UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 37

In Search of Prototypes and Feminist Bank-Tellers: Exploring the Representativeness Heuristic

HÅKAN NILSSON

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2008

ISSN 1652-9030
ISBN 978-91-554-7102-6
urn:nbn:se:uu:diva-8476

Dissertation presented at Uppsala University to be publicly examined in VIII, universitetshuset, Wednesday, March 26, 2008 at 13:00 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract

Nilsson, H. 2008. In Search of Prototypes and Feminist Bank-Tellers: Exploring the Representativeness Heuristic. Acta Universitatis Upsaliensis. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences* 37. x, 60 pp. Uppsala. ISBN 978-91-554-7102-6.

According to the heuristics and biases approach, the representativeness heuristic (RH) is one of the heuristics available for assessing subjective probabilities (A. Tversky & D. Kahneman, 1974). A subjective probability assessed by the RH is determined by how representative the target object is of the target category. Several aspects of the RH are argued to cause systematic biases, for example: (*i*) When the RH is used, the category is represented by one single prototypical exemplar. This feature is argued to cause biases such as misperception of chance and insensitivity to sample size. (*ii*) The RH assesses the inverse rather than the conditional probability. This feature is argued to cause biases such as the conjunction fallacy and base-rate neglect.

The present thesis focuses on the cognitive aspects of the RH. Three studies were conducted. Overall, data indicated that the RH does not play a major role when subjective probabilities are assessed. Study I indicated that subjective probabilities are not typically determined by how representative the target object is of the target category. Study II indicated that the category is not represented by one single prototypical exemplar when subjective probabilities are assessed. Study III indicated that conjunction fallacies are not caused by the RH.

The results presented in Studies I-III cast serious doubts on the claim that subjective probabilities are routinely assessed using the RH. Rather, Studies I-II suggested that subjective probabilities are based on exemplar memory and Study III suggested that the conjunction fallacy is caused by people combining component probabilities in a an inappropriate way. In the General Discussion, it is suggested that people use a weighted average rule when combining component probabilities into conjunction probabilities. A simulation showing the ecological relevance of the weighted average rule is presented.

Keywords: Cognitive psychology, Subjective probability, Representativeness heuristic, PROBEX, Category learning, Cognitive modeling, Conjunction fallacy

Håkan Nilsson, Department of Psychology, Box 1225, Uppsala University, SE-75142 Uppsala, Sweden

ISSN 1652-9030
ISBN 978-91-554-7102-6
urn:nbn:se:uu:diva-8476 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-8476)

Till Filip

List of Papers

This thesis is based on the following studies, which in the following will be referred to by their Roman numerals.

I Nilsson, H., Olsson, H., & Juslin, P. (2005). The cognitive substrate of subjective probability. *Journal of Experimental Psychology: Learning, Memory and Cognition, 31*, 600-620.

II Nilsson, H., Juslin, P., & Olsson, H. (in press). Exemplars in the mist: The cognitive substrate of the representativeness heuristic. *Scandinavian Journal of Psychology*.

III Nilsson, H. (2008). *The conjunction fallacy: A phenomenon not caused by the representativeness heuristic*. Manuscript submitted for publication.

Contents

Abbreviations

CBRF	Cue-Based Relative Frequency
CRH	Combination-Rule Hypothesis
PROBEX	PRObabilities from EXemplars
RHH	Representativeness Heuristic Hypothesis
REP[ESAM]	REPresentativeness as Evidential Support Accumulation
REP[L]	REPresentativeness as relative Likelihood
REP[P]	REPresentativeness as Prototype similarity
RMSD	Root Mean Squared Deviation
WAH	Weighted Average Hypothesis

1. Introduction

Imagine that a posh friend has invited you over for a glass of wine. On your arrival he tells you that he has two bottles, one (ridiculously) expensive and one medium-priced. He wants to taste the exclusive wine but can not decide if it is worth wasting half the bottle on you. Therefore, he has decided to test your knowledge of wines. He covers your eyes and hands you a glass of wine. He tells you that if you can taste whether this wine is red or white then he will let you taste the exclusive wine. You take a sip and absorb the ensemble of flavors that spreads through your mouth. Based on the characteristics of the wine you assess the probability that the wine is red and the probability that it is white and choose the type of wine that it is most likely to be.

Everyday we face situations (more or less) similar to the one above, where we evaluate options and rank them according to how likely they are to bring us to a desired goal. The study of *subjective probability judgment* is the study of how we make such evaluations. In order to understand how probabilities are assessed there are several questions that have to be answered. First, how is the relevant knowledge represented in memory? Is the relevant knowledge represented as, for example, frequencies (e.g., Gigerenzer, Hoffrage, & Kleinbölting, 1991), abstract prototypes (e.g., Kahneman & Frederick, 2002; Koehler, White, & Grondin, 2003), associative links (e.g., Gluck & Bower, 1988a; Lagnado & Shanks, 2002), or memory traces of concrete events (e.g., Dougherty, Gettys, & Ogden, 1999; Juslin & Persson, 2002)? Second, what characterizes the cognitive process behind the subjective probability assessment? Does the process involve, for example, computation of frequency (e.g., Gigerenzer et al., 1991; Juslin, 1994), associative strength (e.g., Gluck & Bower, 1988a; Lagnado & Shanks, 2002), or similarity (e.g., Juslin & Persson, 2002; Kahneman & Frederick, 2002)? Third, what empirical patterns emerge in subjective probability judgments? For example, under which conditions do subjective probability judgments obey the rules of probability theory? These questions are of course not independent. Studies aimed at answering one question will provide partial answers to the others as well. Nonetheless, full understanding of how people assess subjective probabilities can not be accomplished without exploring all three questions.

This thesis focuses on one of the heuristics for assessing probabilities suggested by the *heuristics and biases approach* (Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982), the *representativeness* *heuristic* (Kahneman & Tversky, 1972; Kahneman & Frederick, 2002). According to the heuristics and biases approach, humans have developed a set of heuristics (cognitive shortcuts) enabling us to cope in a world of uncertainty. The scientific strategy of the heuristics and biases approach has been to study frequently occurring data patterns (biases) and to make inferences about which processes (heuristics) that have generated them (Kahneman et al., 1982). Due to the focus on biases, relatively little is known about the cognitive representations and processes involved in heuristics such as the representativeness heuristic. This is the problem addressed in this thesis. In three studies, an attempt is made to more thoroughly explore the cognitive aspects of the representativeness heuristic.

Study I and Study II focus on which cognitive representations and cognitive processes that are involved when subjective probabilities are assessed. Study II, but also Study I to some extent, uses an experimental task designed to encourage the participants to apply the representativeness heuristic. Study III explores the validity of the claim that the representativeness heuristic causes the conjunction fallacy (Tversky & Kahneman, 1983). As will be shown, the results of all three studies cast doubt on whether subjective probabilities are routinely generated by the representativeness heuristic.

2. Subjective Probability

2.1. The Evolution of Probability

The theoretical concept of *probability* emerged in the middle of the seventeenth century (Hacking 1975). The enlightenment was just about to bloom (Skirbekk & Gilje, 1995). Among the intelligentsia, the common assumption was that the events of the world could be reduced to deterministic laws, and consequently, so was the function of the human mind. The human mind (or more exactly, the mind of the educated man) was thought to be governed by the objective mechanism of rationality. The study of probability became the study of rationality (Gigerenzer et al., 1989).

Modern day scholars of probability distinguish between two meanings of the term probability, the *subjective* and the *frequentistic* interpretation (Hacking, 1995). According to the subjective view of probability, probability is a property of the human mind. According to the frequentistic view of probability, probability is a property of the environment.[1] Early scholars of probability, commonly referred to as the classicists (Gigerenzer et al., 1989), did not make this distinction. It was recognized, for example, that the probability of the occurrence of an event, such as the probability that a flipped coin shows heads, can be obtained by observing frequencies of identical events in the environment (an insight relating to the frequentistic view of probability). However, as the mind was assumed to be a "counting machine that automatically tallied the frequencies of past events and scaled the degrees of belief in their recurrence accordingly" (Gigerenzer et al., 1989, p. 9), there was no reason to believe that subjective probability, i.e. the degree of belief in the occurrence of an event assessed by a single individual, and frequentistic probability should differ. Therefore, subjective probability and frequentistic probability was believed to be the same thing (Gigerenzer et al., 1989).

The goal of the classicists was to capture the workings of the human mind. As summarized by Joseph Butler in 1736, it was commonly assumed that probability was "the very guide of life" (Salmon, 1967, p. 64). There-

[1] In latter chapters I will discuss to what extent subjective probabilities correspond to objective probabilities. In those chapters, the term objective probability refers to the probability that an event occurs given perfect knowledge about the environment and given that the algorithms of probability theory are used. Accordingly, when objective probabilities are discussed, the frequentistic interpretation of probability is used.

fore, the claim was that when the laws of probability had been captured, human behavior could be described in the same way as, for example, gravity. Given the classicists' goal of describing the human mind, the early theorems of probability theory were not determined by statistics or mathematics but by observations of human behavior (Gigerenzer et al., 1989).

In the early nineteenth century, the enlightenment was quickly fading away. The human mind was no longer assumed to be driven by definable laws. Rather, it was assumed to be driven by unanalyzable factors such as intuition and sensibility (Skirbekk & Gilje, 1995). This development had a profound effect on how the concept of probability was perceived (Gigerenzer et al., 1989). As the mind was unpredictable, the goal of probability theory could no longer be to describe the human mind. Therefore, the classicists' interpretation of probability was replaced by the frequentistic interpretation of probability.

In the frequentistic era, probability became a part of statistics (Gigerenzer et al., 1989). The universal laws determining the events of the world were no longer assumed to be deterministic, as during the enlightenment, rather they were assumed to be probabilistic. The scientific method was now to study frequencies of events in the environment and the scientific goal was to use frequencies to develop means to describe and predict. In contrast to the development during the enlightenment, the formulation of probability theory was now driven by statistics and mathematics rather than by observations of human behavior.

The next significant event in the history of probability occurred in the middle of the twentieth century when utility theory (Savage, 1954; Von Neuman & Morgenstern, 1947) was launched. At that time, economists were occupied with homo economicus, an idea suggesting that humans constantly act to maximize their own well-being. The problem was that there was no available tool for modeling personal well-being. Therefore, when utility theory, which provided such a tool, was presented, it quickly became popular in economics (Goldstein & Hogarth, 1997). Importantly, utility theory brought back subjective probability. As a result, the theory provided an opportunity to dust off the old idea that the human mind was guided by the laws of probability (which was highly compatible with the theories concerning homo economicus). Today, scholars interested in the concept of probability explore both the subjective and the frequentistic interpretation of probability (Hacking, 1975). As this thesis deals with subjective probability, the rest of this chapter focuses on how the study of subjective probability has evolved since utility theory was presented.

2.2. Modern Research on Subjective Probability

This section briefly reviews the modern literature on subjective probability. The theories on how subjective probabilities are assessed can be clustered into two groups, theories developed in research on judgment and decision making (e.g., Gigerenzer et al., 1991; Juslin, 1994; Kahneman et al., 1982) and theories developed in research on various forms of learning (e.g., Dougherty et al., 1999; Gluck & Bower, 1988a; Juslin & Persson, 2002). In research on judgment and decision making, the main focus has traditionally been on exploring to what extent subjective probability obeys the laws of probability theory. Research has, for example, focused on identifying systematic deviations between subjective probabilities and objective probabilities. As a result, the primary purpose of theories developed in this field has been to explain data patterns rather than the cognitive processes that generate subjective probabilities. In research on learning, on the other hand, the main focus has been on how knowledge about the environment is mentally represented and how these representations are used to make different types of judgments. As a result, the main purpose of theories developed in this field has not been to explain judgment phenomena, but rather to describe the cognitive mechanisms behind subjective probability judgments in general. Below the two groups are discussed separately.

2.2.1. Judgment and Decision Making

The Intuitive Statistician

Early research on judgment and decision making took a normative stance. The focus was on whether normative models, such as utility theory, could serve as good models of human judgments (Goldstein & Hogarth, 1997). The general conclusion was that humans could be described as *intuitive statisticians*. When sufficient information was provided, human judgments were well captured by normative models (for early reviews see e.g., Becker & McClintock, 1967; Peterson & Beach, 1967).

The Heuristics and Biases Approach

In the late 1960s and early 1970s, an increasing amount of evidence indicated that man was probably not an intuitive statistician, or at least, not a very good one (Goldstein & Hogarth, 1997). Though human judgments most often obeyed the rules postulated by normative models, there were deviations. Interestingly, these deviations were often found to be systematic rather than random. The discovery of systematic deviations between human judgments and norms gave raise to a new research program, the *heuristics and biases approach* (Kahneman et al., 1982; Gilovich et al., 2002; Tversky & Kahneman, 1974). The heuristics and biases approach suggests that by studying systematic deviations between human judgments and norms, underlying cognitive mechanisms will be revealed.

The heuristics and biases approach relies on three basic assumptions (Gilovich et al., 2002; Kahneman et al., 1982). (*i*) *Humans have two cognitive systems working in parallel*. *System 1* (also referred to as the intuitive system) is an uncontrolled system that rather effortlessly works with general (i.e., not task specific) heuristics. For assessing subjective probabilities, System 1 is assumed to use, for example, the representativeness heuristic (Kahneman & Tversky, 1972) or the availability heuristic (Tversky & Kahneman, 1973). *System 2* (also referred to as the analytical system) is a capacity consuming, highly controlled and rational system working with abstract rules. Though System 1 always generates a judgment, that judgment is only used when System 2 fails to generate one (Kahneman & Frederick, 2002). (*ii*) *Due to the generality of the heuristics they will produce systematic biases*. The heuristics are argued to be shared by all humans. The rationale is that humans constantly encounter situations where they are unable to apply normative rules. Therefore, general heuristics that work relatively well in a range of environments have evolved. Because the heuristics are not task specific and because the heuristics are shared by all humans, systematic biases will occur in people's judgments. The representativeness heuristic, for example, is argued to cause biases such as misperception of regression, the conjunction fallacy, and base-rate neglect (Tversky & Kahneman, 1974). (*iii*) *By studying systematic biases in empirical data, inferences about the underlying heuristics can be made*.

In the tradition of the heuristics and biases approach, heuristics have the following characteristics (Kahneman et al., 1982; Kahneman & Frederick, 2002). (*a*) *A heuristic involves attribute substitution*. Many tasks include target attributes that are not easily brought to mind. Therefore, the heuristic assesses the target attribute by mapping the value of a more easily accessible attribute, referred to as the heuristic attribute, on the target dimension. By doing so, the heuristic becomes quick and easy to use. (*b*) *Overall, a heuristic generates good judgments*. A heuristic exploits correlations between feature dimensions. To maximize performance, heuristic attributes that correlate strongly with the target attributes are used. (*c*) *A heuristic is general in the sense that it can be used in a range of different situations*.

The heuristics and biases approach has been criticized on several grounds. First, it has been claimed that human judgments are compared against inappropriate norms, and as a result, the observed "biases" can not be seen as errors (e.g. Gigerenzer, 1991, 1996; Baratgin, 2002). The heuristics and biases approach compares human judgments against norms derived from utility theory. Authors such as, for example, Gigerenzer (1991, 1996), argue that, as it is the interpretation of probability adopted by most modern day statisticians, human judgments should be compared against norms derived from the frequentistic interpretation of probability. According to the frequentistic view, the concept of probability in relation to single events has no meaning. As a consequence, because their tasks typically involve probability judgments of single events, it is questionable whether the biases observed by the heuristics and biases approach actually can be seen as errors (e.g., Gigerenzer & Goldstein, 1996; Gigerenzer & Hoffrage, 1995). Though this might be true, the importance of this objection can be questioned. Independently of whether biases are systematic errors or just systematic phenomena, the study of biases has the potential to tell us a great deal about the human cognitive system. I would argue that even if the biases are not biases, the scientific strategy of the heuristics and biases approach is still useful for students of psychology.

Second, it has been argued that the observed biases are experimental artifacts (e.g. Dhami, Hertwig, & Hoffrage, 2004; Gigerenzer, 1996, 2005; Gigerenzer et al., 1991; Juslin, 1994; Juslin, Olsson, & Björkman, 1997; Winman, 1997). When designing an experimental task it is, of course, important to preserve a link with the real world. If this link is lost, generalization from experiment to the real world becomes meaningless. Furthermore, it is important that the experimental task deals with a phenomenon that is ecologically relevant rather than "a bearded lady living at the fringe of reality" (Brunswik, 1955, p. 204). Therefore, starting with Egon Brunswik (1952, 1956), several authors have stressed the importance of *representative design* (e.g., Dhami et al., 2004; Juslin, Winman, & Olsson, 2000). In a representative design, stimuli are randomly sampled from a real world environment. The rationale for using a representative design is twofold. (*a*) As it uses a real world environment the important link between experiment and the real world is preserved. (*b*) As it is randomly generated, the sample of stimuli is unlikely to include only "bearded ladies". It has been shown that biases such as overconfidence (Lichtenstein, Fishhoff, & Philips, 1982) and the hindsight bias (Fishhoff, 1975) are diminished or even eliminated in a representative design (e.g., Gigerenzer et al., 1991; Juslin, 1994; Winman, 1997; however see e.g., Griffin & Tversky, 1992). These studies suggest that at least some of the biases acknowledged by the heuristics and biases approach may partially or wholly be experimental artifacts (read bearded ladies). However, it should be noted that other biases, such as the conjunction fallacy, are observed also in studies adopting a representative design (e.g., Nilsson & Winman, 2007).

The third theme in the critique of the heuristics and biases approach concerns the definitions of the heuristics. The main focus is on identifying biases. When a bias is identified, a heuristic that potentially could explain why it appears is generated. Unfortunately, as the formulation of heuristics has been of secondary importance. The definitions of the heuristics are often so vague that it is impossible to test whether they actually are the cause of the associated biases (e.g. Evans & Pollard, 1982; Gigerenzer, 1996; Gigerenzer & Murray, 1987; Shanteau, 1989; Wallsten, 1983). In my view, this is the major weakness of the heuristics and biases approach. However, recent attempts to strengthen the definitions of, for example, the representativeness heuristic have been made (e.g., Kahneman & Frederick, 2002). This development opens up for a more thorough exploration of the assumptions concerning cognitive mechanisms that is made by the heuristics and biases approach. Therefore, Study I and II of this thesis use Kahneman and Frederick's (2002) definition to study whether subjective probabilities are commonly generated by the representativeness heuristic.

The Fast and Frugal Approach

The 1990s witnessed the birth of a third research program concerning judgment and decision making, the *fast and frugal heuristics approach* (Gigerenzer, Todd, and the ABC-Group, 1999). The fast and frugal heuristics approach was founded on the view that human behavior can only be understood by examining the interplay between the organism and the environment (Brunswik, 1952, 1956). The fast and frugal heuristics approach stress that humans have a set of cognitive tools especially adapted to the environment. The adaptive tools are argued (*a*) to exploit regularities in the environment, (*b*) to exploit the capabilities of the human cognitive system, (*c*) to be frugal in the sense that they put low demands on the cognitive system and (*d*) to be fast to use.

Each adaptive tool is suggested to include a search rule, a stopping rule, and a decision rule (Gigerenzer et al., 1999). The search rule defines how information is gathered in the environment. The stopping rule defines when the search for new information should end. The decision rule defines how the final judgment should be made. The three rules are assumed to be constructed to maximize the speed, the frugality, and the accuracy of the adaptive tool.

Regarding subjective probabilities, the fast and frugal heuristics approach suggests that probabilities are based on so called *cue-validities* (Gigerenzer et al., 1991, 1999). The cue-validity of feature j in regard to category A, $cv_j(A)$, is given by,

$$cv_j(A) = \frac{n(feature_j \& A)}{n(feature_j)} \tag{1}$$

where $n(feature_j)$ represents the frequency of feature j in the environment and $n(feature_j \ \& \ A)$ denotes the frequency of category A members with feature j. Thus, $cv_j(A)$ indicates how good feature j is at predicting whether or not an object belongs to category A. For example, if $cv_j(A)$ is high, most objects with feature j are members of category A. Accordingly, the probability that an object with feature *j* is a member of category A is high when $cv_j(A)$ is high.

To understand how subjective probabilities are assumed to be assessed, imagine a respondent that is to judge the probability that object X belongs to category A. For simplicity, assume that object X can belong to either category A or to category B. When faced with this task, the respondent creates a probabilistic mental model. The probabilistic mental model holds information concerning how good various features are at discriminating between members of category A and members of category B. In other words, it holds information concerning the cue-validities of a set of features. The features are ranked from the feature with the highest cue-validity to the feature with the lowest cue-validity. The respondent searches for the feature with the highest cue-validity that is present in object X (search rule). The search is ended when the feature with the highest possible cue-validity is found (stopping rule). The assessed probability that object X is a member of category A is based solely on the cue-validity of the single most valid feature (decision rule). Imagine that this feature has a cue-validity of .9. If the feature indicates membership in category A, then the subjective probability that object X is a category A member will equal .9. If the feature indicates membership in category B, then the subjective probability that object X is a category A member will equal .1.

The adaptive tool for assessing subjective probabilities exploits the fact that features are not randomly distributed over categories. Some categories are more probable given certain features. Therefore, people will perform relatively well if they base judgments on features with high cue-validities (Gigerenzer & Goldstein, 1996). Furthermore, the adaptive tool for assessing subjective probabilities exploits the fact that people appear to store frequencies automatically even in cases where it is not explicitly called for (Estes, 1976; Zacker & Hasher, 2002). Accordingly, as it uses relative frequencies of categories given features, the tool for assessing subjective probabilities uses information that is often easily available. Finally, it is argued to be fast and frugal because probabilities are based on single features.

Empirically, there is some support for the idea that subjective probabilities might be based on cue-validities of single features. First, Gigerenzer et al. (1991; see also e.g., Björkman, 1994; Juslin, 1994) showed that the overconfidence bias can be accounted for by a model that uses cue-validities of single features to generate probabilities. Second, studies have shown that especially when objects increase in complexity people often adopt strategies focusing on few highly diagnostic cues (e.g., Ford, Schmitt, Schechtman, Hults, & Doherty, 1989; Skov & Sherman, 1986; Trope & Bassak, 1983). On the negative side, there are findings that question the validity of the claim that probabilities are based on cue-validities of single features. The most important of these concerns how features are ranked. To enable judgments based on the single most valid feature, features have to be a priori ranked according to their cue-validities. However, studies have shown that in rather simple environments (with only four to six relevant features), people often have trouble learning the correct cue-validities and consequently are incapable of ranking features according to their cue-validities (e.g., Newell & Shanks, 2003; Rakow, Hinvest, Jackson, & Palmer, 2004; Rakow, Newell, Fayers, & Hersby, 2005). Furthermore, as the number of relevant features increase, the task of ranking them according to their cue-validities quickly becomes so capacity-demanding that it becomes impossible for the human mind to perform (Schmitt & Martignon, 2006; however see, Garcia-Retamero, Takezawa, & Gigerenzer, 2007). Thus, while there are findings suggesting that people sometimes base judgments on just few highly diagnostic features (i.e., features with high cue-validities), it is not clear how people are able to identify them.

2.2.2. Learning

Exemplar Theory

The world includes a myriad of feature-rich objects and events. Despite this complexity, we (humans) are able to extract enough information to enable us to survive within this chaos. In cognitive psychology it is commonly assumed that this is possible because we structure the world into categories and store information concerning these categories (Solomon, Medin, & Lynch, 1999). Our mental representation of a category is referred to as a concept. The knowledge embedded in our concepts enables us to categorize new objects, to understand occurring events, to predict future events and so forth.

A central question is how concepts are formed and used (for an introduction see, Murphy, 2004). One of the most influential explanations is provided by exemplar theory (e.g., Hintzman, 1988; Medin & Shaffer, 1978; Nosofsky & Johansen, 2000). Exemplar theory assumes that people store concrete memory traces of objects encountered in the environment (so called exemplars). For example, when a new object is to be classified or when a new event is to be interpreted, exemplars similar to the new object or the new event are retrieved from memory. The information included in the sample of exemplars is then used to make a categorization or an interpretation. Exemplar models have been found to account for phenomena in fields such as attention (Logan, 2002), social cognition (Smith & Zarate, 1992), categorization (Nosofsky & Johansen, 2000), automatisation (Logan, 1988), memory (Hintzman, 1986, 1988) and language (Daelemans, 1995).

In the last decade, a numbers of authors have argued that subjective probabilities are based on exemplar memory (e.g., Dougherty et al., 1999; Juslin & Persson, 2002; Sieck & Yates, 2001). Exemplar models have been shown to account for phenomena such as conservatism (Dougherty et al., 1999), overconfidence (Sieck & Yates, 2001), subadditivity (Bearden & Wallsten, 2004) and hindsight bias (Dougherty et al., 1999) to name a few. The most notable exemplar models applied to subjective probability are MINERVA-DM (Dougherty, 2001; Dougherty et al.,1999) and PROBEX (Juslin & Persson, 2002). Though MINERVA-DM and PROBEX differ on some accounts (for a discussion see, Juslin & Persson, 2002), they both assume that when the probability that object X belongs to category A is to be judged, a set of exemplars similar to object X is retrieved. The probability is then dependent on the summed similarity between object X and those of the retrieved exemplars that correspond to category A members relative to the summed similarity between object X and all retrieved exemplars. That is, if object X is relatively similar to the exemplars corresponding to category A members, the subjective probability that object X is a member of category A will be high. But, if object X is relatively more similar to the exemplars that do not correspond to category A members, the subjective probability that object X is a member of category A will be low.

Associative Learning Theory

The last presented theory concerning assumes that subjective probabilities are based on the associative strength between outcomes and cues (e.g., Cobos, Almaraz, & García-Madruga, 2003; Gluck & Bower, 1988a; Lagnado & Shanks, 2002; Newell, Lagnado, & Shanks, 2007). The core assumption, suggested already by Pavlov (e.g., Pavlov, 1927) and enthusiastically explored during the behaviorist era (e.g., Watson & Rayner, 1920), is that humans (and animals) learn by associating cues with outcomes (e.g., Gluck & Bower, 1988b). The more often a cue is accompanied by a certain outcome, the stronger the link between them. Connectionist networks based on the assumption that people learn by associating cues with outcomes have been shown to reproduce a range of cognitive behaviors relating to, for example, learning, pattern recognition, speech production and motor control (Gluck & Bower, 1988a). Regarding subjective probability, connectionist networks have been shown to account for base-rate neglect, the disjunction fallacy, and the conjunction fallacy (e.g., Coboz et al., 2003; Lagnado & Shanks, 2002).

The theory that subjective probabilities are based on associative links might appear similar to the theory that subjective probabilities are based on cue-validities. However, there are three key differences. According to the associative view; knowledge is represented by associative links not by frequencies; the subjective probability is based on the contingency between outcome (O) and cue (C), defined as the probability of O when C is present minus the probability of O when C is absent (i.e., $p(O \mid C) - p(O \mid \neg C)$), and not on the conditional probability of outcome given cue, defined as $p(O \mid C)$[2];

[2] Note that $p(O \mid C)$ is identical to $cv_C(O)$ as defined by Equation 1.

in addition, all cues present in object X are taken into account when the subjective probability that object X belongs to category A is to be assessed. As this last theory is not included in any of the studies, it will not be discussed any further.

3. The Representativeness Heuristic

3.1. Assumptions

A probability assessed by the representativeness heuristic is determined by how representative the target object is of the target category. Or in other words, "the subjective probability of an event, or a sample, is determined by the degree to which it: (*i*) is similar in essential characteristics to its parent population; and (*ii*) reflect the salient features of the process by which it is generated" (Kahneman & Tversky, 1972, p. 430). For example, the probability that the present doctoral thesis (the target object) will be approved depends on how representative it is of previously approved doctoral theses (the target category). The assumption is that if the thesis shares many features with previously approved theses then the probability that it will be approved will be high, otherwise it will be low.

Study I of this thesis explores the validity of the assumption that probabilities are dependent on how representative the target object is of the target category. Three different interpretations of this assumption are implemented as cognitive models and compared with one model using cue-validities and one model based on exemplar memory.

The representativeness heuristic meets the heuristics and biases approach's criteria for a heuristic (presented in Chapter 2.2.1). First, it involves attribute substitution. "Prediction by representativeness involves two separate acts of substitution – the substitution of a prototypical exemplar for a category, and the substitution of the heuristic attribute of similarity for the target attribute of probability" (Kahneman & Frederick, 2002, p. 73). These substitutions are argued to create a process that puts limited demands on the cognitive system. To store one single representation per category uses limited storage capacity and to retrieve only one representation at the moment of judgment makes the representativeness heuristic quick and easy. Using similarity is argued to be effective since is a property that is routinely estimated also in situations where it is not directly called for (Tversky, 1977; Tversky & Kahneman, 1983). Therefore, to consider similarity is argued not to include any additional processing. The claim that similarity is substituted for probability is indirectly supported by the finding that there is often a high correlation between the probability judgment that an object belongs to a category and the objects similarity to the typical member of that category (r = .99 in Kahneman et al., 1982, and r = .97 in Kahneman & Tversky, 1973).[3]

Second, the representativeness heuristic is argued to generate good probability judgments. The performance of the representativeness heuristic is mainly determined by two factors. (*a*) The representativeness heuristic performs relatively well in environments with homogenous categories and relatively poor in environments with heterogeneous categories. That is, the representativeness heuristic can only perform well in environments where similarity to the category prototype actually is a good indicator of category membership. (*b*) The representativeness heuristic performs well in environments where the various categories are more or less equally common. However, in environments where one or several categories are either highly common or highly uncommon the performance of the representativeness heuristic deteriorates. For example, even if this thesis shares very few features with the prototypical approved thesis, the objective probability that it is approved will be high if the approval rate is extremely high. This base-rate factor is neglected by the representativeness heuristic. At least concerning the first aspect, the representativeness heuristic appears to be reasonably well adapted to the environment. Members of real-world categories tend to share features (Mervis & Rosch, 1981). Since the prototype incorporates the features that are most common in a category, it holds a great deal of important information. As category members often share features, similarity, and therefore also representativeness, will at least to some extent correlate with probability. Thus, because it exploits regularities in the environment and because it uses a heuristic attribute that is commonly available, it is reasonable to believe that the attribute substitution of the representativeness heuristic would make the heuristic at least moderately efficient (to the author's knowledge there is no study directly testing the ecological validity of the representativeness heuristic).

Third, the representativeness heuristic can be used in a range of situations (Tversky & Kahneman, 1974). It can be used to assess probabilities in any situation that involves one target object and one target category. For example, the representativeness heuristic can be used to assess the probability that a person is a member of a particular category, the probability that a sample was drawn from a particular population, or the probability that an event will cause a certain outcome.

The present thesis focuses on two aspects of the representativeness heuristic argued to cause biases. The first is the attribute substitution. As each category is represented by a "prototypical exemplar" (Kahneman & Freder-

[3] Please note that this correlation is predicted also if subjective probabilities are based on exemplar memory. Hence, though this correlation does indicate that similarity plays an important role in subjective probability judgment, it is debatable whether the correlation actually indicates use of the representativeness heuristic.

ick, 2002, p. 73), information concerning atypical members, within-category variation, and within-category correlation are overlooked. That is, the prototypical exemplar does not incorporate information such as, for example, that the bird category includes members that can not fly, that blue cars can differ in color, or that tall people tend to have relatively large feet while short people tend to have relatively small feet. *Misperception of chance*, *insensitivity to sample size,* and *misconception of regression* are some of the biases ascribed to the attribute substitution (Tversky & Kahneman, 1974). To exemplify one of these biases; the most classic demonstration of the misperception of chance involves sequences of coin tosses. When asked whether it is more likely to obtain H-H-H-T-T-T (H = heads and T = tails) or H-T-H-T-T-H, people tend to argue that the latter is more likely. It is argued that people fail to appreciate that both sequences are equally probable because the H-T-H-T-T-H sequence is more similar to the prototypical random sequence (Tversky & Kahneman, 1974).

Study II of this thesis deals with representativeness effects, which is a judgment pattern that can be ascribed to the attribute substitution. For a demonstration of representativeness effects, consider the following. Object X shares a high number of features with the prototype of category B. As it happens to be the case that objects identical to X more often belong to category A, the normative probability that it belongs to category A is higher than the normative probability that it belongs to category B. However, the representativeness heuristic predicts that p(category A) < p(category B) since object X is more representative of category B. A representativeness effect occurs when an object is assigned to a category to which it never or rarely belongs just because it is representative of that category.

The second addressed aspect of the representativeness heuristic that is assumed to generate biases is that it actually estimates the wrong probability. Rather than reflecting the conditional probability (i.e., p(*target category* given *target object*)), subjective probabilities assessed by the representativeness heuristic reflects the inverse probability (i.e., p(*target object* given *target category*)). In less abstract terms, when the representativeness heuristic is used to assess the probability that the present thesis will be approved, the subjective probability reflects the probability that an approved thesis has the features of the present thesis (i.e., p(that a thesis has the features of the present thesis given the category of approved theses)) rather than the probability that a thesis with the features of the present thesis is approved (i.e., p(the present thesis is approved given the present thesis)). That subjective probabilities assessed by the representativeness heuristic reflect the inverse rather than the conditional probability cause biases such as *base-rate neglect* (Kahneman & Tversky, 1972, 1973) and the *conjunction fallacy* (Tversky & Kahnemman, 1983).

To exemplify base-rate neglect, the probability that this thesis will be approved is a function of the merits of this thesis in relation to other theses and of the proportion of theses that are approved. The last factor is referred to as the base-rate of approved theses. Even if this thesis is extremely bad, the objective probability that it will be approved is high if the base-rate of approved theses is very high and vice versa. It has frequently been found that people often either put to little weight on or simply ignore base-rates (e.g., Gavanski & Hui, 1992; Kahneman & Tversky, 1972, 1973). This phenomenon, known as base-rate neglect, is argued to be due to people estimating the inverse probability. When assessing the probability that an approved thesis has the features of this thesis, the base-rate of approved theses is ignored. Study III focuses on another bias attributed to the representativeness heuristic estimating the inverse probability, the conjunction fallacy (further discussed in Chapter 4).

In sum, when the representativeness heuristic is applied the probability that object X belongs to category A depends on how representative object X is of category A. The representativeness heuristic causes a set of biases, some attributed to the aspect that the category is represented by a prototype and some attributed to the representativeness heuristic assessing the inverse probability. Study I focuses on the assumption that subjective probabilities are based on representativeness. Study II and III focuses on whether the representativeness heuristic is a good account of the biases it is argued to cause. Study II focuses on a phenomenon that can be ascribed to the category being represented by a prototype and Study III is concerned with a bias attributed to the representativeness heuristic estimating the inverse probability.

3.2. Critique

As discussed above, the heuristics and biases approach has been criticized for presenting heuristics that are too vaguely defined (e.g. Evans & Pollard, 1982; Gigerenzer, 1996; Gigerenzer & Murray, 1987; Shanteau, 1989; Wallsten, 1983). Up until recently, the representativeness heuristics was simply defined as a heuristic that assesses subjective probabilities by estimating how representative the target object is of the target category (Kahneman & Tversky, 1972). Exactly how the category is represented or how representativeness is derived was never clearly defined. However, in their 2002 paper, Kahneman and Frederick made an attempt to strengthen the definition of the representativeness heuristic by suggesting both how the category is represented and how representativeness is derived. As stated above, they claimed that the category was represented by a prototypical exemplar and that representativeness was based on the similarity between the target object and the prototypical exemplar representing the target category. Unfortunately, though this certainly increased the testability of the representativeness heuristic, it did not completely solve the problem. The representativeness heuristic is only assumed to be involved when System 1 is used.

Because it never is properly defined when judgments are assessed by System 1 and when they are assessed by System 2, this makes it difficult to test to what extent the representativeness heuristic is used to assess subjective probabilities. The problem is that, as the same task might trigger System 1 for respondent A but System 2 for respondent B, a poor fit of a model implementing the ideas of Kahneman and Frederick (2002) can always be explained by arguing that System 1 was not involved (or that System 1 relied on one of the other heuristics available for assessing subjective probabilities). Thus, given the assumption that the representativeness heuristic is a property of System 1 and given that it never is defined exactly when System 1 is used, the hypothesis that subjective probabilities routinely are assessed using the representativeness heuristic becomes impossible to falsify. One way of limiting this problem is to explore to what extent biases associated with the representativeness heuristic occur in situations where the representativeness heuristics can not be applied (as in Study III of this thesis). Though such studies can not show that the representativeness heuristic never is used, they can show to what extent the associated biases are dependent on the representativeness heuristic. Interestingly, studies using this strategy have indicated that the representativeness heuristic is a rather poor account of both the conjunction fallacy (e.g., Fisk, 2002; Gavanski & Roskos-Ewoldsen, 1991) and base-rate neglect (e.g., Gavanski & Hui, 1992).

4. The Conjunction Fallacy

4.1. The Conjunction Fallacy and the Standard Problem

> Linda is 31 years old, single outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations. Which of the following statements is most probable?
>
> A) **Linda is a bank-teller**
> B) **Linda is a bank-teller** that is active in the feminist movement

Because A is included in B (as highlighted above), the normative response is that A is more probable than B. Obviously, for Linda to be a bank-teller that is active in the feminist movement, she has to be a bank-teller. Despite this, when faced with this mind teaser the majority tend to argue that B is more probable than A (e.g., Beyth-Marom, 1981; Hertwig & Chase, 1998; Tversky & Kahneman, 1983; Wedell & Moro, in press). This error in reasoning is known as the *conjunction fallacy*. More precisely, a conjunction fallacy occurs when it is argued that a conjunction of two events is more probable than either of the component events incorporated in the conjunction.

As mentioned in Chapter 3.1, the conjunction fallacy has been attributed to the representativeness heuristic estimating the inverse probability. The assumption is that if the target object (e.g., Linda) is more representative of the conjunction (e.g., feminist bank-tellers) than of either one of the components (i.e., bank-tellers or feminists) the conjunction fallacy will occur. As Study III is devoted to exploring whether it is likely that the representativeness heuristic is the cause of the conjunction fallacy, I will very briefly go through the important themes of the literature on the conjunction fallacy.

The mind teaser above demonstrates the predominant experimental paradigm for exploring the conjunction fallacy, in this thesis referred to as the *standard problem* (initially introduced by Tversky & Kahneman, 1983). The standard problem includes a personality sketch presenting a particular person and a set of statements about that person. The personality sketch is always designed to be typical of one category of people (e.g., feminists) and atypical of another category of people (e.g., bank-tellers). The statements always include one statement claiming that the person is a member of the category of which he or she is atypical and one statement claiming that he or she is a member of the category corresponding to the conjunction between the category of which he or she is typical and the category of which he or she is atypical. The task is either to estimate the probability that each statement is true or to rank the statements from the most probable to the least probable.

4.2. Methodology Debate

The literature concerning the conjunction fallacy is somewhat puzzling. Though the literature is extensive, it holds relatively few studies actually exploring the causes of the conjunction fallacy. Rather, the majority of studies are either critiques against or defenses of the standard problem. There are two main themes in the critique of the standard problem. First, authors stressing the frequentistic view of probability have argued that since the standard problem deals with probability of single events the laws of probability are not applicable. And accordingly, conjunction fallacies can not be committed (e.g., Gigerenzer, 1996, 2005). Second, it has been argued that the conjunction fallacy occurs because the respondents misinterpret the task. It has, for example, been shown that if possible misinterpretations of the wording of the standard problem are accounted for, the rate of conjunction fallacies drops (e.g., Dulany & Hilton, 1991; Fiedler, 1988; Hertwig & Gigerenzer, 1999).

In defense of the standard problem, a range of claims can be made. First, it is clear that people share a common understanding of what a probability of a single event means. For example, after observing a set of colored balls being sampled from an urn, peoples' judgments of *p*(next sampled ball is blue) tend more or less to equal the proportion of blue balls among the previously sampled balls (Peterson & Beach, 1967). I would argue that independently of whether it is truly possible to violate the laws of probability in the standard problem, it is definitely possible to violate the consensus on what a subjective probability means. Second, the standard problem captures a data pattern that has been frequently observed also using other experimental paradigms, namely that subjective probabilities of single events tend to be well calibrated while subjective probabilities of compound events tend to be overestimated (e.g., Nilsson & Winman, 2007). For example, in studies where participants are asked to bet on gambles, it has been shown that they tend to have a reasonably realistic view of the probability of winning a single gamble but grossly overestimate the probability of winning a compound gamble (e.g., Cohen & Hansel, 1957; Zizzo, 2003). Hence, though it might have been used too often, the standard problem does provide a way to examine a phenomenon that stretches far beyond the standard problem. Third, though it is clear that some conjunction fallacies are caused by people misinterpreting the task, high levels of conjunction fallacies prevail also when this is checked for (e.g., Bonini, Tentori, & Osherson, 2004; Tentori, Bonini, & Osherson, 2004; Wedell & Moro, in press).

4.3. Empirical Evidence

The most striking aspect of the conjunction fallacy is probably its robustness. Though many have set their mind to it, no one has yet been able to eliminate the conjunction fallacy (Stolarz-Fantino, Fantino, Zizzo, & Wen, 2003). The conjunction fallacy is for example observed among children and adults (Davidson, 1995), among sophisticated statisticians and people with limited knowledge about the laws of probability (Hertwig & Chase, 1998) and among experts and laypeople (in relation to the framing of the task; Adam & Reyna, 2005). The conjunction fallacy prevails when the respondents have monetary incitements for avoiding it (Zizzo, Stolarz-Fantino, Wen, & Fantino, 2000), when the participants have received a brief lecture on the laws of probability (Crandall & Greenfield, 1986), when the transparency of the task is maximized (Tversky & Kahneman, 1983), and so on. However, there are factors that have been found to affect the rate of the conjunction fallacy, such as intelligence (Stanovich & West, 2002), statistical sophistication (Hertwig & Chase, 1998), response format (Hertwig & Chase, 1998), likelihood of the components (Wells, 1985), or whether or not the task is framed in terms of frequencies (Hertwig & Gigerenzer, 1999; however see Wedell & Moro, in press).

4.4. Potential Explanations

There is no doubt that if given the standard problem many will commit the conjunction fallacy. This has been shown in numerous studies. Less effort has been put into understanding why people so frequently violate the conjunction rule when asked to solve the standard problem. That is, why people argue that it is more likely that Linda is a feminist bank-teller than that she is a bank-teller. Though this has been explored in relatively few studies (for exceptions see e.g., Fantino, Kulik, Stolarz-Fantino, & Wright, 1997; Gavanski & Roskos-Ewoldsen, 1991; Hertwig & Chase, 1998), some hypotheses concerning the cause of the conjunction fallacy have been presented. Study III of the present thesis contrasts the two most prominent explanations. The first and most wide spread explanation is that the conjunction fallacy is caused by the representativeness heuristic (Tversky & Kahneman, 1983). This explanation is further on referred to as the *RHH* (representativeness heuristic hypothesis). One can not open a text book on psychology or cognitive psychology that presents the conjunction fallacy which does not suggest the representativeness heuristic as the sole cause (interestingly, you are just as unlikely to find a textbook presenting the representativeness heuristic which does not exemplify it with the conjunction fallacy and the standard problem). Despite this, the few studies that have tested the claim that the conjunction fallacy is caused by the representativeness heuristic have failed to find support for the RHH (e.g., Gavanski & Roskos-Ewoldsen, 1991; Fisk, 2002; Fisk & Pidgeon, 1996).

The second common explanation is that people combine component probabilities into conjunction probabilities in a way that generates violations of the conjunction rule (Fantino et al., 1997; Gavanski & Roskos-Ewoldsen, 1991; Nilsson & Winman, 2007; Shanteau, 1974; Zizzo, 2003). This explanation is further on referred to as the CRH (combination rule hypothesis). Normatively, if p(A) and p(B) are independent, p(A&B) is derived by multiplying p(A) with p(B). However, it has been found that people often use other rules for combining p(A) and p(B) into p(A&B), such as, for example, assessing p(A&B) by taking the average between p(A) and p(B) (Shanteu, 1974). Some alternative combination rules produce systematic violations of the conjunction rule, and accordingly, when used they produce conjunction fallacies. In Linda terms, the CRH suggests that it is the way the probability that Linda is a bank-teller and the probability that she is a feminist are combined that causes the conjunction fallacy.

Study III of this thesis focuses on an important difference between the RHH and the CRH. To highlight this difference, let us return to the Linda example presented above. The RHH suggests that the conjunction fallacy occurs because Linda is judged to be more representative of the feminist bank-teller category as compared to the bank-teller category. If a respondent has no understanding of what a feminist bank-teller is like[4], he or she will not be able to judge how representative Linda is of the feminist bank-teller category. Accordingly, an understanding of what a feminist bank-teller is like is a precondition for the conjunction fallacy to occur. The CRH, on the other hand, predicts conjunction fallacies only when an understanding of what a feminist bank-teller is like is absent. The CRH suggests that the conjunction fallacy is caused when the probability that Linda is a bank-teller is combined with the probability that Linda is a feminist. If a respondent has an understanding of what a feminist bank-teller is like, he or she should be able to assess the probability that Linda is a feminist bank-teller directly (i.e., without combining the probability that she is a feminist with the probability

[4] Regarding the RHH, the term "an understanding" refers to the ability to create a prototypical exemplar representing the target category. The prototypical feminist bank-teller can be created in two ways, either directly from encounters with feminist bank-tellers or indirectly from encounters with feminists and encounters with bank-tellers. In the latter case, information about what characterize a feminist on the one hand and what characterize a bank-teller on the other is combined into a mental image of the prototypical feminist bank-teller. Hence, an understanding of what a feminist bank-teller is like is possible even for people that never have encountered any feminist bank-tellers.

that she is a bank-teller). Because no probability combination is involved, the conjunction fallacy should not occur when an understanding of what a feminist bank-teller is like is available. Hence, while the RHH implies that an understanding of what a feminist bank-teller is like is a precondition for the conjunction fallacy to occur, the CRH implies that the conjunction fallacy should only occur when such an understanding is absent.

5. Empirical Studies

5.1. Study I: The Cognitive Substrate of the Representativeness Heuristic

The main theories on how subjective probabilities are assessed make contradictory assumptions concerning the cognitive possesses and representations involved. That is, they disagree on what constitutes the cognitive substrate of subjective probability. The problem motivating Study I was that these theories are never contrasted against each other in experiments with a particular focus on cognitive mechanisms. Therefore, Study I used cognitive modeling, an analytical tool commonly used in studies on, for example, category learning (e.g., Medin & Shaffer, 1978; Nosofsky & Johansen, 2000; Olsson, Wennerholm, & Lyxén, 2004; Smith & Minda, 2000) and multiple-cue judgment (e.g., Helverson & Rieskamp, in press; Juslin, Olsson, & Olsson, 2003; Juslin, Karlsson, & Olsson, 2008), to explore the cognitive mechanisms behind subjective probability judgment. The advantage of cognitive modeling is that it is a well developed tool that allows for detailed comparisons between cognitive models (for a discussion on cognitive modeling see, Myung & Pitt, 1997). The strategy of Study I was to define the three leading theories of the cognitive substrate of subjective probability, the representativeness heuristic theory (e.g., Gilovish et al., 2002; Kahneman & Tversky, 1972; Kahneman et al., 1982; Koehler et al., 2003), the cue-validity theory (e.g., Björkman, 1994; Gigerenzer et al., 1991; Juslin, 1994), and the exemplar memory theory (e.g., Dougherty et al., 1999; Juslin & Persson, 2002; Sieck & Yates, 2001), as cognitive models and contrast them in category learning experiments.

5.1.1. Three Theories, Five Models

Study I compares five models. The first three models correspond to different implementations of the representativeness heuristic: *representativeness as prototype similarity* (referred to as *REP[P]*), *representativeness as evidential support accumulation* (referred to as *REP[ESAM]*; Koehler et al., 2003), and *representativeness as relative likelihood* (referred to as *REP[L]*). The fourth model derives probabilities from cue-validities (referred to *as cue-based relative frequency* or *CBRF*) and the fifth model bases probabilities on exemplar memory (referred to as *PROBEX* or *probabilities from exemplars*;

Juslin & Persson, 2002). The five models will be briefly presented below.[5] The presentation will focus on the following example. Imagine that the task is to assess the probability that object X belongs to category A rather than to category B (i.e., the task is to assess p(A)). Furthermore, imagine that object X has J features and that category A and category B are mutually exclusive.

REP[P] is an implementation of the claim that "representativeness involves two separate acts of substitutions – the substitution of a prototypical exemplar for the category, and the substitution of the heuristic attribute of similarity for the target attribute of probability" (Kahneman & Frederick, 2002, p. 73). REP[P] links the representativeness heuristic to one of the classic theories of categorization, prototype theory (e.g., Homa, Gross, Cornell, Goldman, & Shwartz, 1973; Homa, Rhoades, & Chambliss, 1979; Love, Medin, & Gureckis, 2004; Mervis & Rosch, 1981; Minda & Smith, 2001; Posner & Keele, 1968; Smith & Minda, 2000). In accordance with how prototype models have traditionally been implemented, REP[P] assumes that each category is represented by an abstract prototype incorporating the features most common among the members of the category, that prototypes are abstracted during learning, and that retrieval is triggered by probe-prototype similarity. According to REP[P],

$$p(A) = \frac{sim(prototype_A)}{sim(prototype_A) + sim(prototype_B)} \tag{2}$$

where $sim(prototype_A)$ represents the similarity between object X and the prototype of category A and $sim(prototype_B)$ represents the similarity between object X and the prototype of category B. Hence, the probability that object X belongs to category A is determined by the similarity between object X and the prototype of category A relative to the summed similarity between object X and all relevant prototypes.

REP[ESAM] (Koehler et al., 2003), rooted in support theory (Tversky & Koehler, 1994), assumes that people store the relative prototypicality[6] for

[5] The rationale of this section is to describe the core features of the models. To make this section more easily accessible, simplified versions of the models are presented. For complete versions, please refer to Nilsson, H., Olsson, H., & Juslin, P. (2005). The cognitive substrate of subjective probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 600-620.

[6] I choose the term relative prototypicality to highlight the fact that the value denotes distinguishableness rather than commonness. Prototypicality refers to how common a feature is in a category. For example, the feature of "two legs" has high prototypicality in the category of "men" while the feature of "bearded" has relatively low prototypicality in the category of "men". The relative prototypicality, on the other hand, refers to how relatively common a feature is in a particular category. Therefore, the feature of "two legs" has a low relative prototypicality in the category of "men" (as it does not distinguish men from other types of humans, such as, for example, women or boys) while the feature of "bearded" has a high

each possible feature given each possible category. The relative prototypicality of feature j in relation to category A, $rp_j(A)$, is given by,

$$rp_j(A) = \frac{prop_j(A)}{prop_j(A) + prop_j(B)} \tag{3}$$

where $prop_j(A)$ represents the proportion of category A members with feature j and $prop_j(B)$ represents the proportion of category B members with feature j. Hence, the relative prototypicality of feature j in relation to category A is determined by how common feature j is among category A members relative to how common feature j is among the members of each relevant category ($rp_j(B)$ is derived in the same way). According to REP[ESAM],

$$p(A) = \frac{\sum rp_j(A)}{\sum rp_j(A) + \sum rp_j(B)} \tag{4}$$

where $j = 1 \ldots J$. Thus, according to REP[ESAM], p(A) is determined by the relative prototypicality of each feature incorporated in object X. As noted by Koehler et al. (2003), REP[ESAM] "can be viewed as employing a prototype representation of cue information, in which the interpretation of each cue value is uninfluenced by other cue values in the cue pattern" (p. 192).

REP[L] is based on the claim that "[T]he representativeness heuristic boils down to computing probabilities using only likelihoods, without prior probabilities" (Gigerenzer & Murray, 1987, p. 153; see also Villejoubert & Mandel, 2002). According to REP[L],

$$p(A) = \frac{prop_{identical}(A)}{prop_{identical}(A) + prop_{identical}(B)} \tag{5}$$

where $prop_{identical}(A)$ represents the proportion of category A members that are identical to object X and $prop_{identical}(B)$ represents the proportion of category B members that are identical to object X. Thus, p(A) is determined by the relative frequency of category A members that are identical to X relative to the relative frequency of members identical to X in all relevant categories.

The three implementations of the representativeness heuristic make distinct assumptions concerning how information is stored and concerning how representativeness is assessed. REP[P] assumes that each category in the

relative prototypicality in the category of "men" (as it distinguish men from other types of humans, such as, for example, women or boys).

environment is represented by a holistic representation embodying the features that are most common among the members of a particular category (i.e., a prototype). REP[P] assumes that the degree to which object X is judged to be representative of category A is determined by how similar object X is to the prototype of category A. REP[ESAM] assumes that for each available feature the relative prototypicality is stored. According to REP[ESAM], the relative prototypicality in relation to category A and category B of each feature incorporated in object X are used to produce the probability judgment that object X is a member of category A. REP[L] assumes that the relative frequency of category A members with the features of object X is stored and that this relative frequency determines the degree to which object X is judged to be representative of category A. Importantly, though they make distinct assumptions concerning how information is stored and how representativeness is assessed, REP[P], REP[ESAM], and REP[L] all estimate the inverse probability rather than the conditional probability (i.e., they estimate the probability that a category A member have the features of object X rather than the probability that an object such as object X is a member of category A). As a result, they can all account for biases such as the conjunction fallacy and base-rate neglect.

CBRF assumes that the cue-validity of each of object X′s features in relation to all relevant categories is stored in memory. Cue-validities are derived by Equation 1. According to CBRF, *p*(A) equals the cue-validity of the feature that is best at discriminating between category A members and category B members (i.e., the feature with the highest cue-validity). That is, if feature *H* has the highest cue validity, CBRF assumes that,

$$p(A) = cv_H(A) \tag{6}$$

The exemplar based PROBEX-model (Juslin & Persson, 2002) is a modification of the context model (Medin & Schaffer, 1978; Nosofsky, 1984). As in other exemplar models, PROBEX assumes that memory consists of exemplars corresponding to real-world objects. When the probability that object X belongs to category A is to be assessed, a sample of exemplars similar to object X is retrieved from memory. According to PROBEX,

$$p(A) = \frac{\sum sim(exemplar_i) * c}{\sum sim(exemplar_i)} \tag{7}$$

where $sim(exemplar_i)$ denotes the similarity between object X and exemplar_i (i = 1…number of retrieved exemplars) and c equals 1 if exemplar_i corresponds to a category A member and 0 if exemplar_i correspond to a category B member. Hence, according to PROBEX, the probability that object X belongs to category A is determined by the summed similarity between object X and the retrieved exemplars corresponding to category A members relative to the summed similarity between object X and all retrieved exemplars.

5.1.2. Method

All three experiments in Study I adopt basically the same design. For each experiment an experimental environment was constructed. The environments included a set of stimuli, described along a set of binary feature dimensions, spread across a set of categories. For example, the environment of Experiment 3 (presented in Table 1) included a set of bugs that were either dangerous (category A in Table 1) or harmless (category B in Table 1). The bugs varied in terms of four binary features (F1-F4 in Table 1): leg length (long or short), nose length (long or short), spots on the fore back (absent or present), and back pattern (each bug having one of two possible back patterns).

Table 1. *The Category Structure Used in Experiment 3 of Study I and in Both the Simulation and the Experiment of Study II*

	F1	F2	F3	F4	Freq. in category A	Freq. in category B
E1	0	0	0	0	0	14
E2	0	0	0	1	0	6
E3	0	0	1	0	0	1
E4	0	1	0	0	0	1
E5	1	0	0	0	5	1
E6	0	1	1	0	1	1
E7	1	0	0	1	1	1
E8	0	1	1	1	6	0
E9	1	0	1	1	1	0
E10	1	1	0	1	1	0
E11	1	1	1	0	1	5
E12	1	1	1	1	14	0

The three experiments were divided into blocks including a categorization-with-feedback phase and a probability-estimation phase. In the categorization-with-feedback phase participants learned to predict category given presented features. In the probability-estimation phase participants judged the probability of category given presented features. For example, in the categorization-with-feedback phase of Experiment 3 participants were presented with bugs and were asked to categorize them as either dangerous or harmless. After each categorization feedback concerning its accuracy was provided. In the probability-estimation phase participants were presented with bugs and the task was to assess either *p*(dangerous) or *p*(harmless). Experiment 1 and 3 included four blocks and Experiment 2 included two blocks.

The analysis focused mainly on cognitive modeling. The stimuli presented to the participants were fed into the "memory" of the five models. Then predictions for each feature combination existing in the environment were generated for each model.[7] To determine the fit of the models (i.e., how good they were at accounting for the observed subjective probabilities), model predictions were compared with the probability judgments of the participants. The difference between model predictions and judgments was calculated in terms of root mean squared deviation (*RMSD*). The fit of a model was considered to be good if the *RMSD* was relatively low (relative to the *RMSD* of the competing models) and it was considered to be poor if the *RMSD* was relatively high. This analysis was performed on both individual and group data.

Although the main emphasis was on the model fit, two other important analyses were made. First of all, Experiment 1-2 included so called critical exemplars. These were stimuli for which the models made qualitatively distinct predictions (i.e., predictions that were unaffected by the values on the free parameters). For example, in Experiment 1 there were three critical exemplars (CE1, CE2, and CE3) and the test-phase task was to assess the probability that each presented stimuli corresponded to a company with a stock that had increased or decreased in value during the last year. For these three critical exemplars, the models made the following predictions when the probability of an increasing stock value was to be assessed;

REP[P]: $p(\text{increase} \mid \text{CE1}) = p(\text{increase} \mid \text{CE3}) > p(\text{increase} \mid \text{CE2})$
REP[ESAM]: $p(\text{increase} \mid \text{CE3}) > p(\text{increase} \mid \text{CE2}) > p(\text{increase} \mid \text{CE1})$
REP[L]: $p(\text{increase} \mid \text{CE1}) > p(\text{increase} \mid \text{CE3}) > p(\text{increase} \mid \text{CE2})$
CBRF: $p(\text{increase} \mid \text{CE2}) > p(\text{increase} \mid \text{CE1}) > p(\text{increase} \mid \text{CE3})$
PROBEX: $p(\text{increase} \mid \text{CE1}) > p(\text{increase} \mid \text{CE2}) > p(\text{increase} \mid \text{CE3})$

By studying how the probabilities for the critical exemplars were assessed it was possible to distinguish between the models.

Second, the residual deviations between model predictions and subjective probabilities were studied. A model that captures the underlying cognitive process should be able to predict each subjective probability more or less equally well. Therefore, the residual deviations between model predictions and subjective probabilities should be randomly distributed and not reveal any systematic discrepancies. Or in other words, the residual deviations should be more or less equally large for all test-items. The analysis of the critical exemplars and the analysis of the residuals provided important complements to the model fit.

A problem with cognitive modeling is that due to between-model correlations it is often hard to distinguish between the competing models. There-

[7] Predictions were derived using optimized parameter-values. That is, the parameter values that produces the least deviation between model predictions and data was used. The parameter-values were optimized on data from each block for each participant.

fore, the environment of Experiment 1 was designed to maximize the chances of distinguishing between the five competing models. Experiment 2 tested how an increase in information-load during learning affected how subsequent subjective probabilities were assessed. It could, for example, be that a higher number of unique stimuli would lead to an initial prototype abstraction phase (e.g., Minda & Smith, 2001; Smith & Minda, 1998, 2000) or that more feature dimensions would require people to rely more on the most predictive feature dimensions (e.g., Timmerman, 1993). Therefore, the environment of Experiment 2 included more feature dimensions and more unique stimuli than the environment used in Experiment 1. Experiment 3 tested whether participants would be more inclined to abstract prototypes if the categories where made more homogenous (Minda & Smith 2001). Therefore, Experiment 3 used an environment with high within-category similarity and low between-category similarity.

5.1.3. Results

Across all three experiments, PROBEX provided superior fit. Across every block in each of the three experiments, PROBEX fit data at least as well as any other model. To display the overall advantage of PROBEX, Figure 1 presents the proportion of participants best fit by each model in the last block of each experiment (participants that were equally well fit by more than one model are presented as ties). In the last block of Experiment 1-2, PROBEX was best at accounting for the subjective probabilities for about two thirds of the participants. In Experiment 3, PROBEX together with REP[L] showed best fit to about two thirds of the participants. Given the poor fit of REP[L] in Experiment 1-2 and given the fact that no participant was better fit by REP[L], the relatively good fit of REP[L] in Experiment 3 is likely to be due to the fact that it provided predictions that correlated very strongly with the predictions by PROBEX. If the ties between REP[L] and PROBEX are considered as “wins” for PROBEX, PROBEX best accounts for the subjective probabilities for about two thirds of the participants also in Experiment 3. The fit to group data, that involved cross-validation, showed that the advantage of PROBEX was not due to over-fit.

Concerning the analysis of the critical exemplars, the results were mixed. In Experiment 1, this analysis provided support for PROBEX. In Experiment 2, it was impossible to distinguish between the models. Concerning the analysis of the residual deviations, PROBEX was supported in all three experiments.

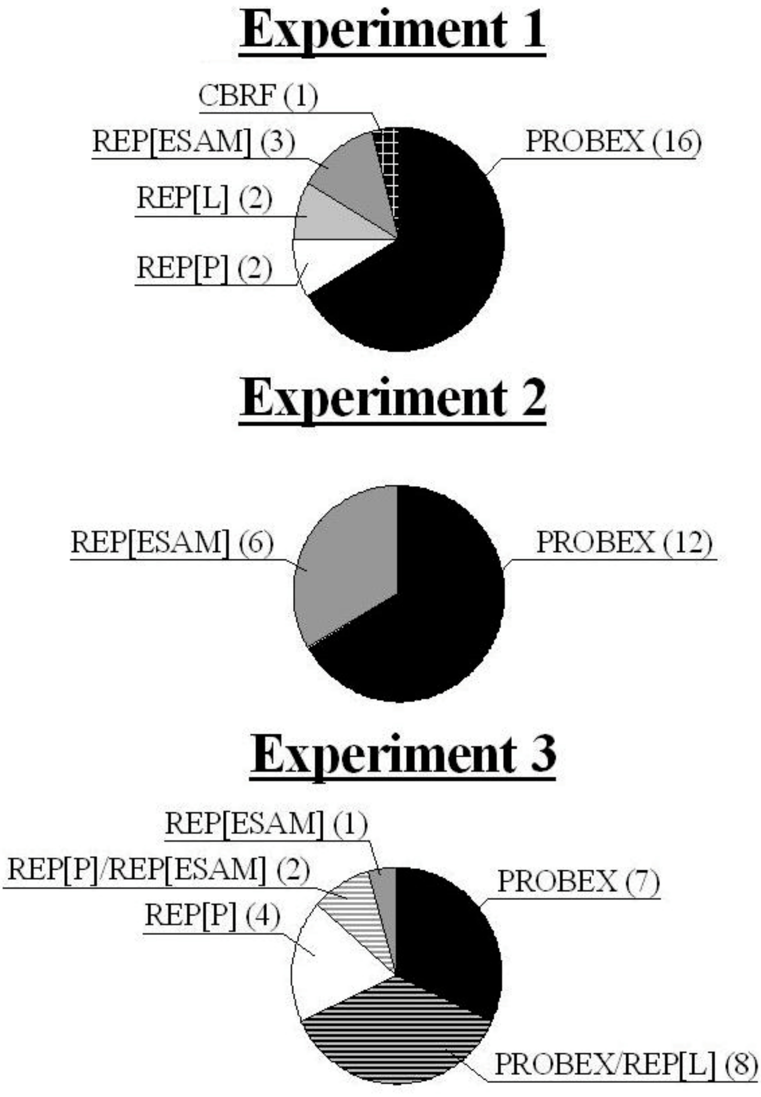

Figure 1. Study I: The proportion of participants best fit by each model in the last block of Experiments 1, 2 and 3 (the number of participants best fit by each model within brackets). Adopted from Nilsson, H., Olsson, H., & Juslin, P. (2005). The cognitive substrate of subjective probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 600-620.

5.1.4. Discussion

Throughout Study I PROBEX was the model that best explained how the subjective probabilities were generated. Despite the increase in information load in Experiment 2 (assumed to make exemplar memorization harder) and the increases in category homogeneity in Experiment 3 (assumed to make exemplar memorization less fruitful), two manipulations aimed to favor the competing models at the expense of PROBEX, the advantage of PROBEX prevailed. Hence, Study I indicated that subjective probabilities are based on exemplar memory.

The goal of this thesis is to explore the validity of the claim that subjective probabilities are routinely assessed by the representativeness heuristic. Study I offers some insights on this account. First, given the poor fit of REP[P], REP[ESAM] and REP[L], it appears as if the representativeness heuristic was not generally used to assess the subjective probabilities in the tasks used in Study 1. Second, despite their poor fit, REP[P] and REP[ESAM] appears as the most plausible implementations of the representativeness heuristic.

5.1.5. Cognitive Modeling and Model Evaluation

The goal of cognitive modeling is to find a model that describes the cognitive mechanisms behind a particular phenomenon. However, to do so is far from uncomplicated. Because cognitive modeling played such a crucial role in Study I, I will now briefly discuss some of the major problems with cognitive modeling and how these problems were addressed. First, it is often hard to determine whether a model that fits well does so because it captures the underlying cognitive processes or if it does so because it captures noise. One of the “problems” with human participants is that they tend to produce noisy data sets. In highly simplified terms, any dataset is generated by a cognitive process and random noise. This means that a model that is good at capturing noise might fit data well even if it does not describe the cognitive process (a phenomenon known as over-fit; see e.g., Myung & Pitt, 1997). When using measures such as *RMSD* or the Pearson’s coefficient of determination (r) to distinguish between models (as is done in Study I), it is impossible to determine the source of the fit. Therefore, if such measures are used carelessly, it is possible to end up with a “winning” model that has no psychological relevance.

The threat of over-fit can be reduced by increasing the number of test items, having participants making repeated judgments for each individual test item, using measures that take a models ability to capture noise into account (e.g., *minimum description length* as suggested by, Pitt, Myung, & Zhang, 2002), or cross-validating data (Olsson et al., 2004). In Study I the following methods for limiting the threat of over-fit were used. In Experiment 2 the number of unique test items was increased. In Experiments 2-3 the probability that test item X belonged to category A was assessed twice in each test-phase. And most importantly, in all three experiments data was cross-validated. If anything, reducing the impact of noise increased the relative advantage of PROBEX.

Second, it is easy to lose the psychological rationale behind the cognitive model. When defining a cognitive theory as a cognitive model, some aspects will be implemented as free parameters. The exact predictions by the model are determined by the exact values on these parameters. In the light of the underlying cognitive theory, some parameter values might make more sense. As a result, sometimes a model will be able to produce predictions that capture data but that does not reflect the essence of the cognitive theory. For example, with particular parameter values PROBEX becomes insensitive to similarity (Juslin & Persson, 2002). When so, if asked to assess the probability that Object X belongs to category A, PROBEX ignores the similarity between X and the retrieved exemplars and makes a response based solely on the relative frequency of the retrieved exemplars that belong to category A. If such a version of PROBEX, known as the "sloppy frequentist" version, is found to fit data, it is questionable whether it actually provides support for exemplar theory. The best method for not losing the underlying rational of the compared models is to generate parameter-free predictions (i.e., predictions that hold for each possible values on each of the models free-parameters), in this thesis referred to as qualitative predictions, distinguishing the models. In Study I this was done for the critical exemplars. Overall, the results were in line with the qualitative predictions of PROBEX.

Third, it is problematic to determine what constitutes good fit. A model that accurately describes the underlying cognitive process should only be able to account for the variations in data that are caused by the underlying cognitive process. As there is noise in data, perfect fit does not mean *RMSD* = 0 (*RMSD* = 0 would indicate extreme over-fit). As a result, it is never possible to determine whether a model accurately describes the underlying cognitive process, it is only possible to determine whether it is relatively good or relatively poor at describing the underlying cognitive process. Therefore, the most important aspect of Study I is not that it provides support for PROBEX, the most important aspect is that it shows that REP[P], REP[ESAM], and REP]L] are relatively poor at accounting for how the subjective probabilities were assessed.

5.2. Study II: Are Representativeness Effects Caused by the Representativeness Heuristic?

5.2.1. Introduction

One consequence of the representativeness heuristic using a prototypical exemplar to represent the category, as suggested by Kahneman and Frederick (2002), is that it predicts representativeness effects (described in Chapter 4.2). Imagine an object that is common in category A but uncommon in category B. Furthermore, imagine that the object shares few features with the prototype of category A but that it shares many features with the prototype of category B (i.e., the object is unrepresentative of category A and representative of category B) . In this case, a representativeness effect occurs when the probability that the object belongs to category B is judged to be higher than the probability that it belongs to category A.

The representativeness heuristic is only suggested to be used in a subset of all situations where subjective probabilities are assessed (e.g., Kahneman et al., 1982). One interpretation of the lack of support for the models implementing representativeness heuristic in Study I could be that the tasks did not particularly trigger the representativeness heuristic. Therefore, Study II focused on representativeness effects. If it is the case the representativeness heuristic is a commonly used tool for assessing subjective probabilities, then it is highly likely (I would argue) that the heuristic has been involved if representativeness effects occur. Study II explores whether REP[P] (the implementation of the claim made by Kahneman & Frederick, 2002) provides a good account of subjective probabilities when representativeness effects are present. To do so, REP[P][8] was contrasted with PROBEX (Juslin & Persson, 2002) using a task specially designed to elicit representativeness effects.

Besides its good performance in Study I, PROBEX is particularly interesting as exemplar based categorization models have been shown to explain so-called prototype enhancements effects (e.g., Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky & Johansen, 2000; Shin & Nosofsky, 1992). The most classic prototype enhancement effect is when “prototypes that are not experienced during training are nevertheless classified as well as or sometimes better than the old training exemplars” (Nosofsky & Johansen, 2000, p. 375). Prototype enhancement effects have been found to be most pronounced if participants are tested either early in learning (e.g., Minda & Smith, 2001; Smith & Minda, 1998; Smith, Murray, & Minda, 1997) or some days after the learning have occurred (Posner & Keele, 1970; Strange, Keeney, Kessel, & Jenkins, 1970). That is, prototype enhancement effects are most pronounced when knowledge is relatively low. Interestingly, repre-

[8] REP[ESAM] was also tested. However, because REP[P] and REP[ESAM] provided indistinguishable predictions, only REP[P], the model that implements the claim by Kahneman and Frederick (2002), was included in the manuscript.

sentativeness effects can be seen as a form of prototype enhancement effect. And accordingly, as exemplar models have been found to account for prototype enhancement effects, it is possible that PROBEX can account for representativeness effects.

Study II used the environment from Experiment 3 of Study 1 (presented in Table 1). Regarding representativeness effects, stimuli E5 and E11, referred to as critical exemplars, are most interesting. Both critical exemplars share more features with the prototype for the category to which they belong less often. Therefore, representativeness effects were expected for the critical exemplars.

Study II was divided into two parts, a simulation and an experiment (both using the environment in Table 1). The simulation explored two questions. First, does both REP[P] and PROBEX predict representativeness effects? Second, if they predict representativeness effects, when do the models predict them to occur? The experiment explored three questions. First, will representativeness effects occur? Second, when do they occur? Third, which model shows best fit to data when representativeness effects are present?

5.2.2. Simulation

The goal of the simulation was to test (*a*) if both REP[P] and PROBEX predict representativeness effects and, if they do, (b) when they predict them to be most pronounced. To do so, the level of knowledge was manipulated by varying the amount of information available to the models. The simulations showed that both models predicted representativeness effects. REP[P] predicted that representativeness effects should be more and more accentuated as the level of knowledge increases. PROBEX predicted representativeness effects to be most accentuated when knowledge is relatively low, and that the representativeness effects should then gradually disappear as knowledge increases. If representativeness effects are similar to prototype enhancement effects, the prediction by PROBEX, i.e. that representativeness effects should be most pronounced when knowledge is low, is more plausible.

5.2.3. Experiment

As in Experiment 3 in Study I, the task was to learn to categorize bugs as dangerous or harmless. On each trial a bug was presented, and the task was to first categorize it and then to assess the probability that the bug actually belonged to the category it had been categorized under. After the probability assessment, feedback concerning the accuracy of the categorization was provided. The experiment was divided into four blocks with 60 trials in each block (the exemplars presented in Table 1 was presented once in each block). There were two conditions, one *standard* and one *complex*. The key difference between the two conditions was that the relevant feature dimensions were easier to identify in the standard condition. It was assumed that this made the complex condition relatively harder to learn.

The simulation revealed that REP[P] and PROBEX predicted representativeness effects to be accentuated at different levels of knowledge. Given the design of the experiment, the models provided conflicting predictions. While REP[P] predicted representativeness effects to be most accentuated in the final blocks and especially so in the standard condition, PROBEX predicted representativeness effect to be most accentuated early in learning and that they should prevail longer in the complex condition.

Data from the experiment was in line with both the qualitative and the quantitative predictions by PROBEX. Representativeness effects were most accentuated early in learning and gradually disappeared as a function of learning. The relationship between low levels of knowledge and strong representativeness effects was further displayed by the finding that the representativeness effects only vanished in the standard condition. Regardless of whether representativeness effects were present or not, PROBEX showed at least as good quantitative fit as REP[P]. Cross-validation showed that the advantage of PROBEX, observed in both individual and group data, was not due to over-fit.[9]

5.2.4. Discussion

Independently of whether representativeness effects were present or absent, PROBEX was better than REP[P] at describing how the subjective probabilities were assessed. Thus, as the results in Study I, these results suggest that subjective probabilities are based on exemplar memory. The key finding of Study II was that REP[P] (the model implementing the ideas of Kahneman & Frederick, 2002) showed a poor fit to data when representativeness effects were observed. Remember, the representativeness heuristic is only suggested to be involved in some subjective probability judgments (Kahneman et al., 1982). Other subjective probabilities are argued to be assessed either by different heuristics or by some algorithm belonging to System 2. This means that from the heuristics and biases perspective, if representativeness effects disappear, it could be due to participants shifting to a different method of assessing subjective probabilities. If so, then even if REP[P] correctly describes the cognitive process generating representativeness effects, it should

---

[9] Interestingly, as shown by Nilsson, Juslin, and Olsson (2003), if the data from the experiment in Study II was not cross-validated, REP[P] showed best quantitative fit in the early blocks (i.e., when representativeness effects were most pronounced). That is, if data was not cross-validated, it appeared as subjective probabilities were assessed by the representativeness heuristic in the early blocks of the experiment. Hence, Nilsson et al. (2003) demonstrates the importance of using methods such as cross-validation to limit the threat of over-fit.

fit poorly when representativeness effects are absent. However, the finding that REP[P] fits data poorly also when representativeness effects are present effectively shows that representativeness effects are unlikely to be caused by a process similar to the one implemented by REP[P].

Though conflicting with the ideas of Kahneman and Frederick (2002), the results were not overly surprising. The prototypical enhancement effect, closely related to the representativeness effect, is a well documented phenomenon in the literature on categorization (e.g., Mervis & Rosch, 1981). As for the representativeness effects in Study II, prototype enhancement effects are most accentuated when knowledge is low. Furthermore, as for the representativeness effects in Study II, prototype enhancement effects are best accounted for by exemplar models (Nosofsky & Johansen, 2000).

Study II pinpoints a weakness of the heuristics and biases approach. A central idea in the heuristics and biases approach is that studies of biases will reveal underlying heuristics. Critics have often argued that this research method produces weak cognitive theories that do not hold for stringent tests (e.g., Gigerenzer, 1996). After observing representativeness effects it might be tempting to conclude that they are caused by the representativeness heuristic. After all, representativeness effects carry more characteristics associated with the representativeness heuristic than characteristics associated with, for example, exemplar memory. However, when the assumptions of the representativeness heuristic were thoroughly tested, it turned out that the representativeness effects could not be attributed to the representativeness heuristic.

5.3. Study III: Are Conjunction Fallacies Caused by the Representativeness Heuristic?

5.3.1. Introduction

Study III focused on one of the biases attributed to the representativeness heuristic assessing the inverse rather than the conditional probability, namely the conjunction fallacy (Tversky & Kahneman, 1983). The two major theories on the cause of the conjunction fallacy, the RHH (Tversky & Kahneman, 1983) and the CRH (e.g., Gavanski & Roskos-Ewoldsen, 1991), were contrasted. According to the RHH, conjunction fallacies occur when the target object (e.g., Linda) is more representative of the category targeted in the conjunction (e.g., feminist bank-tellers) as compared to the category targeted in the component (e.g., bank-teller). Hence, according to the RHH, a mental representation of the category targeted in the conjunction (e.g., a representation of feminist bank-tellers) is a precondition for the conjunction fallacy to appear. According to the CRH[10], people combine p(A) and p(B) into p(A&B) in a way that causes conjunction fallacies. An interesting aspect of the CRH is that it only predicts conjunction fallacies when p(A&B) can not be directly assessed. That is, the CRH only predicts conjunction fallacies when respondents lack a representation of the category targeted in the conjunction (e.g., of feminist bank-tellers). It is this difference that is addressed in Study III.

Two manipulations were made in Study III. First, whether or not the participants had a representation of the category targeted in the conjunction was manipulated. In Linda terms, it was manipulated whether or not the participants had knowledge about feminist bank-tellers. Second, whether or not the representativeness heuristic was applicable for computing the probability of the conjunction was manipulated. According to the RHH, the conjunction fallacies should be committed only when the representativeness heuristic is applicable and the rate of conjunction fallacies should be highest among the participants with a representation of the category targeted in the conjunction. The CRH predicts a high rate also when the representativeness heuristic is not applicable and more conjunction fallacies among those who have no representation of the category targeted in the conjunction.

5.3.2. Method

A category learning design was used. The strategy was to create a category structure that could be experienced as either containing four component categories or as containing four conjunction categories. Two binary feature dimensions were used to create the four component categories. The component categories were *employed in private sector*, *employed in public sector*, *has university degree* and *has no university degree*. The conjunction categories were created by combining the features of the two binary feature dimensions. The conjunction categories were *teacher* (has university degree and employed in the public sector), janitor (has no university degree and employed in the public sector), engineer (has university degree and employed in the private sector), and welder (has no university degree and employed in the private sector).

Participants were assigned to either the *component condition* or to the *conjunction condition*. In the learning phase, participants were trained to categorize two-featured personality sketches. In 50% of the trials participants in the component condition had to judge whether or not the person described in the sketch had a university degree and in 50% of the trials they had to judge in which sector the person described in the sketch was em-

[10] What is described as the CRH is actually a family of hypotheses suggesting different combination-rules. Since no attempt is made to differentiate between these I found it more useful to consider them as one unitary hypothesis.

ployed. In every trial participants in the conjunction condition judged the profession (i.e., teacher, janitor, engineer, or welder) of the person described in the personality sketch. In both conditions feedback concerning the accurate response was provided after each judgment.

For clarity, the logic of the learning phase was to create two distinct representations of the same environment. In Linda terms, it was assumed that participants in the conjunction condition, but not participants in the component condition, should create representations corresponding to feminist bank-tellers. According to the RHH, an ability to create a mental representation of the category targeted in the conjunction is a precondition for the conjunction fallacy. Therefore, the RHH predicts at least as many conjunction fallacies among participants in the conjunction condition as among participants in the component condition. The CRH, on the other hand, only predicts conjunction fallacies when a representation of the category targeted in the conjunction is absent. Hence, in contrast to the RHH, the CRH predicts more conjunction fallacies among participants in the component condition than among participants in the conjunction condition.

To manipulate whether or not the representativeness heuristic could be applied, the post learning test-phase included standard problems as well as so called *mixed problems*. Standard problems were designed to be equivalent to the standard problems of Tversky and Kahneman (1983). For example, participants could be presented with person X and the task could be to asses *p*(employed in private sector), *p*(has university degree), and *p*(employed in the private sector and has university degree). *Mixed problems* (adapted from Gavanski & Roskos-Ewoldsen, 1991) included two personality sketches (described as P1 and P2) and the task could, for example, be to assess *p*(P1 is employed in private sector), *p*(P2 has university degree), and *p*(P1 is employed in the private sector and P2 has university degree). Remember, a probability assessed by the representativeness heuristic is determined by how representative the target object is of the target category. Since the statement "P1 is employed in the private sector and P2 has university degree" has two target objects, the representativeness heuristic can not be used to assess the probabilities of the conjunctions in the mixed problems. Accordingly, since it is the representativeness heuristic that is argued to cause the conjunction fallacy, the RHH predicts no conjunction fallacies on mixed problems. However, since the mixed problems involve probability combination, the CRH predicts high levels of conjunction fallacies on mixed problems. In sum, the RHH predicts conjunction fallacies only on standard problems (and at least as many in the conjunction condition as in the component condition) while the CRH predicts high levels of conjunction fallacies on mixed problems in both conditions and on standard problems in the component condition.

5.3.3. Results[11]

As in the original standard problem, each problem (both standard and mixed) included one likely component and one unlikely component. Figure 2 shows the mean probabilities assessed to likely components (upward arrows), unlikely components (downward arrows), standard conjunctions (open circles), and mixed conjunctions (closed circles). Three aspects of Figure 2 are particularly interesting. First, both component learners and conjunction learners had a clear view of which features were predictive of which component category. Second, while the probabilities assessed to standard and mixed conjunctions were similar among component learners, the standard conjunctions were assessed to be less likely than the mixed conjunctions among conjunction learners. Third, for conjunction learners the probabilities assessed to standard conjunctions were only marginally higher than the probabilities assessed to the unlikely conjuncts.

The rate of conjunction fallacies was found to be high both for standard and mixed problems. The lowest rate of conjunction fallacies was observed on standard problems solved by conjunction learners. Hence, conjunction fallacies were highly frequent when the RHH predicted them to be absent and they were least frequent when the RHH predicted them to be most common. Overall, data showed strong support for the hypothesis that an inability to properly combine probabilities is an important factor behind conjunction fallacies.

5.3.4. Discussion

The conjunction fallacy is often attributed to the fact that the representativeness heuristic assesses the inverse rather than the conditional probability. If so, then (*a*) knowledge about the members of the conjunction category should increase the rate of conjunction fallacies and (b) an inability to apply the representativeness heuristic should decrease the rate of conjunction fallacies. Study III failed to find support for these two assumptions. On the contrary, knowledge about the members of the conjunction category decreased the rate of conjunction fallacies and an inability to apply the representativeness heuristic increased the rate of conjunction fallacies. These results indicate that an inability to properly combine probabilities, but not the representativeness heuristic, is an important factor driving conjunction fallacies.

[11] Study III included two experiments. In the result presentation below, data from these two experiments are combined. In the presentation below a *component learner* is a participant that only encountered component trials during learning (i.e., a participant from the component condition of Experiment 1) and a *conjunction learner* is a participant that encountered conjunction trials during learning (i.e., a participant from the conjunction condition of Experiment 1 or a participant from Experiment 2).

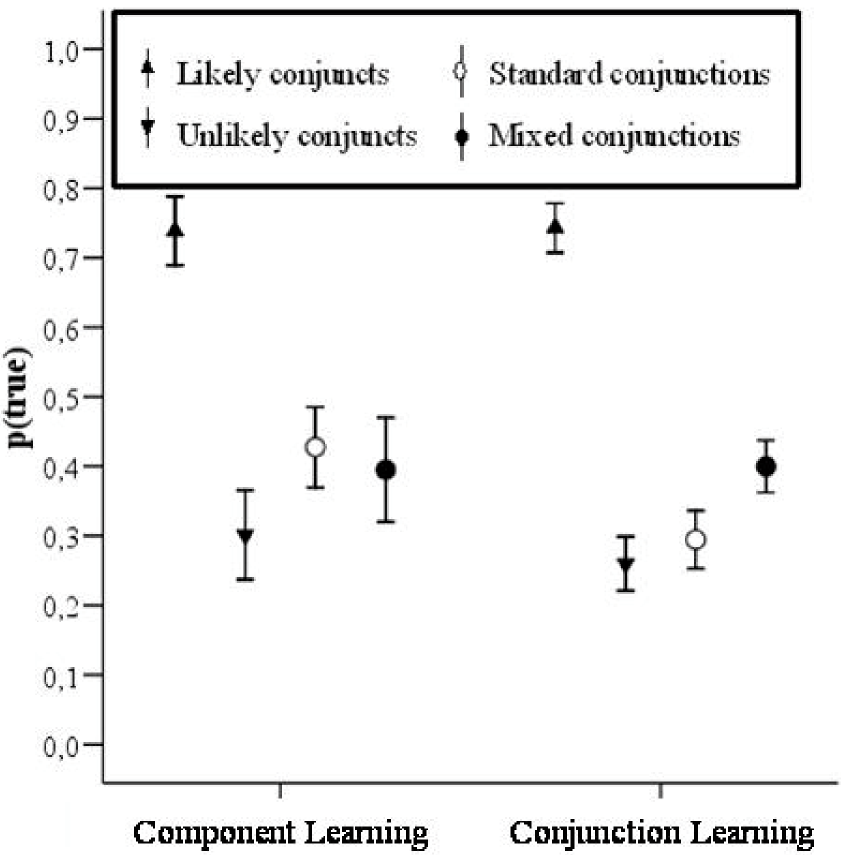

Figure 2. Study III: The probabilities assessed to likely components (upward arrows), unlikely components (downward arrows), standard conjunctions (empty circles), and mixed conjunctions (closed circles) for component learners and conjunction learners separately.

Study III is not the first to question whether the representativeness heuristic is one of the main sources of the conjunction fallacy (see e.g., Fantino et al., 1997; Fisk, 2002; Fisk & Pidgeon, 1996; Zizzo, 2003). In the few studies contrasting the RHH with alternative explanations concerning the cause of the conjunction fallacy, the results have always been negative for the RHH (Gavanski & Roskos-Ewoldsen, 1991; Thüring & Jungermann, 1990). Add the fact that the representativeness heuristic is unlikely to cause base-rate neglect (Gavanski & Hui, 1992), and you end up with the situation that there is no bias that can be seriously attributed to the representativeness heuristic assessing the inverse probability.

The findings in Study III relate to Study I and Study II in a very important way. One of the major arguments against the theory that subjective probabilities are based on exemplar memory is that it can not account for conjunction fallacies. According to exemplar memory theory, the presentation of Linda triggers the retrieval of a set of exemplars. Because it is the description of Linda and not the label of the target category (i.e., *bank-teller* or *feminist bank-teller*) that determines retrieval, the same sample of exemplars will be formed when *p*(Linda is a bank-teller) and when *p*(Linda is a feminist bank-teller) is to be assessed. Because Linda can not be more similar to the exemplars corresponding to feminist bank-tellers than she is to the exemplars corresponding to bank-tellers, exemplar theory has trouble accounting for conjunction fallacies. However, as data in Study III indicates that the conjunction fallacy is not due to the way people assess the subjective probabilities but to the way they combine them, Study III provide an explanation of the conjunction fallacy that is compatible with exemplar theory. Thus, the finding in Study III neutralizes the major argument against the theory that the cognitive substrate of subjective probability is exemplar memory.

6. General Discussion

6.1. Summary of Results

The aim of this thesis was to make a thorough analysis of the representativeness heuristic. According to the heuristics and biases approach, the representativeness heuristic is one of the heuristics available for assessing subjective probabilities (Kahneman et al., 1982; Kahneman & Tversky, 1972). Several aspects of the representativeness heuristic have been argued to cause biases. For example, biases such as misperception of chance, insensitivity to sample size, and misconception of regression are associated with the alleged attribute substitution of the representativeness heuristic (Kahneman & Frederick, 2002) and biases such as the conjunction fallacy and base-rate neglect are ascribed to the representativeness heuristic assessing the inverse rather than the conditional probability (Tversky & Kahneman, 1974, 1983).

Study I focused on the cognitive representations and processes argued to be involved when the representativeness heuristic is used. Three distinct interpretations of the representativeness heuristic (Kahneman & Frederick, 2002; Koehler et al., 2003; Villejoubert & Mandel, 2002) were implemented as cognitive models and contrasted against one model using cue-validities (Gigerenzer et al., 1991) and one model based on exemplar memory (Juslin & Persson, 2002). Across three experiments, no support for the models based on the representativeness heuristic was found. Data showed that at least in the types of tasks used in Study I, subjective probabilities are not typically assessed using the representativeness heuristic. Hence, Study I raised concerns about the validity of the assumptions concerning cognitive representations and processes imbedded in the theoretical package of the representativeness heuristic.

Study II focused on the representativeness effect, a judgment phenomenon that can be linked to the proposed attribute substitution of the representativeness heuristic. The prototype based REP[P] model, implementing the representativeness heuristic as defined by Kahneman & Frederick (2002), was contrasted against PROBEX. An initial simulation showed that both models predicted representativeness effects, but that they predicted them at different stages of learning. In a subsequent experiment, representativeness effects were observed. Interestingly, PROBEX was the model that best predicted when the representativeness effects would be most pronounced and it was the model that showed best quantitative fit to data when representativeness effects were observed. Study II showed that although the representativeness effect is a phenomenon that incorporates many key aspects of a bias associated with the attribute substitution of the representativeness heuristic, it is unlikely to be caused by the representativeness heuristic. Study II highlights the major problem with basing cognitive theories solely on observations of biases.

Study III focused on the conjunction fallacy. The hypothesis that the conjunction fallacy is caused by the representativeness heuristic (RHH; Kahneman & Tversky, 1983) was contrasted with a hypothesis suggesting that the conjunction fallacy is caused by an inability to combine probabilities (CRH; e.g., Gavanski & Roskos-Ewoldsen, 1991). Data from two experiments showed that the rate of conjunction fallacies peaked when the RHH predicted conjunction fallacies to be least common and that the rate of conjunction fallacies hit bottom when the RHH predicted conjunction fallacies to be most common. Interestingly, this data pattern was predicted by the CRH. Study III clearly showed that the RHH provides an inadequate explanation of why conjunction fallacies are so frequently observed. In sum, studies I-III failed to find any substantial support for the theory that the representativeness heuristic is routinely used when subjective probabilities are assessed.

6.2. Implications for the Representativeness Heuristic

As previously discussed, the scientific strategy of the heuristics and biases approach has been to study biases and to reconstruct the heuristics that have generated them (Kahneman et al., 1982). During the last thirty years, a number of biases have been identified and a set of heuristics argued to cause them have been proposed (for a review see, Gilovich et al., 2002). The critics of the heuristics and biases approach have questioned (*a*) whether the observed biases are anything more than experimental artifacts (e.g., Gigerenzer, 2005; Juslin, 1994; Winman, 1997), (b) whether the heuristics truly are models of real cognitive processes or only redesriptions of data (e.g., Gigerenzer, 1996), and (c) whether the observed biases actually are caused by the suggested heuristics (e.g., Gavanski & Hui, 1992; Hertwig, Pachur, & Kurzenhäuser, 2005; Thüring & Jungermann, 1990). In the following section I will discuss how the results of this thesis relate to each of these three points.

First, is the conjunction fallacy an experimental artifact or a true behavioral bias? This thesis does not deal with issues concerning whether the conjunction fallacy necessarily has to be seen as an error (as discussed by e.g., Gigerenzer, 1996; Kahneman & Tversky, 1996; Bonini et al., 2004). However, the thesis does show that the conjunction fallacy is a phenomenon that is not restricted to the standard problem. Though the task used in Study III borrows features from the standard problem, it does deviate from the standard problem in two important respects. (*i*) The relevant frequencies are experienced by repeated sampling from the environment. Though such a procedure has been shown to sometimes change judgment behavior (e.g., Hertwig, Barron, Weber, & Erev, 2006), the conjunction fallacy was still observed. (*ii*) The rate of conjunction fallacies was varied by manipulating the cognitive representations of the respondents and not only by manipulating the formulation of the task. If the conjunction fallacy only had been an artifact of the formulation of the standard problem (as argued by e.g., Dulany & Hilton, 1991) the rate of conjunction fallacies should have been equally high among component learners and conjunction learners (since they performed identical tasks). Considering the overall results, Study III joins a group of studies showing that the conjunction fallacy is a phenomenon that expands beyond the standard problem (e.g., Nilsson & Winman, 2007; Zizzo, 2003).

Second, does the formulation of the representativeness heuristic capture a true cognitive process? This question is rather difficult to answer. The reason is that the representativeness heuristic is only argued to be used in some situations; namely, in those situations where probabilities are generated by System 1 and where System 1 does not use any of the other heuristics available for assessing probabilities. Data indicated that the representativeness heuristic was not typically triggered by the tasks used in Study I and Study II. However, this, of course, does not prove that the representativeness heuristic is not included in the human cognitive repertoire. Still, I would argue that at least as defined by Kahneman and Frederick (2002), it is quite unlikely that the representativeness heuristic plays a key role in human cognition. To date, there is no study focusing on the cognitive mechanisms behind subjective probability that validates their claim and studies on categorization frequently have shown the shortcomings of prototype models (e.g., Nosofsky & Johansen, 2000).[12]

Third, is the conjunction fallacy caused by the representativeness heuristic? The key factor in the heuristics and biases approach is the link between heuristics and biases. Given the adopted research strategy, the main proof of existence for heuristics such as representativeness or availability is obtained from observations of their associated biases. One thing that often appears to be neglected is that the mere existence of the biases does in no way prove the existence of the heuristics. Only when the heuristics can be used to explain variations in the occurrences of the biases (i.e., when the rate of biases increases or decreases) can the observations of the biases be considered to

[12] Some authors have attributed the lack of support for prototype models to limitations of the tasks used in studies on categorization (e.g., Markman & Ross, 2003; Minda & Smith, 2001; Smith & Minda, 1998). The typical task involves learning to categorize a small number of relatively simple stimuli (simple in a sense that they can vary on very few feature-dimensions) into either of two categories. It has been argued that prototype models would assemble more support if more complex learning tasks were used. In the spirit of this argument, it is possible that more complex learning tasks would induce usage of the representativeness heuristic.

indicate that the heuristics have been used. Interestingly, the representativeness heuristic does not account for any of the stable variations in the rate of the conjunction fallacy that have been found. For example, the representativeness heuristic can not be used to explain why the rate of conjunction fallacy is affected by response mode and by statistical sophistication (e.g., Hertwig & Chase, 1998; Wedell & Moro, in press) or why the probability of a conjunction is mainly determined by the probability of the relatively unlikely component probability (e.g., Fantino, et al., 1997; Fisk, 2002). Regarding Study III, the representativeness heuristic can not be used to explain why conjunction fallacies occur in mixed problems or why relatively more conjunction fallacies in standard problems was observed among component learners. Given the representativeness heuristic's inability to account for crucial aspects of the conjunction fallacy phenomenon, it is safe to say that the existence of the conjunction fallacy does not provide any evidence for the representativeness heuristic. Given the representativeness heuristic's inability to account for variations in biases such as the conjunction fallacy or base-rate neglect (Gavanski & Hui, 1992) and given the lack of support for the existence of the cognitive mechanisms argued to be involved in the representativeness heuristic, one can not but question the claim that the representativeness heuristic is routinely used to assess subjective probabilities.

6.3. Exemplar Memory, the Cognitive Substrate of Subjective Probability

Throughout this thesis an underlying theme has been to increase the knowledge about which cognitive representations and processes are involved in subjective probability judgment. Conflicting with influential theories suggesting that subjective probabilities are based on heuristics (e.g., Kahneman & Tversky, 1972; Tversky & Kahneman, 1973) or cue-validities (e.g., Gigerenzer et al., 1991; Juslin, 1994), the results have consequently suggested exemplar memory as the cognitive substrate of subjective probability (for similar findings see, Bearden & Wallsten, 2004; Dougherty et al., 1999; Sieck, 2003; Sieck & Yates, 2001).

To explain how exemplar memory is involved in subjective probability assessment, let us return to the wine-example from the Introduction. Exemplar theory assumes that as you encounter different wines, exemplars containing the tastes, colors, smells and so on of these wines are stored in memory. When you take a sip of the wine, the smell, taste, and texture of the wine trigger the retrieval of exemplars corresponding to wines with similar attributes. Some of the retrieved exemplars will correspond to white wines and some to red wines. The subjective probability that it is red is determined by the summed similarity between the tasted wine and the retrieved exemplars corresponding to red wines relative to the summed similarity between the tasted wine and all retrieved exemplars. The probability that the wine is white will be assessed in the same way. If the probability that the wine is red exceeds the probability that it is white, you decide that it is red. Otherwise you decide that it is white. And if you are lucky, you will soon have one exemplar corresponding to your friend being disgruntled and one exemplar corresponding to the experience of the luxurious wine engraved in memory.

Besides subjective probability assessment, a range of other cognitive activities appear to be influenced by exemplar memory. Nosofsky and colleagues have shown that the generalized context model can account for a range of perceptual-classification phenomena and provide excellent fit to human classification data in both environments with linear and non-linear separable categories (Nosofsky, 1984, 1986; Nosofsky & Johansen, 2000; Nosofsky & Zaki, 2002; Stanton, Nosofsky, & Zaki, 2003). The connectionist model ALCOVE (Kruschke, 1992) has proven to be a good model of category learning. Juslin and colleagues have shown that multiple-cue judgments are often based on exemplar memory (e.g., Juslin et al., 2003; Juslin et al., 2008; Karlsson, Juslin, & Olsson, in press). Add that exemplar models have been successfully applied in fields such as attention (Logan, 2002), social cognition (Smith & Zarate, 1992), memory (Hintzman, 1986, 1988), and language (Daelemans, 1995) and it is reasonable to believe that exemplar memory plays a key role in the human cognition system.

6.4. Which Combination-Rule is Causing the Conjunction Fallacy?

Study III indicated that people commit the conjunction fallacy because they use an improper rule for combining component probabilities. But what characterizes this combination-rule? Several hypotheses have been proposed (e.g., Abelson, Leddo, & Gross, 1987; Einhorn, 1985; Fantino et al., 1997; Wyer, 1976). In the following section I will discuss one of these, the weighted average hypothesis (further on referred to as *WAH*; e.g., Fantino et al., 1997; Gavanski & Roskos-Ewoldsen, 1991; Zizzo, 2003). The WAH suggests that people assess p(A&B) by taking a weighted average between p(A) and p(B). The average is weighted in the sense that the probability of the relatively unlikely component is more heavily weighted.

The WAH is interesting for at least two reasons. First, among the hypotheses suggesting different combination-rules, the WAH has attracted most empirical support (e.g. Fantino et al., 1997; Fisk, 2002; Gavanski & Roskos-Ewoldsen, 1991; Wells, 1985; Zizzo, 2003). Second, there appear to be several advantages, both cognitive and ecological, of using the weighted average of the component probabilities as a substitution for the product of the component probabilities. Let us begin with the cognitive side. Studies have indicated that people have problems with learning multiplicative relationships (e.g., Juslin et al., 2008). The suggestion has been that this difficulty might be due to the human mind being predisposed to addition (e.g., Juslin et al., 2003, 2008; Karlsson et al., in press). While the normative rule for combining component probabilities into conjunction probabilities is multiplicative, the rule suggested by the WAH is additive. It could be that the strategy to use a weighted average rule is the best possible choice for an additive system.

On the ecological side, the WAH has several valuable features. It could be argued that the exact subjective probabilities are not overly important in real life. Rather, the most important role of subjective probabilities is that they help humans to evaluate options. By assessing subjective probabilities, options can be ranked according to how likely they are to lead to a desired goal (for a similar discussion see, Gavanski & Roskos-Ewoldsen, 1991). Interestingly, as will be demonstrated below, if conjunction probabilities are compared, rankings based on the weighted average of p(A) and p(B) tends to correlate to a very high degree with rankings based on the product of p(A) and p(B).

Noise (caused by cognitive and environmental factors) will sometimes make it more useful to base the rankings on the weighted average of p(A) and p(B) as compared to the product of p(A) and p(B). This is illustrated by Figure 3, which presents the results from a simulation (Winman, Nilsson, Juslin, & Hansson, 2008). Table 2 illustrates how the simulation was performed (data in Table 2 are not taken from the simulation). Initially, a set of component probabilities were randomly generated (referred to as original component probabilities). These were randomly paired, creating a set of conjunctions. The probabilities of the conjunctions (referred to as original conjunction probabilities) were calculated using the normative combination-rule:

$$p(\mathrm{A\&B}) = p(\mathrm{A})*p(\mathrm{B}) \tag{8}$$

where p(A) and p(B) are the two component probabilities and p(A&B) is the conjunction probability. These original component and conjunction probabilities are presented in the three first columns of Table 2 (labeled original).

In the second part of the simulation, noise was added to the original component probabilities (the simulation was repeated using different levels of noise). The normative rule and the weighted average rule were used to combine the noisy component probabilities into conjunction probabilities. The normative combination rule was implemented by,

$$norm = p_n(\mathrm{A})*p_n(\mathrm{B}) \tag{9}$$

where p_n(A) equals p(A) plus noise and p_n(B) equals p(B) plus noise. The weighted average rule was implemented by a linear model. If $p_n(A) \leq p_n(B)$, the combination rule suggested by the WAH was implemented by,

$$add = \alpha + [\beta_{low} * p_n(A)] + [\beta_{high} * p_n(B)] \qquad (10)$$

where *add* equals the conjunction probability, α represents the intercept, β_{low} represents the weight put on the relatively unlikely conjunct probability and β_{high} represents the weight put on the relatively likely conjunct probability. To minimize the deviation between the probabilities calculated by Equation 8 and 10, α, β_{low} and β_{high} were optimized using standard linear regression (medians: -.10, .57, and .21 for α, β_{low} and β_{high} respectively). As shown by the weights, *add* is relatively more dependent on the unlikely component probability. The upper half of Table 2 present component and conjunction probabilities generated when the level of noise was relatively low and the lower half of Table 2 present component and conjunction probabilities generated when the level of noise was relatively high. Finally, the deviations (in terms of *RMSD*) and correlations between the original conjunction probabilities and *norm* and *add* were computed.

Table 2. *A Descriptive Example of how Data in Figure 3 was Generated*

p(A)	p(B)	p(A&B)	p_n(A)	p_n(B)	*norm*	*add*
			Noise Level 1			
.90	.70	.63	.91	.72	.66	.74
.80	.70	.56	.81	.68	.55	.69
.50	80	.40	.48	.80	.38	.51
.40	.50	.20	.40	.51	.20	.41
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.10	.00	.00	.09	.01	.00	.02
			Noise Level 2			
.90	.70	.63	.99	.65	.53	.68
.80	.70	.56	.70	.75	.32	.71
.50	80	.40	.43	.75	.28	.46
.40	.50	.20	.51	.55	.53	.51
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.10	.00	.00	.00	.10	.00	.01

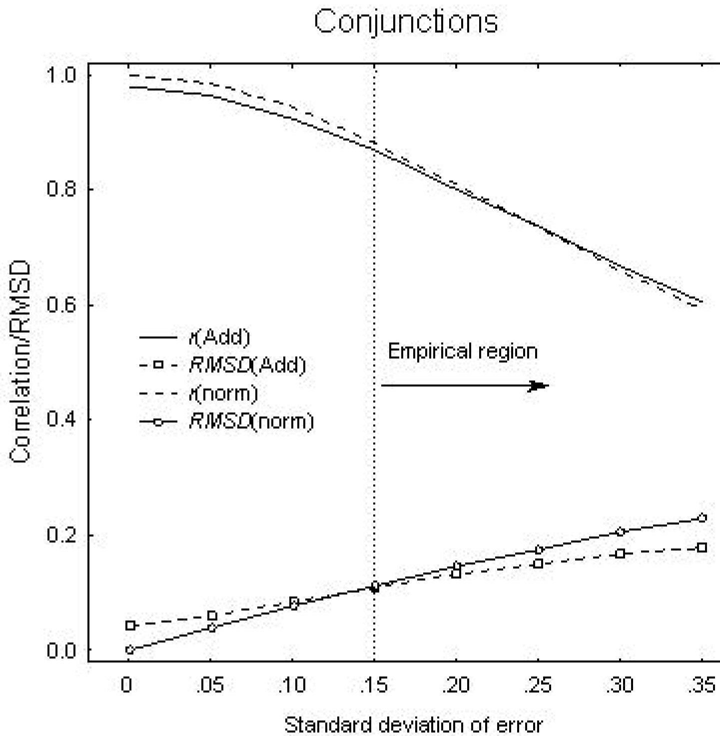

Figure 3. Correlations and deviations, in terms of RMSD, between true conjunction probabilities and noisy conjunction probabilities for the normative rule, norm, and the weighted average rule, add, separately. The x-axis presents the size of the added noise. The empirical region is based on data from Nilsson & Winman (2007). Data are taken from Winman et al., (2008).

Figure 3 shows how noise affects the deviations and correlations between original conjunction probabilities and the conjunction probabilities generated by the additive weighted average rule (i.e., *add*) and the multiplicative normative rule (i.e., *norm*). When the level of noise is low, *norm* deviates less and correlates stronger with the original conjunction probabilities. This is hardly surprising given that when noise equals 0, Equation 8 and Equation 9 are identical. The interesting aspect is what happens when the level of noise increases. As noise increases, the relative advantage of the normative rule is turned into a relative advantage of the weighted average rule. Thus, in a noisy environment it is actually better to use the weighted average rule than it is to use the normative rule.

All environments are more or less noisy. Therefore, the most effective combination-rule will differ (for a descriptive purpose, the noise from Nilsson & Winman, 2007, is represented by the dotted vertical line in Figure 3). However, the important feature of Figure 3 is the rather limited difference between the two combination rules throughout the entire range. This shows that, if the human mind is incapable of using the multiplicative normative rule, then the weighted average rule provides an excellent additive substitute.

In sum, if the human mind is predisposed to addition (as suggested by e.g., Juslin et al., 2003, 2008; Karlsson et al., in press), the (additive) weighted average rule could serve as an ecologically well suited alternative to the (multiplicative) normative combination-rule. Considering this, together with some empirical support, my suggestion is that the conjunction fallacy is caused by people assessing the probability of the conjunction by taking a weighted average between the component probabilities. A set of studies exploring this hypothesis are currently being performed in our lab. So far, the results from these studies are highly promising.

6.5. The Ecological Relevance of the Tasks

Behavior is constructed in the interaction between the environment and the organism. If the experimental environment is not ecologically relevant, behavior observed in the lab might deviate from behavior in the world to which the results are generalized (Dhami et al., 2004). Therefore it is important to construct experiments that capture situations that actually occur in the "real world" (e.g., Brunswik, 1955).

All studies included in this thesis used tasks where a limited amount of stimuli are repeatedly presented. In the learning phase of all experiments, participants received outcome feedback after every categorization, they could consider the feedback for as long as they needed, and as they were not penalized for making wrong decisions, they could freely explore the validity of their hypotheses. Obviously, this is highly uncommon in the every day life of the human being (Einhorn & Hogarth, 1978; Elwin, Juslin, Olsson, & Enkvist, 2007). As a result, the ecological relevance of the experiments could be questioned. It could be that if tasks more common in every day life had been used, evidence for usage of the representativeness heuristic would have been found.

In defense of the studies I-III, I would like to point to the following. First, to enable a study of how people represent and use knowledge, it is important to know which information has been available. Since it is quite impractical to follow participants from birth up to the time of the study, total control over what the participants can and can not know is only possible in these kinds of controlled experimental settings. Therefore, the type of tasks used in the present thesis is a necessity for an evaluation of the assumptions concerning the cognitive aspects of the representativeness heuristic.

Second, Study III replicated a phenomenon that has been found also in studies using more ecologically relevant tasks (e.g., Nilsson & Winman, 2007). Of all designs used in this thesis, the design of Study III was probably the most artificial. Each of the four distinct two-featured stimuli was repeated 30 times to conjunction learners and 60 times to component learners. Despite this, the conjunction fallacy was observed. And even more importantly, as was illustrated by for example the replication of the findings in Gavanski and Roskos-Ewoldsen (1991), the conjunction fallacies observed in Study III appeared to have the same characteristics as the conjunction fallacies observed in more ecologically normal tasks. My conclusion has to be that though the designs were necessarily restricted in an ecological sense, the results obtained have great relevance also for situations outside the laboratory.

6.6. Final Remarks

The goal of this thesis was to explore one of the major theories on how subjective probabilities are assessed, namely the idea that subjective probabilities are often assessed using the representativeness heuristic (Kahneman & Tversky, 1972). Two aspects were targeted. In studies I-II it was explored to what extent subjective probabilities can be attributed to the cognitive mechanisms suggested to be employed by the representativeness heuristic. Study III focused on whether it was likely that the conjunction fallacy, one of its most tightly associated biases, is caused by the representativeness heuristic. In neither of the studies were there any signs that the representativeness heuristic had been frequently used. Though the search for prototypes and feminist bank-tellers will certainly continue, the results of this thesis suggest that ultimately this search may prove futile.

Acknowledgement

First of all, I would like to thank my main supervisor, Professor Peter Juslin, and my assistant supervisor, PhD. Henrik Olsson. Peter you truly excel in the art of tutoring. Whatever the problem, you have the answer. With you as a supervisor, completing a thesis is a walk in the park. Henrik, I do not know exactly what an assistant supervisor is supposed to do. But, I am certain that you have done much more. Peter and Henrik, without you this thesis would never had been completed. Thank you for everything!

I would like to thank the distinguished members of the "Juslin-group": Ebba Elwin, Tommy Enkvist, Göran Hansson, Patrik Hansson, Maria Henriksson, Linnea Karlsson, Anna-Carin Olsson, and Anders Winman. You have improved the quality of both my work and my everyday life.

This work has benefited from insightful comments by PhD. Anders Winman at the Department of Psychology at Uppsala University, PhD. Patric Andersson at the Stockholm School of Economics, Professor Gunilla Bohlin at the Department of Psychology at Uppsala University, Professor Ralph Hertwig and the member of his research-group at the University of Basel, and the members of the ABC-Group in Berlin. I am also grateful to Edward Long at the Department of English at Uppsala University for improving the readability of this thesis.

Finally, I would like to thank my family. Anne: thank you! Filip: thank you for reminding me everyday about the important things in life. Margareta, Leif, Johan, and Elsie: thank you for all your love and support.

Uppsala, February 2008
Håkan Nilsson

References

Abelson, R. P., Leddo, J., & Gross, P. H. (1987). The strength of conjunctive explanations. *Personality and Social Psychology Bulletin*, *13*, 141-155.

Adam, M. B., & Reyna, V. F. (2005). Coherence and correspondence criteria for rationality: Experts estimation of risk of sexually transmitted infections. *Journal of Behavioral Decision Making, 18*, 169-186.

Baratgin, J. (2002). Is the mind definitely not Bayesian? A review of the various arguments. *Cahiers de Psychologie Cognitive/ Current Psychology of Cognition, 21*, 653-680.

Bearden, J. N., & Wallsten, T. S. (2004). MINERVA-DM and subadditive frequency judgments. *Journal of Behavioral Decision Making*, *17*, 349-363.

Becker, G. M., & McClintock, C. G. (1967). Value: Behavioral decision theory. *Annual Review of Psychology, 18*, 239-286.

Beyth-Marom R. (1981). The subjective probability of conjunctions. *Decision Research Report No 81-12,* Eugene, Oregon: Decision Research.

Björkman, M. (1994). Internal cue theory: Calibration and resolution of confidence in general knowledge. *Organizational Behavior and Human Decision Processes, 25*, 386-405.

Bonini, N., Tentori, K., & Osherson, D. (2004). A different conjunction fallacy. *Mind and Language, 19*, 199-210.

Brunswik, E. (1952). *The conceptual framework of psychology*. Chicago: University of Chicago Press.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review, 2*, 193-217.

Brunswik, E. (1956). *Perception and the representative design of psychological experiments* (2nd ed.). Berkeley: University of California Press.

Cobos, P. L., Almaraz, J., & García-Madruga, J. A. (2003). An associative framework for probability judgment: An application to biases. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 80-96.

Cohen, J., & Hansel, C. E. M. (1957). The nature of decision in gambling. *Acta Psychologica, 13*, 357-370.

Crandall, C. S., & Greenfield, B. (1986). Understanding the conjunction fallacy: A conjunction of effects. *Social Cognition, 4*, 408-419.

Daelemans, W. (1995). Memory-based lexical accusation and processing. In P. Steffens (Ed.), *Machine Translation and the Lexicon*, (Springer Lecture Notes in Artificial Intelligence No. 989, pp. 85-98). Berlin: Springer.

Davidson, D. (1995). The representativeness heuristic and the conjunction fallacy effect in children's decision making. *Merrill-Palmer Quarterly, 41*, 328-346.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin, 130*, 959-988.
Dougherty, M. R. P. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General, 130*, 579-599.
Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999) MINERVA-DM: A memory process model for judgments of likelihood. *Psychological Review, 106*, 180-209.
Dulany, D. E., & Hilton, D. J. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition, 9*, 85-110.
Einhorn, H. J. (1985). *A model of the conjunction fallacy*. Chicago: University of Chicago, Center for Decision Research.
Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of illusion of validity. *Psychological Review, 85*, 395-416.
Elwin, E., Juslin, P., Olsson, H., & Enkvist, T. (2007). Constructivist encoding: Learning from selective feedback. *Psychological Science, 18*, 105-110.
Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review, 83,* 37-64.
Evans, J. S. B. T., & Pollard, P. (1982). Statistical judgment: A further test of the representativeness construct. *Acta Psychologica, 51*, 91-103.
Fantino, E., Kulik, J., Stolarz-Fantino, S., & Wright, W. (1997). The conjunction fallacy: A test of the averaging hypotheses. *Psychonomic Bulletin & Review, 4*, 96-101.
Fischoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance, 1*, 288-299.
Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research: Psychologische Forschung, 50*, 123-129.
Fisk, J. E. (2002). Judgments under uncertainty: Representativeness or potential surprise? *British Journal of Psychology, 93*, 431-449.
Fisk, J. E., & Pidgeon, N. (1996). Component probabilities and the conjunction fallacy: Resolving signed summation and the low component model in a contingent approach. *Acta Psychologica, 94*, 1-20.
Ford, J. K., Schmitt, N., Schechtman, S. L., Hults, B. M., & Doherty, M. L. (1989). Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes, 43,* 75-117.
Garcia-Retamero, R., Takezawa, M., & Gigerenzer, G. (2007). *How to learn good cue orders: When social learning benefits fast and frugal heuristics.* Paper presented at the biannual meeting of the European Association for Decision Making (SPUDM), Warsaw, PL.
Gavanski, I., & Hui, C. (1992). Natural sample spaces and uncertain belief. *Journal of Personality and Social Psychology*, *63*, 766-780.
Gavanski, I., & Roskos-Ewoldsen, D. R. (1991). Representativeness and conjoint probability. *Journal of Personality and Social Psychology, 61*, 181-194.
Gigerenzer, G. (1991). How to make cognitive illusion disappear: Beyond "heuristics and biases". *European Review of Social Psychology, 2*, 83-115.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103*, 592-596.

Gigerenzer, G. (2005). I think, therefore I err. *Social Research, 72*, 1-24.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103*, 650-669.

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

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review, 98*, 506-528.

Gigerenzer, G., & Murray, D. J. (1987). *Cognition as intuitive statistics.* Hillsdale: Lawrence Erlbaum.

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

Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart.* New York: Oxford University Press.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.) (2002). *Heuristics and biases: The psychology of intuitive judgment.* New York: Cambridge University Press.

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

Gluck, M. A., & Bower, G. H. (1988b). Evaluating an adaptive network model of human learning. *Journal of Memory and Language, 27*, 166-195.

Goldstein, W. M., & Hogarth, R., M. (Eds.) (1997). *Research on judgment and decision making: Currents, connections, and controversies.* New York: Cambridge University Press.

Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology, 24*, 411-435.

Hacking, I. (1975). *The emergence of probability: A philosophical study of early ideas about probability, induction and statistical inference.* London: Cambridge University Press.

Helversen, B. von, & Rieskamp, J. (in press). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General.*

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006). The role of information sampling in risky choice. In K Fiedler & P. Juslin (Eds.). *Information sampling and adaptive cognition* (pp. 72-91). New York: Cambridge University Press.

Hertwig, R., & Chase, V. M. (1998). Many reasons or just one: How response mode affects reasoning in the conjunction fallacy. *Thinking and Reasoning, 4*, 319-352.

Hertwig, R., & Gigerenzer, G. (1999). The conjunction fallacy revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making, 12*, 275-305.

Hertwig, R., Pachur, T., & Kurzenhäuser, S. (2005). Judgments of risk frequencies: Tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 621-642.

Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review, 95*, 528-551.
Homa, D., Cross, J., Cornell, D., Goldman, D., & Shwartz, S. (1973). Prototype abstraction and classification of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology*, *101*, 116-122.
Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory, 5*, 11-23.
Juslin, P. (1994). The overconfidence phenomenon as a consequence of informal experimenter-guided selection of almanac items. *Organizational Behavior and Human Decision Processes*, *57*, 226-246.
Juslin, P. Karlsson, L., & Olsson, H. (2008). Information integration in multiple-cue judgment: A division of labour hypothesis. *Cognition, 106,* 259-298.
Juslin, P, Olsson, H., & Björkman, M. (1997). Brunswikian and Thurstonian origins of bias in probability assessment: On the interpretation of stochastic components of judgment. *Journal of Behavioral Decision Making, 10*, 189-209.
Juslin, P., Olsson, H., & Olsson, A-C. (2003). Exemplar effects in categorization and multiple judgment. *Journal of Experimental Psychology: General, 132*, 133-156.
Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science, 26,* 563-607.
Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review, 114*, 678-703.
Juslin, P., Winman, A., & Olsson, H. (2000). Naïve empiricism and dogmatism in confidence research: A critical examination of the hard-easy effect. *Psychological Review, 107*, 384-396.
Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.). *Heuristics and biases: The psychology of intuitive judgment* (pp. 49-81). New York: Cambridge University Press.
Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristic and biases*. New York: Cambridge University Press.
Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology, 3*, 430-454.
Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80*, 237-251.
Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 103, 582-591.
Karlsson, L., Juslin, P., & Olsson, H. (in press). Exemplar-based inference in multi-attribute decision making: Contingent, not automatic, strategy shifts? *Judgment and Decision Making.*
Koehler, D. J., White, C. M., & Grondin, R. (2003). An evidential support accumulation model of subjective probability. *Cognitive Psychology, 46*, 152-197.

Kruschke, J. K. (1992). An exemplar-based connectionist model of category learning.
Psychological Review, 99, 22-44.
Lagnado, D. A., & Shanks, D. R. (2002). Probability judgment in hierarchical learning: a conflict between predictiveness and coherence. *Cognition, 83*, 81-112.
Lichtenstein, S., Fischoff, B., & Philips, L. (1982). Calibration and probabilities: The state of the art to 1980. In *Judgment under uncertainty: Heuristics and Biases* (pp. 206-334), D. Kahneman, P. Slovic, & A. Tversky (Eds.). Cambridge, UK: Cambridge University Press.
Logan, G. D. (1988). Towards an instance theory of automatization. *Psychological Review, 95*, 492-527.
Logan, G. D. (2002). An instance theory of atention and memory. *Psychological Review, 109*, 376-400.
Love, B. C., & Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of categgory learning. *Psychological Review, 111*, 309-332.
Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin, 129*, 592-613.
Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207-238.
Mervis, C. B., & Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology, 32*, 89-115.
Minda, P. J., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 775-799.
Murphy, G. L. (2004). *The big book of concepts*. Cambridge: MIT Press.
Myung, I. J., & Pitt, M. A.(1997). Applying Occam´s razor in modeling cognition: A bayesian approach. *Psychonomic Bulletin & Review, 4*, 79-95.
Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). *Straight choices: The psychology of decision making.* New York: University Press.
Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest: Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 82-96.
Nilsson, H., Juslin, P., & Olsson, H. (2003). From prototypes to exemplars: Representational shifts in a probability judgment task. In R. Alterman & D. Kirsh (Eds.), *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*.
Nilsson, H. & Winman, A. (2007). *A study of the conjunction fallacy using an ecologically representative design where the representativeness heuristic can not be applied*. Paper presented at the biannual meeting of the European Association for Decision Making (SPUDM), Warsaw, PL.
Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 104-114.
Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.
Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review, 7*, 375-402.

Nosofsky, R. M., & Zaki, S. R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory and, Cognition, 28,* 924-940.
Olsson, H., Wennerholm, P., & Lyxzén, U. (2004). Exemplars, prototypes, and the flexibility of classification models. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
Pavlov, I. P. (1927). *Conditioned reflexes*. New York: Oxford University Press.
Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin, 68*, 29-46.
Pitt, M. A., Myung, I. J., & Zhang, S. B. (2002). Towards a method of selecting among computational models of cognition. *Psychological Review, 109*, 472-491.
Posner, I. M., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 353-363.
Posner, I. M., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology, 83*, 304-308.
Rakow, T., Hinvest, N., Jackson, E, & Palmer, M. (2004). Simple heuristics from the adaptive toolbox: Can we perform the requisite learning. *Thinking and Reasoning, 10*, 1-29.
Rakow, T., Newell, B. R., Fayers, K., & Hersby, M. (2005). Evaluating three criteria for establishing cue-search hierarchies in inferential judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 1088-1104.
Salmon, W. C. (1967) The *foundations of scientific inference*. Pittsburgh: University of Pittsburgh Press.
Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
Schmitt, M., & Martignon, L. (2006). On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research, 7*, 55-83.
Shanteau, J. (1974). Averaging versus multiplying combination-rules of inference judgment. *Acta Psychologica, 39*, 83-89.
Shanteau, J. (1989). Cognitive heuristics and biases in behavioral auditing: Review, comments, and observations. *Accounting Organizations and Society, 14*, 165-177.
Shin, H. J., & Nosofsky, R. M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121*, 278-304.
Sieck, W. R. (2003). Effects of choice and relative frequency elicitation on overconfidence: Further test of an exemplar-retrieval model. *Journal of Behavioral Decision Making*, *16*, 127-145.
Sieck, W. R., & Yates, J. F. (2001). Overconfidence effects in category learning: A comparison of connectionist and exemplar memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*, 1003-1021.
Skirbekk, G., & Gilje, N. (1995). *Filosofins historia* [The history of philosophy]. Stockholm: Daidalos.
Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirming strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology, 22*, 93-121.
Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99*, 3-21.

Smith, J. D., & Minda, J. P. (1998). Prototype in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1411-1436.
Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 3-27.
Smith, J. D., Murray, M. J., & Minda, J. P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology: learning Memory, and, Cognition, 23*, 659-680.
Solomon, K. O., Medin, D. L., & Lynch, E. (1999). Concepts do more than categorize. *Trends in Cognitive Sciences, 3*, 99-105.
Stanovich, K. E., & West, R. F. (2002). Individual differences in reasoning: Implications fort he rationality debate? In T. Gilovich, D. Griffin, & D. Kahneman (Eds.). *Heuristics and biases: The psychology of intuitive judgment* (pp. 421-440). New York: Cambridge University Press.
Stanton, R. D., Nosofsky, R. M., & Zaki, S. R. (2002). Comparisons between exemplar similarity and mixed prototype models using a linearly separable category structure. *Memory & Cognition, 30*, 934-944.
Stolarz-Fantino, S., Fantino, E., Zizzo, D. J., & Wen, J. (2003). The conjunction effect: New evidence for robustness. *American Journal of Psychology, 116*, 15-34.
Strange, W., Keeney, T., Kessel, F. S., & Jenkins, J. J. (1970). Abstraction over time of prototypes of distortion of random dot patterns: A replication. *Journal of Experimental Psychology, 83*, 508-510.
Tentori, K., Bonini, N., & Osherson, D. (2004). The conjunction fallacy: A misunderstanding about conjunction? *Cognitive Science, 28*, 467-477.
Thüring, M., & Jungermann, H. (1990). The conjunction fallacy: Causality vs. event probability. *Journal of Behavioral Decision Making, 3*, 61-74.
Timmermans, D. (1993). The impact of task complexity on information use in multiattribute decision-making. *Journal of Behavioral Decision Making, 6*, 95-111.
Trope, Y., & Bassok, M. (1983). Information-gathering strategies in hypothesis-testing. *Journal of Experimental Social Psychology, 19*, 560-576.
Tversky, A. (1977). Features of similarity. *Psychological Review, 85*, 327-352.
Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207-232.
Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*, 1124-1131.
Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*, 293-315.
Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextentional representation of subjective probability. *Psychological Review*, *101*, 547-567.
Villejoubert, G., & Mandel, D. R. (2002). The inverse fallacy: An account for deviations from bayes theorem and the additivity principle. *Memory and Cognition, 30*, 171-178.
Von Neuman, J., & Morgenstern, O. (1947). *Theory of games and economic behaviour*. Princeton: Princeton University Press.

Wallsten, T. S. (1983). The theoretical status of judgmental heuristics. In R. W. Scholz (Ed.), *Decision making under uncertainty (pp.* 21-39). Amsterdam: Elsevier.
Watson, J. B., & Rayner, R. (1920). Conditioned emotional reactions. *Journal of Experimental Psychology, 3*, 1-14.
Wedell, D. H., & Moro, R. (in press). Testing boundary conditions for the conjunction fallacy: Effects of response mode, conceptual focus, and problem type. *Cognition*.
Wells, G. L. (1985). The conjunction error and the representativeness heuristic. *Social Cognition, 3*, 266-279.
Winman, A. (1997). The importance of item selection in "knew-it-all-along" studies of general knowledge. *Scandinavian Journal of Psychology, 38*, 63-72.
Winman, A., Nilsson, H., Juslin, P., & Hansson, G. (2008). [Simulation of the ecological rational of the weighted average hypothesis]. Unpublished raw data.
Wyer, R. S., Jr. (1976). An investigation of the relations among probability estimates. *Organizational Behavior & Human performance, 15*, 1-18.
Zacker, R. T., & Hasher, L. (2002). Frequency processing: A twenty-five year perspective. In P. Sedelmeier & T. Betsch (Eds.). *Etc. frequency processing and, Cognition* (pp. 21-36). London: Oxford University Press.
Zizzo, D. J. (2003). Verbal and behavioral learning in a probability compounding task. *Theory and Decision, 54*, 287-314.
Zizzo, D. J., Stolarz-Fantino, S., Wen, J., & Fantino, E. (2000). A violation of the monotonicity axiom: Experimental evidence on the conjunction fallacy. *Journal of economic behavior and organization, 41*, 263-276.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 37

Editor: The Dean of the Faculty of Social Sciences

A doctoral dissertation from the Faculty of Social Sciences, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-8476

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2008