed.), *Animal learning and cognition*. San Diego, CA: Academic Press.

Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology, 52*, 111-139.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review, 87*, 532-552.

Pearce, J. M., George, D. N., & Aydin, A. (2002). Summation: Further assessment of a configural theory. *Quarterly Journal of Experimental Psychology, 55*, 61-73.

Pearl, J. (2000). In *Causality: Models, reasoning, and inference*. London: Cambridge University Press.

Perales, J. C., & Shanks, D. R. (2003). Normative and descriptive accounts of the influence of power and contingency on causal judgement. *The Quarterly Journal of Experimental Psychology, 56A*(6), 977-1007.

Perales, J. C., & Shanks, D. R. (2007). Models of covariation-based causal judgment: A review and synthesis. *Psychonomic Bulletin & Review, 14*(4), 577-596.

Perales, J. C., & Shanks, D. R. (2008). Driven by power? Probe question and presentation format effects on causal judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*(6), 1482-1494.

Peters, E., Hart, S., & Fraenkel, L. (2011). Informing patients: The influence of numeracy, framing, and format of side-effect information on risk perceptions. *Medical Decision Making, 31,* 432-436.

Peters, E., Västfjäll, D., Slovic, P., Mertz, C. K., Mazzocco, K., & Dickert, S. (2006). Numeracy and decision making. *Psychological Science, 17,* 407-413.

Peterson, C. (1980). Recognition of non contingency. *Journal of Personality and Social Psychology, 38,* 727-734.

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin, 68*, 29-46.

Platzer, C., & Bröder, A. (2012). Most people do not ignore salient invalid cues in memory-based decisions. *Psychonomic Bulletin & Review, 19*, 654-661.

Real, L. A. (1991). Animal choice behavior and the evolution of cognitive architecture. *Science, 253*, 980-986.

Rescorla, R. A. (1970). Reduction in the effectiveness of reinforcement after prior excitatory coditioning. *Learning and Motivation, 1*, 372-381.

Rescorla, R. A., & Coldwell, S. E. (1995). Summation in Autoshaping. *Animal learning and behavior, 23,* 314-326.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement. In A. H. Black, & W. F. Prokasky, *Classical conditioning II: Current theory and research* (pp. 64-99). New York: Appleton-Century-Crofts.

Reyna, V. F., & Brainerd, C. J. (2007). The importance of mathematics in health and human judgment: Numeracy, risk communication, and medical decision making. *Learning and Individual Differences, 17,* 147-159.

Reyna, V. F., & Brainerd, C. J. (2008). Numeracy, ratio bias, and denominator neglect in judgments of risk and probability. *Learning and Individual Differences, 18,* 89-107.

Reyna, V. F., Lloyd, F. J., & Whalen, P. (2001). Genetic testing and medical decision making. *Archives for Internal Medicine, 161,* 2406-2408.

Reyna, V. F., & Mills, B. A. (2007). Converging evidence supports fuzzy-trace theory's nested sets hypothesis (but not the frequency hypothesis). *Behavioral and brain sciences, 30,* 278-280.

Reyna, V. F., Nelson, W. L., Han, P. K., & Dieckmann, N. F. (2009). How numeracy influences risk comprehension and medical decision making. *Psychological Bulletin, 135,* 943-973.

Shanks, D. R. (1985). Continuous monitoring of human contingency judgment across trials. *Memory and Cognition, 13*, 158-167.

Shanks, D. R. (1987). Acquisition functions in causality judgment. *Learning and Motivation, 17*, 27-34.

Shanks, D. R. (1989). Selectional processes in causality judgment. *Memory & Cognition, 17*, 27-34.

Shanks, D. R. (1991). A connectionist account of base-rates biases in categorization. *Connection Science, 3*(2), 143-162.

Shanks, D. R. (1991). Categorization by a connectionist network. *Journal of Experimental Psychology: Learning, Memory and Cognition, 17*, 433-443.

Shanks, D. R. (1995). Is human learning rational? *The Quarterly Journal of Experimental Psychology, 48A*(2), 257-279.

Shanks, D. R. (2002). Test of the Power PC theory of causal induction with negative contingencies. *Experimental Psychology, 49*(2), 81-88.

Shanks, D. R. (2010). Learning: From association to cognition. *Annual Review of Psychology, 61*, 273-301.

Shanks, D. R. (s.f.). The psychology of associative learning. New York: Cambridge University Press.

Shanks, D. R., Charles, D., Darby, R. J., & Azmi, A. (1998). Configural processes in human associative learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24*(6), 1353-1378.

Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavioral Processes, 24,* 405-415.

Shanks, D. R., & Dickinson, A. (1987). Associative accounts of causality judgment. En G. H. Bower, *The Psychology of Learning and Motivation* (Vol. 21, págs. 229-261). New York: Academic Press.

Sheridan, S. L., & Pignone, M. (2002). Numeracy and the medical students' ability to interpret data. *Effective clinical practice, 5*, 35-40.

Siegler,R. S. (1987). The perils of averaging over strategies: An example from children's addition. *Journal of Experimental Psychology: General, 116,* 250-264.

Simkin, D., & Hastie, R. (1987). An information processing analysis of graph perception. *Journal of the American Statistical Association., 82,* 454-465.

Simmons, D. (1989). Hyperion Cantos. ISBN 1-56865-175-9.

Sloman, S. A. (2005). In *Causal models: How we think about the world and its alternatives*. Cambridge, MA: Oxford University Press.

Sloman, S. A., Over, D. E., Slovak, L., & Stibel, J. M. (2003). Frequency illusions and other fallacies. *Organizational behavior and human decision processes, 91*, 296-309.

Soto, F., Vogel, E., Castillo, R., & Wagner, A. (2009). Generality of the summation effect in human causal learning. *Quarterly Journal of Experimental Psychology, 62,* 877-889.

Spellman, B. A. (1997). Crediting causality. *Journal of Experimental Psychology: General, 126*(4), 323-348.

Spence, I. (1990). Visual psychophysics of simple graphical elements. *Journal of Experimental Psychology: Human perception and performance, 16,* 683-692.

Spence, I., & Lewandowsky, S. (1991). Displaying proportions and percentages. *Applied cognitive psychology, 5,* 61-77.

Spirtes, P., Glymour, C., & Scheines, R. (2001). In *Causation, Prediction, and Search*. Cambridge, MA: MIT Press.

Stigler, S. M. (1990). *The history of statistics: The measurement of uncertainty before 1900*. Cambridge, MA: The Belknap press of Harvard University Press.

Tangen, J. M., & Allan, L. G. (2003). The relative effect of cue interaction. *The Quarterly Journal of Experimental Psychology, 56B*(3)*,* 279-300.

Taynes, E. T. (2003). *Probability theory: The logic of science*. New York: Cambridge University Press.

Tenenbaum, J. B., & Griffiths, T. L. (2003). Theory-based causal induction. In S. Becker, S. Thrun, & K. Obermayer, *Advances in Neural Information Processing Systems* (pp. 35-42). Cambridge, MA: MIT Press.

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences, 10*(7), 309-318.

Tolman, E. C., & Brunswick, E. (1935). The organism and the causal texture of the environment. *Psychological review, 42*, 43-77.

Treisman, M. (1998). Combining information: Probability summation and probability averaging in detection and discrimination. *Psychological* Methods, 3(2), 252-265.

Tufte, E. R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.

Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbeing, *Progress in Social Psychology*. Hillsdale,

Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (1998). Causal induction in the presence of a perfect negative cue: Contrasting predictions from associative and statistical models. *Quarterly Journal of Experimental Psychology, 51B*, 173-191.

Vallée-Tourangeau, F., Murphy, R. A., & Baker, A. G. (2005). Contiguity and the outcome density bias in action-outcome contingency judgments. *The Quarterly Journal of Experimental Psychology, 58B*(2), 177-192.

Vallée-Tourangeau, F., Murphy, R. A., & Drew, S. (1997). Causal judgments that violate the predictions of the power PC theory of causal induction. En M. G. Shafto, & P. Langley, *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (págs. 775-780). Mahwah, NJ: Erlbaum.

Vallée-Tourangeau, F., Murphy, R. A., Drew, S., & Baker, A. G. (1998). Judging the importance of constant and variable candidate causes: A test of the

Power PC Theory. *The Quarterly Journal of Experimental Psychology, 51A*(1), 65-84.

Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: the role of nonpresentation of compound stimulus elements. *Learning and Motivation, 25,* 127-151.

Van Osselaer, S. J. J., Janiszewski, C., & Cunda, M., Jr. (2004). Stimulus generalization in two associative learning processes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30,* 626-638.

Wagner, A. R. (2003). Context-sensitive elemental theory. *Quarterly Journal of Experimental Psychology, 56B,* 7-29.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology, 76,* 171-180.

Waldmann, M. R. (1996). Knowledge-based causal induction. En D. R. Shanks, K. J. Holyoak, & D. L. Medin, *Psychology of Learning and Motivation* (págs. 47-88). San Diego, CA: Academic Press.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 53-76.

Waldmann, M. R., & Hagmayer, Y. (2001). Estimating causal strength: The role of structural knowledge and processing effort. *Cognition, 82*, 27-58.

Waldmann, M. R., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General, 121*(2), 222-236.

Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. In N. Chater, & M. Oaksford, *The probabilistic mind: Prospects for rational models of cognition* (pp. 453-484). Oxford, England: Oxford University Press.

Ward, W. C., & Jenkins, H. M. (1965). The display of information and the judgment of contingency. *Canadian Journal of Psychology, 1*9, 231-241.

Wasserman, E. A. (1974). Stimulus-reinforcer predictiveness and selective discrimination learning in pigeons. *Journal of Experimental Psychology, 103,* 284-297.

Wasserman, E. A. (1990). Detecting response-outcome associations: Toward an understanding of the causal texture of the environment. In G. H. Bower, *The psychology of learning and motivation* (Vol. 26, pp. 27-82). San Diego: Academic Press.

Wasserman, E. A.; Berglan, L. R. (1998). Backward clocking and recovery from overshadowing in human causal judgment: The role of within-compound associations. *Quarterly Journal of Experimental Psychology, 51*(2), 121-138.

Wasserman, E. A., Elek, S. M., Chatlosh, D. L., & Baker, A. G. (1993). Rating causal relations: The role of probability in judgments of response-outcome contingency. *Journal of Experimental Psychology: Learning, Memory and Cognition, 19*, 174-188.

Waters, E. A., Weinstein, N. D., Colditz, G. A., & Emmons, K. (2006). Formats for improving risk communication in medical tradeoff decisions. *Journal of Health Communication, 11*, 167-182.

White, P. A. (1989). A theory of causal processing. *British Journal of Psychology, 80,* 431-454.

White, P. A. (2005). The Power PC Theory and causal powers: Comment on Cheng (1997) and Novick and Cheng (2004). *Psychological Review, 112*(3), 675-684.

White, P. A. (2009). Not by contingency: Some arguments about the fundamentals of human causal learning. *Thinking & Reasoning, 15*(2), 129-166.

Williams, D. A., & Braker, D. S. (1999). Influence of past experience on the coding of compound stimuli. *Journal of Experimental Psychology: Animal Behavior Processes, 25*, 461-474.

Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 129-142.

Williams, D. A. (1995). Forms of inhibition in animal and human learning *Journal of Experimental Psychology: Animal Behavior Processes, 21*, 129-142.

Woloshin, S., Schwartz, L. M., Black, W. C., & Welch, H. G. (1999). Women's perceptions of breast cancer risk: How you ask matters. *Medical decision making, 19*, 221-229.

Woodward, J. (2003). In *Making things happen: A theory of causal explanation*. New York: Oxford University Press.

Wu, M., & Cheng, P. W. (1999). Why causation need not follow from statistical association. *Psychological Science, 10*(2), 92-97.

Yamagishi, K. (2003). Facilitating normative judgments of conditional probability: Frequency or nested sets? *Experimental Psychology, 50*, 97-106.

Yarritu, I., Matute, H., & Vadillo, M. A. (2014). Illusion of control: The role of personal involvement. *Experimental Psychology, 61,* 38-47.

Young, M. E. (1995). On the origin of personal causal theories. *Psychonomic Bulletin & Review, 2*(1), 83-104.

Zeaman, D., & House, B. J. (1963). The role of attention in retardate discrimination learning. In N. R. Ellis (Ed.). *Handbook of mental deficiency: Psychological theory and research*. New York: Mc Graw-Hill.

Chapter 9. Appendices

9.1. Appendix 1: Instructions for Pink Eyes

Please, answer the questions in the order they are written. Once you have answered one question, please do not come back to this question later to change your answer. We are interested in your first answers.

Try to imagine the following situation. The cosmetics industry has recently been revolutionized by the discovery of a number of natural substances that can change the color of the people's eyes. Interestingly, the color most demanded by consumers is not any of the normal ones (blue, green, brown…), but pink. Unfortunately, there are just a few known substances that produce this eye color.

For the time being, scientists have identified only five substances that may cause the pink eyes. The results of these studies are as follows:

-If you inject substance Alfa, eye color turns pink 40% of the time.

-If you inject substance Beta, eye color turns pink 40% of the time.

-If you inject substance Gamma, eye color turns pink 80% of the time.

-If you inject substance Delta, eye color turns pink 64% of the time.

-If you inject substance Omega, eye color turns pink 0% of the time.

There are several cosmetic companies that sell products to change the color of the eyes. Most of these products are based on different compounds of one or several of these substances. Imagine that you also want to change the color of your eyes and answer the following questions accordingly:

1. If there is a product that contains a compound of the Alfa and Beta substances, and another that contains a compound of the Gamma and Omega substances, which one would you choose? (Please circle your choice)

Alfa and Beta / Gamma and Omega

2. If there is a product that contains a compound of the Alfa and Beta substances, and another that contains a compound of the Delta and Omega substances, which one would you choose? (Please circle your choice)

Alfa and Beta / Delta and Omega

9.2. *Appendix 2: Instructions for physical beauty*

Ponte por un momento en la siguiente situación. La ciencia descubre una serie de proteínas que tienen la característica de impedir la percepción de la belleza física, en pro de la interior. Por convicciones y principios, decides someterte a este tratamiento proteínico, con el objetivo de comenzar a percibir la esencia de las personas que te rodean, librándote de la influencia de su apariencia física.

Por el momento, los científicos sólo han estudiado cinco proteínas que podrían provocar que dejes de percibir la belleza física. Los resultados de estos estudios son los siguientes:

- Si se inyecta la proteína Alfa, 120 de cada 200 personas dejan de percibir la belleza física.
- Si se inyecta la proteína Beta, 130 de cada 200 personas dejan de percibir la belleza física.
- Si se inyecta la proteína Gamma, 190 de cada 200 personas dejan de percibir la belleza física.
- Si se inyecta la proteína Delta, 160 de cada 200 personas dejan de percibir la belleza física.
- Si se inyecta la proteína Omega, 60 de cada 200 personas dejan de percibir la belleza física.

Existen diversas empresas farmacológicas que comercializan productos para modificar la percepción de la belleza. La mayor parte de estos productos se basan en compuestos de varias de estas proteínas.

Imagina que tú también deseas dejar de percibir la belleza física inyectándote un compuesto de estas proteínas, y responde a las siguientes preguntas:

1. Si hay una marca comercial que vende un compuesto de las proteínas Alfa y Beta, y otra marca comercial que vende un compuesto de las proteínas Gamma y Omega, ¿cuál elegirías? (Marca tu respuesta con un círculo.)

Alfa y Beta Gamma y Omega

2. Y si tuvieras que elegir entre una marca con el compuesto Alfa y Beta y otra con el compuesto Delta y Omega, ¿cuál elegirías? (Marca tu respuesta con un círculo.)

Alfa y Beta Delta y Omega

9.3. Appendix 3: Instructions for Coins

Imagine you are playing a coin tossing game. You win if you get at least one head. There are three different coins. Each coin is biased, so the chance you get a head will not be 50%. The real chances are:

- Tossing COIN A produces heads 40% of the time.
- Tossing COIN B produces heads 80% of the time.
- Tossing COIN C produces heads 64% of the time

Answer the following questions taking into account that your goal is to maximize your chances of winning:

1. Which option would you prefer if given the option of tossing COIN A TWICE (you win if you get a head on either toss or on both tosses) or the option of tossing COIN B ONCE (you win if you get a head)? Please circle your preferred option.

   TOSS COIN A TWICE TOSS COIN B ONCE

2. If you are given the option of tossing COIN A TWICE (again you win if you get at least one head) or tossing COIN C ONCE (you win if you get a

head), which option would you prefer? Please circle your preferred option.

TOSS COIN A TWICE   TOSS COIN C ONCE

What is the probability of getting at least one head if you toss COIN A twice?

____

9.4. Appendix 4: Instructions for Hyperion with graphical banners

Imagina que eres un científico encargado de terraformar un planeta para que el resto de la humanidad pueda vivir en un nuevo mundo cuando la Tierra deje de ser habitable. El planeta en el que estás destinado, Hyperión, tiene unas características muy peculiares que hacen que la supervivencia humana sea muy difícil hasta que se logre dominar el terreno. Dispones de la información que te suministran los técnicos sobre los peligros que contiene el planeta y la tasa de supervivencia de los primeros pobladores.

Los técnicos han concluido que los peligros más habituales en Hyperión son:

1- Ingerir liquen de la superficie del planeta, lo cual puede provocar muerte por intoxicación.

2- Acercarse a océanos de metano líquido, lo cual puede provocar muerte por abrasión.

3- Exponerse a una tormenta eléctrica, lo cual puede provocar muerte por electrocutamiento.

4- Adentrarse en cuevas subterráneas, lo cual puede provocar muerte por aplastamiento.

5- Pisar arenas movedizas de mercurio, lo cual puede provocar muerte por ahogamiento.

Por el momento sólo disponemos de los datos sobre el número de exploradores que han sobrevivido al encontrarse con cada uno de estos peligros. En la siguiente página verás grupos de exploradores que se enfrentaron a alguno de estos peligros. Se representa gráficamente cuántos de estos exploradores lograron sobrevivir:

	Representa a un explorador que sobrevivió
	Representa a un explorador que murió

Estos exploradores ingirieron liquen de la superficie del planeta:

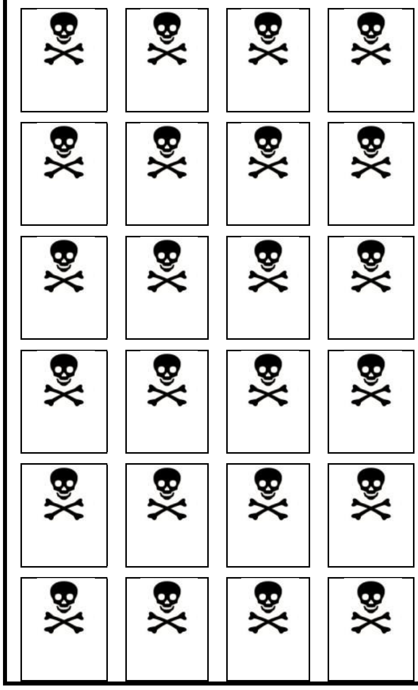

Estos exploradores se acercaron a océanos de metano líquido:

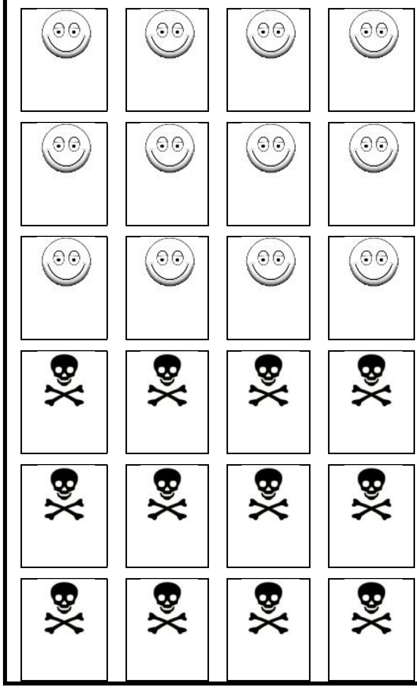

Estos exploradores se expusieron a una tormenta eléctrica:

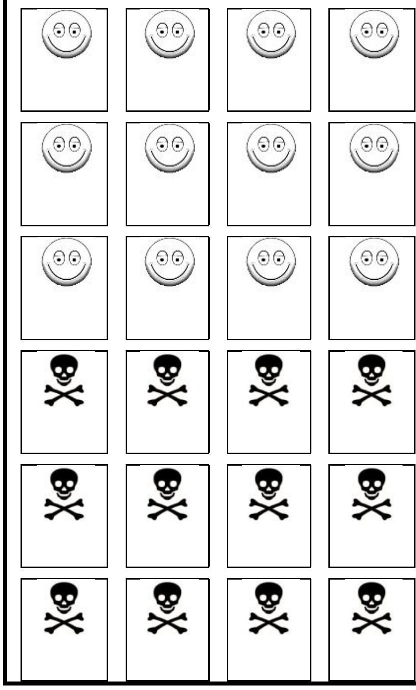

Estos exploradores se adentraron en cuevas subterráneas:

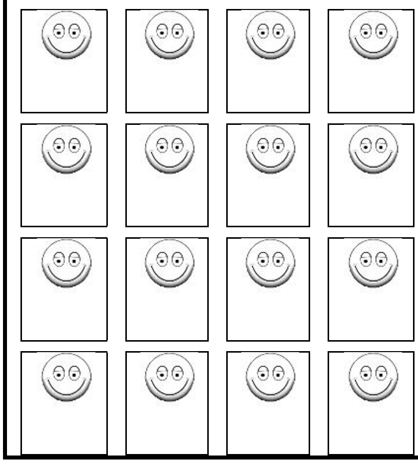

Estos exploradores pisaron arenas movedizas de mercurio:

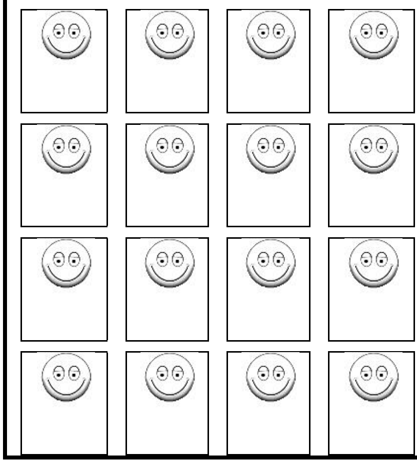

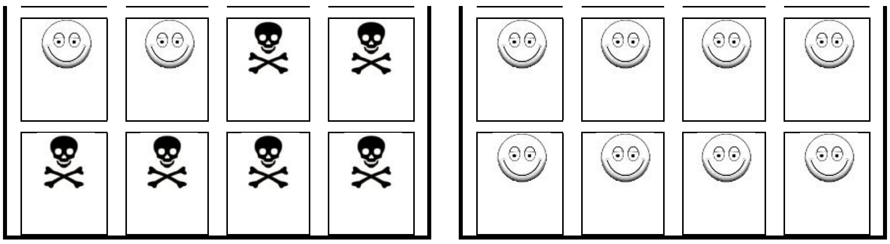

Ahora, imagina que tu misión consiste en colocar y activar la primera planta depuradora de atmósfera, que os permitirá respirar oxígeno sin necesitar máscara. Desgraciadamente, los puntos idóneos para colocar la maquinaria están situados en zonas de riesgo. Tu tarea es elegir la ruta más segura. Por favor, responde a las preguntas en orden, y no vuelvas atrás para cambiarlas. Estamos interesados en tus primeras respuestas.

Con el fin de maximizar tus posibilidades de sobrevivir, ¿preferirías tomar una ruta en la que te acercarás a océanos de metano y te expondrás a tormentas eléctricas o por el contrario una ruta en la que pisarás arenas movedizas? Redondea la respuesta.

océanos y tormentas arenas

Con el fin de maximizar tus posibilidades de sobrevivir, ¿preferirías tomar una ruta en la que te acercarás a océanos de metano y te expondrás a tormentas eléctricas o por el contrario una ruta en la que te adentrarás en cuevas subterráneas?

océanos y tormentas cuevas

Con el fin de maximizar tus posibilidades de sobrevivir, ¿preferirías tomar una ruta en la que te acercarás a océanos de metano y te expondrás a tormentas eléctricas o por el contrario una ruta en la que pisarás arenas movedizas e ingerirás liquen?

océanos y tormentas arenas y liquen

Con el fin de maximizar tus posibilidades de sobrevivir, ¿preferirías tomar una ruta en la que te acercarás a océanos de metano y te expondrás a tormentas eléctricas o por el contrario una ruta en la que te adentrarás en cuevas subterráneas e ingerirás liquen?

océanos y tormentas cuevas y liquen

De 100 nuevas personas que toman una ruta donde se acercan a océanos de metano y se exponen a tormentas eléctricas, ¿cuántas dirías que van a sobrevivir?

____

De 100 nuevas personas que toman una ruta donde pisan arenas movedizas e ingieren liquen, ¿cuántas dirías que van a sobrevivir?

____

De 100 nuevas personas que toman una ruta donde se adentran en cuevas subterráneas e ingieren liquen, ¿cuántas dirías que van a sobrevivir?

____

Edad:

Sexo:

Estudios: ctrb1