UPPSALA
UNIVERSITET

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 51

Learning With Selective Feedback

Effects on Performance and Coding of Unknown Outcomes

EBBA ELWIN

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2009

ISSN 1652-9030
ISBN 978-91-554-7568-0
urn:nbn:se:uu:diva-106880

Dissertation presented at Uppsala University to be publicly examined in Auditoriet, Gustavianum, Akademigatan 73, Uppsala, Saturday, September 19, 2009 at 13:00 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract
Elwin, E. 2009. Learning With Selective Feedback. Effects on Performance and Coding of Unknown Outcomes. Acta Universitatis Upsaliensis. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 51. Uppsala.*
ISBN 978-91-554-7568-0.

In experiential learning, one important source of information is the feedback that people receive on the outcomes of their decisions. Typically, however, feedback is systematically absent for many decisions and the actual experience of people may therefore be highly selective. It is thus surprising that research on the cognitive processes involved in human judgement and categorisation has not addressed the effects of learning with selective feedback. In this thesis, three studies are presented in which the effects of learning with systematically selective feedback were investigated.

The hypothesis of *constructivist coding* was introduced in Study I, suggesting a cognitive mechanism for the processing of selective information. In the absence of external feedback people infer the most likely outcome, and then code this inference into memory as “internal feedback”. This internally generated feedback is stored and processed in the same manner as externally presented feedback and is used as a basis for beliefs about the characteristics of the environment. Results from Studies I, II, and III demonstrated support for constructivist coding under varied learning conditions.

Study III investigated the effects on the beliefs of participants when they learn from feedback received only for positive decisions. Results indicated that the participants’ beliefs well reflected their actual, however selective, experience. When participants aimed to achieve good immediate outcomes, their experience became restrictive and biased, resulting in biased beliefs. In contrast, when the focus of participants was on long-term learning, their decisions produced a more representative experience and their beliefs came to reflect the actual structure of the environment. Biased beliefs were thus demonstrated to result from a sensitivity of participants to selectively available information.

The present findings offer an understanding of the cognitive processes involved in learning from selectively absent feedback. The conclusions propose a sensitivity of participants to objectively experienced information in the forming of knowledge and beliefs. Further, when external information is absent, participants appear to rely on their knowledge and expectations to infer and code the most likely outcome, and use these stored inferences to form a coherent representation of the environment.

Ebba Elwin, Department of Psychology, Box 1225, Uppsala University, SE-75142 Uppsala, Sweden

ISSN 1652-9030
ISBN 978-91-554-7568-0
urn:nbn:se:uu:diva-106880 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-106880)

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals. Reprints were made with permission from the respective publishers.

I Elwin, E., Juslin, P., Olsson, H., & Enkvist, T. (2007). Constructivist coding: Learning from selective feedback. *Psychological Science, 18*, 105-110.

II Henriksson, M. P., Elwin, E., & Juslin, P. (2009). What is coded into memory in the absence of outcome feedback? Manuscript submitted for publication.

III Elwin, E. (2009). Living and learning: The interplay between beliefs, sampling behaviour, and experience. Manuscript submitted for publication.

Contents

Introduction

Anyone who is interested in the human mind and in the general ability of people to learn and make good decisions, will find in cognitive science – like in everyday life – contradictory views of human cognitive capacity. On the one hand, much research on learning, categorisation, and judgement report good performance on the part of participants. With repeated exposure to objects in a new and unfamiliar task, people are typically highly competent in learning to make correct judgements, predictions, or categorisations (e.g., Ashby & Maddox, 2005; Minda & Smith, 2001; Juslin, Jones, Olsson, & Winman, 2003). Similarly, everyday life provides frequent demonstrations of people's astonishing ability to process large amounts of information and make decisions that involve highly complex information. Consider, for example, the skills required when driving a car in heavy traffic.

On the other hand, much research on intuitive judgement and decision making emphasises the role of heuristic processing of information, resulting in a range of judgement and decision biases. Even extensive learning and expertise have been found to be associated with poor decisions made with a level of confidence significantly above the actual accuracy. For example, a review of studies on professional judgements within the mental health field demonstrated that experienced clinicians rarely achieved more accurate clinical predictions than novices (Garb, 1989). Indeed, research findings of inadequate judgements in many areas are not surprising considering the prevalence of discrimination, prejudice, superstition, and other unjustified beliefs that people entertain and encounter in their everyday lives.

These conflicting descriptions of the achievement of the human mind provide an incongruous picture of people's ability to learn and make good decisions. In the attempt to understand the apparent contradiction, this thesis will begin with a closer examination of methodology.

Experimental Control vs. Ecological Validity

One important aim of cognitive research is to understand and describe the mental processes that underlie learning from experience. Models that describe such processes are developed from data collected in well-controlled experimental studies. Research on the cognitive processes in categorisation learning or judgements needs to reduce the everyday learning or decision situation to its essential components in order to identify the important aspects of learning or decision making for the behaviour under study. In a complex "real-world" decision situation, it would be impossible to disentangle the history of learning from the characteristics of the decision situation when determining the unique influence of various factors on decision behaviour and performance. Much research on categorisation learning or judgements thus involves laboratory studies that typically use tasks with a large number of categorisations or judgements of outcomes made by the participant based on a few characteristics of each object. Each judgement is followed by feedback on the correct answer (e.g., Ashby & Maddox, 2005; Kruschke, 2005; Minda & Smith, 2001; Nosofsky, 1986).

Obviously, the aim of such experimental research is to explain what occurs in a "natural" environment. The benefits of reliability of results from the establishing of control over possibly relevant factors, however, come with the risk of reducing ecological validity. Research performed in the laboratory needs to be generalisable to the natural situations that we want to understand (Brunswik, 1943; Dhami, Hertwig, & Hoffrage, 2004; Hammond, 1955, 1996; Hogarth, 1981). Typically, in everyday experiential learning, information is not systematically collected, information may be vague and ambiguous, some information may be lacking or ignored, and other information inferred. Such aspects of everyday experiential learning make it quite more complex than the well-controlled laboratory task that involves repeated exposure to objects similar in all but some aspects and where every decision is followed by the correct answer (Johnson & Bruce, 2001).

Availability of Feedback

The contrast between research findings of high performance in the laboratory as compared with inadequate expert decision making could possibly be explained by some consistent aspect of the tasks that faces the decision makers in these different environments. One such factor is outcome feedback, i.e. the information received after a decision has been made. As noted, in laboratory learning tasks, feedback is typically received after each decision made by the participant. In contrast, everyday life offers little outcome feedback (Hogarth, 2006). For everyday decisions, whether personal or professional, people may often be informed of the outcome only for the accepted alternatives, rarely receiving feedback for rejected options. When the decision is to buy a product or engage in an activity, the outcome will be known and may be evaluated, whereas the outcomes of rejected alternatives will not be experienced. Learning from such systematically selective feedback may be an explanation of poor performance and biased decisions.

In fact, the biased and poor decisions often seen among experts and in everyday judgements have been suggested to be the result of the amount and quality of available feedback (e.g., Brehmer, 1980; Einhorn & Hogarth, 1978; Dawes, 1994). Surprisingly, however, the effects of learning from incomplete feedback have not received much attention in cognitive research, and the hypothesis that poor performance and bias are the result of learning from selective feedback has rarely been empirically investigated.

When people learn from *selective feedback*, their beliefs are updated based on selective experience. Outcome feedback on decisions is systematically absent on some occasions, such as for specific decisions or for certain outcomes. In the machine learning area, one such structure of systematically selective feedback is investigated in reinforcement learning networks, in which feedback is received only for positive decisions (e.g., Sutton & Barto, 1998). For example, an employer receives information concerning the results of hired job applicants, but no information on the potential of rejected applicants. Learning from experience without any feedback has been investigated in areas such as language acquisition by means of statistical learning (e.g., Saffran, Aslin, & Newport, 1996) and unsupervised learning (e.g., Barlow, 1989; Vallabha, McClelland, Pons, Werker, & Amano, 2007). Without feedback information, knowledge is extracted from regularities and organisation in the observed features. In contrast, selective feedback provides some feedback information that may be used in learning, and when feedback is systematically absent, this information is likely to be skewed.

A few experimental studies have compared the effects of learning from selective feedback received only for positive decisions, with learning from complete outcome feedback. Results indicate that the effects of selective and biased feedback on the level of performance may actually be quite small (Fazio, Eiser, & Shook, 2004; Griffiths & Newell, 2009; Grosskopf, Erev, & Yechiam, 2006; Yechiam & Busemeyer, 2006). In fact, it has been noted that (complete) outcome feedback actually does not always benefit learning (e.g., Hammond, Summers & Deane, 1973; Jacoby, Mazursky, Troutman, & Kuss, 1984; Kluger & DeNisi, 1996; Remus, O'Connor, & Griggs, 1996).

Not only are the effects of selective feedback on performance unclear, but questions concerning the coding of experiences with systematically selective feedback have been highly neglected in cognitive research. There have been no efforts to identify the processes that underlie learning from systematically absent feedback information and the resulting representations in memory. As shown later in this thesis, however, knowledge from various areas of cognitive research may offer interesting contributions to our understanding of what people learn from their experiences when feedback is systematically absent for some outcomes.

Outline and Aims of the Thesis

The subject of this thesis is the role of selective feedback in learning and decision making. After the presentation of influential contemporary perspectives on judgement and decision making, a brief account will be given of research on the cognitive processes involved in judgement and categorisation learning. Most research on the processes that underlie judgement and learning from experience involves the study of general cognitive processes which are (implicitly) assumed quite similar across individuals. Likewise, the research presented in this thesis aims to investigate and understand the cognitive processes common to most people, i.e. differences between individuals are not predicted or addressed. The relevant processes may possibly vary with cultural conditions, as well as with more or less stable individual differences, in terms of intellectual ability or personality. Today, however, research in this area does not enable substantiated predictions as to differences linked to the characteristics of individuals regarding the specific cognitive processes focused on in this thesis (see e.g., Weber & Johnson, 2009, for an overview of individual differences in decision making and Weber & Hsee, 2000, for cultural influences).

Further, an overview of research on learning from feedback will be given and specifically the effects of learning with selective feedback on decisions and performance. In addition, the mental processes that may deal with incomplete information will be outlined. Three studies will be presented in this thesis in which the aims were the study of (a) the effects of learning with selective feedback on performance, (b) the coding in memory of selective feedback, and (c) selective feedback as a cause of decision bias. The presentation will be completed with a discussion of the results in relation to other research, as well as to everyday learning and decision making.

Terminology

Before the presentation of research areas relevant for the issues addressed in this thesis, some central concepts will be outlined. The terms *judgement* and *decision* are often used in an undifferentiated manner. However, Harvey (2001) attempted at a defining distinction by describing decisions as judgements with consequences. Judgements concern inferences or predictions of an outcome that may be more or less accurate. Judgements thus involve the beliefs or knowledge of the individual and the processes by which such knowledge is represented, integrated and used. A decision is associated with action and may be assessed on the basis of its consequences. Thus, decision-making research not only involves the beliefs of the individual about the situation and the likely outcome but may also consider desires, goals, and preferences.

A related research area concerns *categorisations*, or classifications, which may be defined as involving judgements on a nominal scale (Juslin, Olsson, & Olsson, 2003). In this thesis, categorisations of objects into one of two categories will therefore often be referred to as binary judgements. Further, no clear distinction will be made between judgements, decisions, or choices, and any of the terms may be used to denote the response of an individual assumed to express a belief about the probable outcome based on available information. Notably, in any decision situation, the expectations regarding the consequences of a response will most likely also influence decisions. However, the research presented in this thesis does not seek to investigate the aspects of the decision situation specifically associated with emotions or preferences; instead, the focus lies on the beliefs and the knowledge of the individual.

In this context, emotions may be regarded as one source of information upon which a decision is based (Schwarz, 2002). Accordingly, the feelings of the individual concerning the qualities of a specific product may be one feature among others (such as the price and practical aspects) that affect the decision of whether to purchase the product. Further, the outcome criterion on which the decision is evaluated by the individual may be of emotional quality, such as the enjoyment of a chosen product (for a more general presentation of the role of emotions in decision making, see, e.g., Bechara, Damasio, & Damasio, 2000; Weber & Johnson, 2009).

Decisions and judgements are likely to be based, at least in part, upon knowledge and beliefs acquired in *learning*. Without aiming at a general definition of the term, the research presented in this thesis involves learning as a process in which the knowledge and beliefs of the individual develop with experience in a specific environment. Improved performance as a result of learning may be indicated by a significantly increased proportion of correct decisions, in which case the knowledge of the individual is evaluated against a normative criterion defined by the actual structure of the environment. In the Method section of this thesis, several measures of different aspects of knowledge will be outlined.

Perspectives on Judgement and Decision Making

The subject of intuitive judgement and decision making concerns everyday decisions made without extensive reasoning or decision tools. The research area has been the focus of considerable attention and debate. Many studies have demonstrated suboptimal decisions under various circumstances, even in areas in which people have extensive experience and expertise (e.g., Dawes, Faust, & Meehl, 1989). The demonstrated poor decision making and overconfidence, however, have been met with different explanations. Moreover, important objections have been raised concerning what is perceived as an unjustified emphasis on cognitive limitations and defective decisions. Today, three perspectives on human judgement and decision making stand out and influence much research. The most senior of these perspectives is the *heuristics and biases* programme initiated by Amos Tversky and Daniel Kahneman in the 1970s. This perspective has achieved widespread acknowledgement of the limitations and fallibility of human judgement. In contrast, Gerd Gigerenzer and colleagues emphasise the *ecological rationality* of heuristics in decision making and the capacities of the human mind. Recently, Klaus Fiedler introduced a "sampling approach" (Fiedler, 2000), later recast by him and Peter Juslin in the metaphor of the *naive intuitive statistician* (Fiedler & Juslin, 2006). This view focuses on the interaction between the individual and the environment and the significance of individual experience in knowledge and judgements. Each of these perspectives is briefly presented in the following sections.

Heuristics and Biases

Research within the heuristics and biases approach has identified a number of heuristics or "rules of thumb" that may facilitate decision making, but that may also lead to systematic errors or biases. For example, the availability heuristic is used to describe how people judge the probability of an event based on the ease with which examples of the event are retrieved from memory (Tversky & Kahneman, 1973). Availability in memory often provides a good indication as to the relative frequency or probability of an event; however, bias may appear when the correlation between availability and probability is low. For example, participants have been demonstrated to judge the probability that a randomly sampled English word starts with an R as larger than the probability that R is the third letter, even though the reverse is true. This is understood in terms of availability or the ease with which instances of the categories come to mind, where it is easier to think of words that begin with a certain letter than words with a certain letter in the third position (Tversky & Kahneman, 1973).

More recent developments within this perspective present a broader conceptualisation on human judgement and decision making as governed by two systems of reasoning, where System 1 is intuitive and automatic and System 2 is more deliberate and controlled. The quick and intuitive judgements from System 1 are supervised - and may be corrected - by the effortful judgements of System 2. Bias is thus the result of the failure of System 1 to produce a correct answer and of System 2 to detect and correct the error (Kahneman & Frederick, 2002).

In the heuristics and biases approach, people's judgements are evaluated against normative standards of rationality and are shown to deviate from these standards in systematic ways. The normative role of statistics and probability theory as standards of rationality is accepted, although probability and statistics are concluded to be inadequate as descriptive models of human judgement and decision making. Instead, decisions in general, and biases in particular, are described as the products of the decision situation (e.g., the framing of questions) and processes involved at the decision stage (such as attention, retrieval, or affect; e.g., Gilovich, Griffin, & Kahneman, 2002; Slovic, Fischhoff, & Lichtenstein, 1977; Tversky & Kahneman, 1973, 1974). In this research tradition, little attention is devoted to the development of knowledge and bias with experience.

Ecological Rationality

A contrasting perspective on human decision making views people as "boundedly rational" and competent in making good use of limited time, knowledge, and cognitive capacity to make adaptive decisions (Gigerenzer, 2007; Gigerenzer & Goldstein, 1996; Simon, 1956). The notion of "ecological rationality" further develops this view and involves the study of mental heuristics for their benefits rather than their biases (e.g., Gigerenzer, Czerlinksi, & Martingnon, 2002). A suitable heuristic is chosen from a range of alternative decision strategies available in the "adaptive toolbox". These simple and robust algorithms in the form of ecologically rational heuristics are demonstrated to exploit the structure of a specific environment and improve decision quality as compared with more complex inference models. For instance, the recognition heuristic relies on people's automatic registering of frequencies of objects in the environment. When a criterion (such as importance or size) is correlated with frequency, recognised objects are adequately judged to have higher values on the criterion than unrecognised objects (Goldstein & Gigerenzer, 2002).

From the perspective of ecological rationality, the biases demonstrated in much research are regarded as artefacts of the specific properties of the tasks (such as presentation format) and inappropriate normative standards from probability theory and statistics (Gigerenzer & Goldstein, 1996; Todd & Gigerenzer, 2007). Instead, people are considered sensitive to the structure of the environment, which is assumed to be reflected in the experience of the individual.

The Naive Intuitive Statistician

In the 1960s, the notion of man as an "intuitive statistician" emphasised people's impressive ability to survive and prosper even with imperfect information in the environment (Peterson & Beach, 1967). According to this view, just like the scientist who relies on statistics and probability theory to make inferences from imperfect information, so do people in their everyday life though it is done intuitively. Several studies have compared human inferences and performance to normative standards in statistical theory, concluding that such formal models may serve as a good approximation for understanding human inferences (e.g., Cosmides & Tooby, 1996; Peterson & Beach, 1967; Peterson, Hammond, & Summers, 1965).

Recently, a modification of this view was introduced. The metaphor of the *naive intuitive statistician* acknowledges the ability of people to accurately describe the properties of the environment, but with meta-cognitive limitations (Fiedler, 2000; Fiedler & Juslin, 2006). The individual's capacity to process information is accompanied by an inability to evaluate the representativity of experienced information (Juslin, Winman, & Hansson, 2007). Decision makers are "intuitive statisticians" because of their ability to register basic distributional properties of the environment (such as frequencies) and accurately describe the properties of the experienced sample (Estes, 1976; Hasher & Zacks, 1979, Zacks & Hasher, 2002). However, they are "naive" in their inability to take into account the limitations of their individual experience and in relation to the more sophisticated statistical properties of the information. When the sample is skewed or unrepresentative, learning will be biased because people lack the meta-cognitive capacity to detect and correct such bias (Fiedler, 2000; Unkelbach, Fiedler, & Freytag, 2007). Judgements are based on information that has been encountered and stored in memory and bias is manifested to the extent that this information is biased.

The two approaches outlined in previous sections describe human judgement as a reliance on heuristic processes with the aim to simplify decisions: their focus is not on the specific experience of the individual. In contrast, the perspective of the naive intuitive statistician emphasises the importance of experience for the development of the knowledge and beliefs that underlie judgement. Learning from experience, often based on decisions with feedback, enables the development of knowledge in relation to an environment, which is why the understanding of judgement benefits from the study of the experience of the individual and the processes associated with learning (Barron & Erev, 2003; Hertwig, Barron, Weber, & Erev, 2004, 2006; Hogarth, 1981; Newell, Lagnado, & Shanks, 2007).

Cognitive Processes in Judgement and Categorisation Learning

A brief overview will be presented of research designed to investigate and describe the cognitive processes and memory representations associated with judgement and categorisation. The cognitive mechanisms involved in learning from experience when interacting with the environment have been the focus of much research on categorisation learning. In this research, learning is typically described as the storing of information provided in the features of categorised objects along with feedback on categorisations, allowing the coding of concrete examples, or exemplars (e.g., Nosofsky, 1986; Nosofsky & Johansen, 2000), or of category prototypes (e.g., Ashby & Maddox, 2005; Kruschke, 2005; Minda & Smith, 2001). The stored representations are retrieved in later categorisations when a new object is categorised based on the similarity to the relevant representations. For instance, a physician may diagnose (categorise) a new patient based on the similarity of the patient's symptoms (features) to those of previously encountered patients with known diagnoses (exemplars).

In recent decades, associative learning models have been developed describing categorisation learning as the updating of associations in a connectionist network. Connectionist models of supervised learning generally involve error-correcting mechanisms in which associations between features and outcomes are adjusted based on feedback (e.g., Gluck & Bower, 1988; Shanks, 1991). Models that combine error-driven learning mechanisms with exemplar-based representations depict learning as the association of exemplars to categories (e.g., Kruschke, 1992; Nosofsky, Kruschke, & McKinley, 1992).

As indicated in the previous overview of research paradigms in judgement and decision making, considerable research within this area has usually been concerned with the cognitive processes of the decision situation (e.g., availability or recognition) rather than the cognitive mechanisms involved in the underlying learning process. In this research, several accounts of the processes associated with judgements describe how information may be used in decision making by means of heuristic processing such as the take-the-best heuristic (Gigerenzer, Hoffrage, & Kleinbölting, 1991) or the mapping model (von Helversen & Rieskamp, 2008). These models, however, specify the processes involved in decisions without reference to how cognitive representations are formed (however, see Rieskamp & Otto, 2006, for a learning model for the *selection* of decision strategy, SSL).

The processes and representations associated with learning from judgements with feedback have been investigated in research on multiple-cue learning. In this research, learning is described as the abstraction of rules or relationships between characteristics (cues) of objects and their outcomes (e.g., Cooksey, 1996; Doherty & Brehmer, 1997; Einhorn, Kleinmuntz, & Kleinmuntz, 1979; Hammond & Summers, 1972; Juslin, Jones, et al., 2003). A physician may thus learn the rule-governed relationships between symptoms and diagnoses from her experience with patients. For a new patient, these rules are integrated in a diagnosis (judgement).

One issue that has been addressed in research on multiple-cue judgement is the insight of individuals into their own decision strategies. In general, such research has demonstrated little insight by the decision maker, which has been interpreted as the reliance on implicit and non-declarative processes (Evans, Clibbens, Cattani, Harris, & Dennis, 2003; Gluck, Shohamy, & Myers, 2002). Recent findings, however, suggest that these results may be an effect of the method of investigation rather than a reflection of the knowledge of the individuals. If asked in a proper manner, the participants may give accurate reports on their decision strategies (Lagnado, Newell, Kahan, & Shanks, 2006). The general procedure when investigating the processes that underlie judgements or categorisations, however, involves the analysis of the judgement behaviour of participants, rather than their introspective reports.

Today, a prevalent view of the processes involved in human judgement is that of two separate cognitive systems working in parallel. These systems are governed by different processes; often one is based on associative and implicit processes while the other is governed by rules (e.g., Evans, 2008; Hammond, Hamm, Grassia, & Pearson, 1997; Sloman, 1996). As indicated earlier in this thesis in relation to System 1 and System 2 processing (Kahneman & Frederick, 2002), the processes may have different advantages depending on the context and circumstances.

Specifically, the appropriate or dominating process may depend on characteristics of the task (Ashby & Maddox, 2005; Juslin, Jones, et al., 2003; Kruschke, 2004; Minda & Smith, 2001; Smith, Patalano, & Jonides, 1998). In research on multiple-cue learning, which involves judgements and feedback on a continuous criterion, learning is often described as the abstraction of rules regarding the relationship between features (or cues) and outcomes. On the other hand, exemplar theories have a prominent role in categorisation research. In categorisation studies, judgements (or categorisations) are made, and feedback is received, regarding a binary or nominal outcome (i.e. the category label). Juslin, Olsson, and Olsson (2003) demonstrated this association between task and cognitive process. Their participants encountered potentially poisonous bugs (on a computer screen) and were instructed to judge the level of toxicity based on the characteristics of the bug. When participants were instructed to judge the outcomes on a continuous scale (the level of toxicity in the blood), they tended to rely on abstract rules that connect the characteristics of the bug to the level of toxin. In contrast, when the task involved the categorisation of the bugs as either poisonous or not poisonous, information tended to be stored as concrete experiences of different bugs along with their category label, and later categorisations were made based on the similarity of a new bug to previously encountered bugs.

Thus, research that aims to specify the cognitive mechanisms involved in learning from judgements and their outcomes offers several models that rely on different processes and representations. One aim of this thesis is the investigation of the coding of selective feedback, where a well-supported categorisation learning model will be used to benefit the analyses. The *Generalised Context Model* (GCM, Nosofsky, 1986; Nosofsky & Johansen, 2000) is a widely applied model that formalises categorisation learning as the storing of concrete experiences of features along with the category label provided in outcome feedback (as described above). One advantage of a model that relies on the storing of exemplars rather than the abstraction of rules or a central tendency is that the distribution of events is naturally preserved in memory. When individual experiences are stored, the decision maker may directly infer distributional properties, such as the variability of observations or the experienced proportions of different categories (Hasher & Zacks, 1979, Zacks & Hasher, 2002). The GCM may be used as a tool when investigating different aspects of the cognitive processes associated with judgement and categorisation learning. The model will be returned to in this context later in the thesis.

Learning From Feedback

Outcome Feedback

As outlined above, research on judgement or categorisation learning generally involves learning from complete outcome feedback on decisions, i.e. participants receive information after each decision on the actual state of affairs. The coding of exemplars, the abstraction of rules or prototypes, or the updating of associations are made based on information received in features and in outcome feedback. Other forms of feedback have been investigated, including cognitive feedback that offers information to the judge concerning the regularities of her decisions, and task information feedback that informs the judge of the properties of the specific task. The latter information is often described as feed-forward because it provides information beforehand that may be used in later judgements rather than feedback on already made decisions. When such task information is given to participants it has been demonstrated to benefit performance (Balzer, Doherty, & O'Connor, 1989; Karelaia & Hogarth, 2008). In general, cognitive feedback and outcome feedback appear to be beneficial to the extent that they allow participants to elicit information of the task (Newell et al., 2007).

In fact, outcome feedback has repeatedly been demonstrated not to aid, and may even hinder, learning in probabilistic or complex environments (Hammond & Summers, 1972; Hammond et al., 1973; Jacoby et al., 1984; Kluger & DeNisi, 1996; Maddox, Love, Glass, & Filoteo, 2008; Remus et al., 1996; Schroth, 1997). Several explanations have been offered for such negative effects of outcome feedback on performance. People may be unable to use the information provided in outcome feedback or they may become inconsistent in the use of acquired task knowledge (Hammond & Summers, 1972; Karelaia & Hogarth, 2008). Full outcome feedback may encourage the selection of risky alternatives, which can be detrimental to performance in some environments (Grosskopf et al., 2006; Yechiam & Busemeyer, 2006). Feedback may trigger explicit learning strategies that are inadequate for more complex task environments (Hammond, Summers & Deane, 1973; Maddox et al, 2008). Moreover, the evaluative aspect of outcome feedback may direct the locus of attention of the decision maker from learning toward self-focus concerns, which would impair performance (Kluger & DeNisi, 1996).

In contrast to studies that demonstrate detrimental effects of outcome feedback, considerable research has reported high performance when learning from outcome feedback. As noted, the research areas of categorisation learning and multiple-cue judgement usually involve tasks in which participants learn from experience with complete outcome feedback. Such research generally demonstrates good performance on the part of participants (e.g., Ashby & Maddox, 2005; Fiedler & Juslin, 2006; Kruschke, 2005; Minda & Smith, 2001). Possibly, the high performance reported in this research may be understood in terms of the lower complexity and deterministic nature of the tasks.

In this thesis, the effects of learning from outcome feedback on performance are investigated. The focus on outcome feedback is motivated not only by the strong position held by this form of feedback in research on categorisation and multiple-cue judgement. The study of outcome feedback is also important from an ecological perspective. In everyday experience, people are unlikely to receive information in the form of organised presentations of the specific relationship between cues and criterion, or the weight of each cue. In learning from experience, if any information is received after a decision, it will often be in the form of (selective) outcome feedback.

Sensitivity to Outcomes

When complete outcome feedback is received by the individual, all feedback information may not be processed in the same manner or produce equal effects on the beliefs held by the decision maker. A general finding in research on people's assessments of covariation between variables is that judgements are most influenced by the joint presence of the variables and least influenced by their joint absence (e.g., Allan, 1993; Kao & Wasserman, 1993; Lipe, 1990; McKenzie & Mikkelsen, 2007). Similarly, in judgement research, it has been suggested that people rely more heavily on information received in hits (positive decisions with positive outcomes, Einhorn and Hogarth, 1978) or events that are present on positive trials (Newman, Wolff, & Hearst, 1980). For example, if an individual wants to evaluate whether a particular test is predictive of an illness, she is likely to focus on cases in which the test results were positive and the illness present, and devote less attention to instances where the test results were negative and the illness absent.

Further, a large body of research on hypothesis testing and social judgement indicate that people tend to rely more on outcome information that confirms a previous assumption (e.g., Crocker, 1981; Klayman, 1995; Nickerson, 1998; Ross, Lepper, & Hubbard, 1975). These effects have been explained in functional terms, where the differential influence of outcome information on the beliefs of the individual is explained in terms of the relationship between the prior belief and the new information. Bayesian inference provides a normative framework for learning as the updating of a prior belief in light of new information. The subjective probability of the prior belief may be more or less adjusted in response to an observation depending on the strength of the prior belief, the probability of the observation given that the prior belief is true, and the probability of the observation given that the prior belief is false (Edwards, Lindman, & Savage, 1963; Fischhoff & Beyth-Marom, 1983; McKenzie & Mikkelsen, 2007). The updating of a prior belief takes place gradually in light of conflicting new information. When the prior belief is strong, inconsistent information may not achieve immediate discernible influences on the knowledge of the individual.

The Bayesian view on the revision of a prior belief may be illustrated by results reported by Fugelsang, Stein, Green and Dunbar (2004) in the context of hypothesis testing in scientific practice. The authors demonstrated how data inconsistent with theory were initially discounted by participants, and theory was modified only with repeated exposure to disconfirming information. By a Bayesian account, the individual is unlikely, as well as unadvised, to reject theory on the first disconfirming evidence. If a theory is supported by previous findings and inconsistent information is inconclusive and uncertain, learning may be achieved by the revision of a belief with accumulated evidence. In this manner, the perseverance of a belief in the face of conflicting information may be adaptive.

From a Bayesian perspective, the processing of outcome feedback involves the interaction between the strength of the prior belief and the characteristics of the information gained in feedback. Many laboratory tasks are designed to minimise the influence of prior beliefs. However, in experiments where prior knowledge or expectations are present (such as in Study III in this thesis), they are likely to affect the decisions of the individual whether outcome feedback is complete or selective.

Selective Feedback

An important task of psychology is the study of behaviour that is representative of a natural environment outside of the laboratory (e.g., Brunswik, 1943). In this context, Hogarth (2006) performed an empirical investigation of people's everyday decisions. Participants reported their current decision-making activities at random moments during the day, including the expectation of feedback on the decisions. For about 40 % of decision, participants expected no feedback. Specifically, feedback was hardly ever expected for actions not taken. It seems that people have learned not to anticipate feedback information when they decide to reject an alternative. As noted previously, when someone says "no" to a product or an activity, she will rarely experience the actual outcome.

Clearly, such feedback structures in an everyday environment are very different from the complete outcome feedback usually encountered in a laboratory learning experiment. Indeed, the findings of good performance when learning from (complete) feedback in experimental laboratory research, as contrasted with research on intuitive judgement and decision making that demonstrates poor decisions and biases, have been explained by reference to the amount and quality of feedback received on decisions. Selective experience has been offered as an explanation of biased judgement (Brehmer, 1980; Dawes, 1994; Denrell, 2005; Einhorn & Hogarth, 1978; Fazio et al., 2004; Fiedler & Juslin, 2006; Fischer & Budescu, 2005; Hogarth, 1981; Nickerson, 1998), and specifically, of poor expert judgement and overconfidence (Dawes, 1994; Dawes, Faust, & Meehl, 1989; Garb, 1989).

The "illusion of validity" associated with high confidence in erroneous judgement is exemplified by Einhorn and Hogarth (1978) in the recruitment of job applicants. Because no outcome information is available for rejected applicants, only the outcomes of accepted applicants may be evaluated. If the majority of recruited applicants are successful, this information will appear to confirm the validity of the decision policy and induce a high level of confidence in decisions. Possibly, however, many of the rejected applicants are also suitable, in which case the validity of the decision policy is actually low. Selective feedback thus allows the development of overconfidence in decisions.

If the higher performance observed in laboratory settings compared with everyday judgements are caused by such differences in feedback, complete outcome feedback should benefit learning and improve the quality of decisions also in an everyday environment. Weather forecasting and racetrack betting are examples of areas of expert decision making characterised by complete outcome feedback that are actually associated with higher accuracy and less overconfidence than many other areas, such as clinical or economical decision making, where feedback is selective (Dawes, 1994; Johnson & Bruce, 2001).

Further, selective feedback should also be expected to produce detrimental effects on performance in a laboratory environment. Surprisingly few studies have been made to test this hypothesis. Actually, the empirical basis for the idea that selective feedback causes bias is quite sparse, not least when considering the prevalence of selective feedback in our everyday environments. A few studies have compared decision behaviour resulting from selective feedback and complete feedback (the latter sometimes referred to as information on foregone payoffs), and even fewer have investigated the effects on performance of learning with selective feedback.

It has been predicted and demonstrated that selective feedback tends to reduce risk-taking behaviour as compared with full feedback (Denrell & Le Mens, 2007; March, 1996; Yechiam & Busemeyer, 2006). When feedback is received only after positive decisions, any mistaken assumption that an alternative produces good outcomes will be corrected by experience (of negative feedback). However, because feedback is not received after negative decisions, a mistaken assumption that an alternative will produce negative outcomes tends to persist. This results in a conservative tendency to prefer to say "no" rather than "yes", i.e. to avoid rather than approach (Denrell, 2005; Eiser, Fazio, Stafford, & Prescott, 2003; Fazio et al., 2004). The higher risk behaviour with full feedback may have different consequences on performance depending on the characteristics of the learning environment. In general, however, effects of complete feedback on performance tend to be only weakly positive or even negative (Griffiths & Newell, 2009; Grosskopf et al., 2006).

Fazio et al. (2004) investigated the formation of attitudes toward objects when feedback is selective. In an experiment that directly compared effects of learning from complete and selective feedback contingent on positive decisions, participants' level of performance was not affected by feedback condition. However, in line with the conservative choice behaviour resulting from selective feedback, there was a larger proportion of correct positive decisions than negative decisions when feedback was selective. The tendency to say "no" rather than "yes" will naturally produce more faulty negative decisions than positive ones.

As evident in this overview, the effects of learning with selective feedback on performance have received little attention in experimental research. Further, findings have demonstrated that the effects of selective feedback on performance may be surprisingly small. In this thesis, the effects of systematically selective feedback will be investigated in a laboratory task traditionally associated with complete outcome feedback on each decision. The availability of feedback will be manipulated and the effects on performance and decision behaviour evaluated. We thus introduce an aspect of many everyday environments in the laboratory task in order to determine its significance for the process and products of learning.

Sources of Selectivity

When feedback is selective, it is absent on some occasions. A further specification is needed to indicate whether there is a systematic source of selectivity. Feedback randomly absent is a theoretical possibility (and a random element is likely to be present in any feedback structure). A large amount of randomly presented feedback is likely to be representative of the environment, in which case learning should mirror that from complete outcome feedback.

Two forms of selectivity that seem more ecologically relevant will be presented. These two forms are associated with different sources of selectivity, depending either on the actions of the individual or on the structure of the environment. These types of selective feedback are outlined in Figure 1, which is an illustration of a task that involves decisions whether to invest in companies. The oval area indicates a positive correlation between decisions and outcomes, where a positive decision to invest is usually associated with a positive outcome and a negative decision with a negative outcome. In terms of *Signal Detection Theory* (Swets, 1964), a correct positive decision is a *hit* and a correct negative decision is a *correct rejection*. An incorrect positive decision is a *false alarm* and an incorrect negative decision is a *miss*. High performance is achieved by maximising hits and correct rejections and avoiding incorrect decisions, i.e. false alarms or misses.

Decision-contingent feedback is outcome feedback received contingent on the decisions of the individual. An example would be when information is received only for positive decisions and is absent for rejected alternatives, i.e. the consumer is informed of the quality of a product only when she actually tries it, and not if the product is rejected. Results are known only for investments that are made while nothing is known about the outcomes for declined investments. The suitability of an applicant that was accepted and employed may be evaluated, whereas no information is gained concerning the suitability of rejected applicants. In Figure 1, when feedback is received contingent on positive decisions, the decision maker is informed of the outcomes for hits and false alarms but receives no information for misses or correct rejections.

In contrast, for *outcome-contingent feedback*, information depends not on the decision but on the actual outcome. This may be illustrated in the case of a critical reviewer reporting all the deficiencies of a product while not mentioning any positive qualities. Feedback is provided only when the outcome is negative. Denrell (2003) described the effects of outcome-contingent feedback only for positive outcomes as a case of sampling bias in organisational learning. The characteristics of surviving and successful firms are available as learning material, whereas information regarding firms that have failed is not considered. Specifically, because risky practices tend to result in either success or failure (and the closing of the firm), risk will appear to be associated with success even when there is no actual correlation between risky practices and success in the population. Feedback contingent on positive outcomes in Figure 1 suggests that outcome information is received only for successful companies, i.e. only for hits and misses (no feedback is received for false alarms or correct rejections).

Decision-contingent and outcome-contingent selective feedback are presented here in their "ideal" forms. In an everyday environment, different structures of selectivity may be combined and are probably rarely absolute. In this research, however, we begin the investigation of the effects of these kinds of selective feedback in their pure forms.

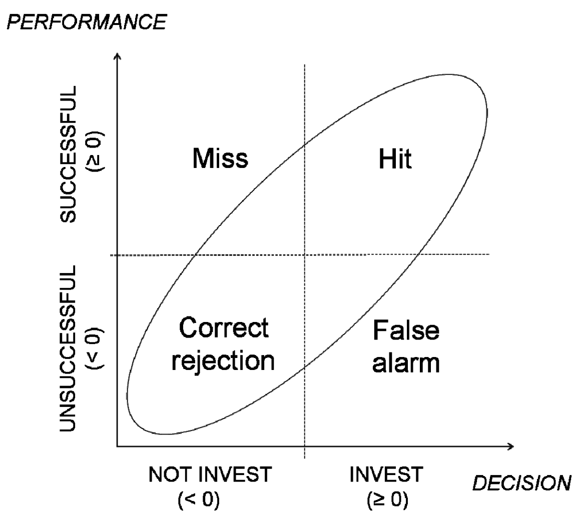

Figure 1. An illustration of a task involving decisions whether to invest in companies, where decisions are positively correlated to the outcomes. The outcome of a decision falls into one of four outcome categories depending on the decision and the actual outcome of the company. A positive decision (to invest) produces a *hit* if correct or a *false alarm* if incorrect and a negative decision is a *miss* if incorrect or a *correct rejection* if correct.

Sampling

Whatever the source of the selectivity, selective feedback implies that the experienced feedback is a sample of all outcomes, a sample that may be more or less representative of all possible outcomes. In this thesis, particular attention will be devoted to decision-contingent feedback, where the actual behaviour of the decision maker influences the knowledge that she gains from the environment. If feedback is contingent on positive decisions (i.e. one will not know the outcome unless one tries), the individual who makes a larger amount of positive decisions will receive more information. A sampling approach to learning, as outlined earlier in the metaphor of the naive intuitive statistician, emphasises how the individual's sample of experiences constitutes the learning material, where only sampled experiences may be used in learning (Fiedler, 2000; Fiedler & Juslin, 2006). In this manner, when people learn from their sample of experiences and their decisions determine whether feedback is received, the decisions affect what may be learned. In this context, the objectives of the decision maker may become highly relevant. If the decisions of the individual are motivated by the aim to gather information and learn from the outcomes, a larger sample is likely to be collected. In contrast, when the decisions are motivated by a focus on immediate performance and the aim of the individual is to produce favourable outcomes and avoid negative outcomes, the decisions are likely to result in a smaller sample of outcomes from which to learn.

Focus on Performance vs. Learning

Many everyday decisions are made with a focus on the immediate consequences of our decisions and the intention to produce beneficial outcomes. When people decide between alternative experiences, they often choose an alternative with an expected positive outcome, rather than the alternative that offers a good opportunity for learning. The employer wants to hire the best job applicants and decline unsuitable applicants, and any everyday consumer wants to buy only good products and avoid the poorer alternatives. Under such circumstances, learning will merely be a by-product of the aim to seek positive outcomes and avoid negative ones. The goal to produce beneficial outcomes at limited expenses in terms of time, money, and other resources may severely reduce the information gained from a restricted sample and thus the opportunity to learn (Brehmer, 1980; Einhorn & Hogarth, 1978, 1981; Fazio et al., 2004). Fiedler (2008) notes how the sampling strategies associated with the aim to make optimal decisions will determine the sample of information to be used in later decisions. The author refers to this as "the ultimate sampling dilemma", where sampling bias is the rule and representative sampling the exception.

Denrell (2005) illustrated the effects of sampling behaviour in a model on attitude formation, where biased social attitudes may arise from a correctly described observed sample; however, this results from biased sampling. Naturally, people will systematically try to avoid the negative experiences of interacting with disagreeable persons, seeking, instead, the company of people they like. An initial negative impression of an outgroup member is therefore likely to lead to the choice not to interact further with this person or others in that group, with the result that no more information is received that may disconfirm the initial negative impression. An initial negative attitude, whether correct or not, will therefore last and the only way for the decision maker to discover that a negative impression is false is if she is required to interact with the people that she expect to dislike. Interaction enables positive encounters, which serve to correct false negative impressions and thereby generate more favourable attitudes. The attitudes of an individual are therefore likely to be more favourable toward people with whom she is obliged to interact, such as close colleagues, as compared with her attitudes toward people who are more easily avoided.

The model presented by Denrell illustrates how decision behaviour may become restricted as an effect of negative initial experiences. In a similar manner, the tendency to reject most alternatives may be an effect of previous knowledge or expectations, as based on experiences in a similar environment or on theoretical knowledge (Alloy & Tabachnik, 1984). In any situation where decisions are highly restricted because of fear, prejudice, or superstition, selective decision-contingent feedback results in a smaller, and often unrepresentative, sample of experiences from which to learn.

On the other hand, when the expectations of the individual promote liberal decision behaviour, the resulting sample is likely to be more representative of the environment, which should benefit performance. With or without previous knowledge, a learning focus on the part of the decision maker indicates that the aim is to gather and evaluate information rather than to produce the most beneficial outcomes in the short term. When the aim is to use information for long-term learning, a negative outcome may be valuable if informative. A focus on learning may thus result in more liberal response behaviour and a sample of experiences more representative of the environment. This information-collecting procedure is the logic of the scientific method. Representative experience is achieved by means of representative sampling, which in turn enables representative knowledge. Researchers, exposure therapy patients, and strategic everyday learners, however, are all aware that such hypothesis-testing behaviour may be highly time-consuming, cognitively and emotionally demanding, and costly. Although representative sampling of participants may possibly be achieved in much research, representative sampling of stimuli is rarely even considered (for exceptions, see Brunswik, 1943; Dhami, Hertwig, & Hoffrage, 2004; Fritzsche & Brannick, 2002; Hogarth, 2005; Hurlburt; 1997).

The sampling perspective emphasises the association between the actual experience of the individual and her knowledge. Whether sampling behaviour produces a more representative sample or is restrictive and biased, the knowledge of the individual will reflect her specific experience. The conflict between the gathering of information for long-term learning vs. immediate performance is likely to be present in most situations where information is contingent on action. When information is received for positive decisions only, such as the purchase of a product, the decision maker may have to balance any ambitions to receive valuable information against more short-term outcome goals. In machine learning, this dual role of choices is apparent in reinforcement learning systems. The outcome is known only when a positive action is taken and therefore choices are a basis for selection as well as for learning, resulting in a trade-off between exploration and exploitation (i.e., between the gathering of information vs. the maximisation of immediate positive outcomes, Sutton & Barto, 1998).

The sampling effects outlined above demonstrate the significance for learning of the size and representativeness of the experienced sample. The issue of sampling effects is thus related to the topic of the strategies by which people learn through hypothesis testing. The focus of this thesis however, is the investigation of the effects of learning with selective feedback on performance and decision behaviour, rather than the decision strategies that may underlie the specific sample of the individual. There is a host of studies on hypothesis testing and the interested reader is referred to any of several presentations that present different views on the topic (e.g., Fischhoff & Beyth-Marom, 1983; Klayman & Ha, 1987, 1989; Oaksford & Chater, 1994; McKenzie & Mikkelsen, 2007).

Coding of Selective Feedback

As outlined previously in this thesis, cognitive models of human learning describe the process of acquisition and representation of knowledge in memory. Surprisingly however, these models have not paid any attention to the consequences for the acquisition of knowledge when outcome feedback is systematically absent on some occasions. Markman (1989), in a comment on a network model outlined by Gluck and Bower (1988), noted the dilemma associated with the coding of feature information that is absent. The binary coding used in the model indicates the absence of feature information by the same activation level as negative information (-1). Similarly, the presence of a feature is indicated by a positive activation (1). Markman (1989) proposed an alternative means of coding, involving neutral coding (0) for absent information, which would reserve negative values for negative feature information. It was illustrated how the alternative coding rules may yield quite different predictions concerning the resulting categorisations. The coding of absent information thus influences what is learned from experiences where feedback is absent. However, the author did not link these different means of coding to the possible consequences for coding of systematically absent information or the psychological processes associated with such coding.

Nosofsky (1986) presented an adapted version of the GCM (the *augmented GCM*), where categorisations of exemplars without feedback were assumed to be used as “inferred exemplars” in later categorisations. The augmented GCM provided a substantially better fit for some of the category structures used in the experiment. The rationale for the augmented version, however, is unclear and the model does not involve the processing of feedback that is systematically absent on some occasions. Unfortunately, the augmented version of the GCM has hardly received any attention in later research.

We may conclude that a large number of studies on human judgement and categorisation learning have not resulted in any discussions of the processes associated with the coding of feedback that is systematically absent on some occasions. In general, it appears to be assumed that if no outcome feedback is received, nothing can be coded and hence nothing is learned from that experience. This view on learning as taking place only in response to actually presented feedback is evident in reinforcement learning networks in which feedback is received only for positive decisions and learning (modification of weights) occurs only for these decisions (Eiser et al., 2003; Sutton

& Barto, 1998). For negative decisions, where feedback is absent, no learning occurs. This "default" assumption that learning only occurs in response to the actual experience of the learner will be referred to in this thesis as *positivist coding*, emphasising the reliance exclusively on externally presented information. We will see how this may actually be a highly problematic assumption.

As outlined previously in this thesis, empirical findings demonstrate that selective feedback contingent on positive decisions tends to produce conservative decision behaviour, where participants prefer to make negative decisions (Denrell, 2007; Fazio et al., 2004; March, 1996; Yechiam & Busemeyer, 2006). When participants learn to avoid negative outcomes, their actual experience will contain a larger proportion of positive than negative outcomes. Griffiths and Newell (2009) compared the effects of learning with selective and complete feedback with different objective base rates of positive and negative outcomes in a study using the experimental paradigm introduced in Elwin et al. (2007; Study I of this thesis, further described below). Results demonstrated that the more conservative decision behaviour seen with selective as compared with complete feedback was accompanied by the actual belief - explicitly reported by the participants - that the base rate of positive outcomes in the material was lower. It appears that although external feedback reveals a large proportion of positive outcomes, participants come to perceive the proportion of positive outcomes in the environment as low.

The explanation that is offered in this thesis suggests that participants' perceptions of proportions are influenced not only by their actual experiences with feedback but also by non-feedback experiences. Experiences made without feedback may thus be involved in experiential learning. The suggestion that will be outlined and tested is referred to as *constructivist coding* because it involves the coding in memory of not only externally presented experience but also of the inferences made by the individual. The constructivist coding hypothesis is described in more detail below.

The Constructivist Coding Hypothesis

When participants learn from repeated trials and feedback is absent after negative decisions, it appears unlikely that they simply ignore non-feedback trials. In a recruitment task, for example, when no feedback is received on the suitability of rejected applicants, the recruitment officer may have an idea concerning the outcomes. The officer may even be sure what the outcomes must be, assuming that she has enough knowledge on which to base such inferences. When no external feedback is received for rejected applicants, such feedback may be unnecessary for the officer to "know" the outcomes. The applicants were rejected because the recruitment officer judged them unsuitable for the job, and thus the outcomes were expected to be poor.

If, on a later occasion, the recruitment officer is asked to estimate the proportion of applicants who are suitable (or unsuitable) for the job, she is unlikely to base the answer solely on the accepted and evaluated applicants. Although no feedback was received concerning rejected applicants, the likely outcomes were inferred and acted upon in the rejection decisions. If many applicants were rejected because they were believed unsuitable, the inferred experience contains a large proportion of unsuitable applicants. If the recruitment officer's inferences inform her that most applicants are unsuitable, she will consequently believe that most applicants are unsuitable and hence prefer to say "no" rather than "yes". Because the same officer's objective experience of external feedback is likely to consist of a high proportion of suitable applicants, this example highlights how experiences without feedback cannot be assumed to be ignored but are likely to be cognitively processed and stored in memory.

The argument will be made in this thesis that when no external feedback is received after a decision, the inference by the individual will be coded into memory as "internal" feedback. This process is referred to as *constructivist coding,* expressing the idea that inferred information is coded into memory and processed in a similar manner as externally presented feedback. We return to the recruitment officer who makes decisions based on her knowledge and receives external feedback only for the employed job applicants. Each time the officer expects a positive outcome, she will employ the applicant and thus the actual outcome will be revealed in feedback. When she expects a negative outcome and makes a rejection decision, no feedback is received. Here, constructivist coding suggests that the expected (negative) outcome is coded into memory in the form of "internally constructed" feedback.

Learning, in the form of updating of knowledge or beliefs, thus occurs not only for external feedback but also in response to inferred outcomes. In this manner, constructivist coding implies that the internally generated feedback coded in the absence of external feedback has a direct causal effect on the beliefs and knowledge of the individual. For each negative decision by the recruitment officer, internally generated and coded feedback of an expected negative outcome strengthens her beliefs, as well as her confidence in these beliefs. Further, as outlined above in the context of perceived proportions of outcomes, the coded internal feedback has a direct effect on the perception of the distribution of outcomes in the environment. If the recruitment officer rejects a large proportion of the job applicants and therefore infers and codes a large number of negative outcomes, her perception will be that a large proportion of the applicants was unsuitable. Some of the inferred negative outcomes may be mistaken, which is why the rejection of fewer applicants is likely to result in the perception of a higher proportion of suitable applicants.

When fewer applicants are rejected, the officer is likely to experience a larger number of positive outcomes in feedback, and she will make fewer negative inferences.

Psychology has long endorsed the idea that people base their beliefs on the (re)construction of knowledge and inferences, as will be outlined presently. The novel suggestion implied by constructivist coding, however, involves the hypothesis that an inference is actually coded into memory and processed in a manner similar to external feedback. That is, the inference (e.g., that a particular applicant is unsuitable) is coded and processed in a manner similar to an externally experienced event. On a later occasion, when the recruitment officer evaluates the outcomes of the applicants, the inferred outcomes are retrieved from memory along with (and indistinguishable from) the actually experienced outcomes.

Although the idea of constructivist coding of unseen outcomes has not previously been explicated, a wide body of research supports the hypothesis. Inferential and constructive processes have long played a central role in the understanding of human cognition. Further, the limited ability of individuals to monitor the sources of their memories indicates that, once an inference has been made, it may be impossible to distinguish it from externally presented information. In addition, research on the neural mechanisms associated with false memories demonstrates processes suggestive of the coding of internally generated information. Studies on the role of expectancy in avoidance learning also suggest learning from internally generated feedback. Finally, the coding of inferences when data are missing is a procedure that benefits the analyses of research data when it is put to systematic use in statistical imputation procedures. These arguments will now be presented in more detail.

Constructive Cognitive Processes

Human cognition has long been viewed as the interplay between mental structures and external information (e.g., Bartlett, 1932; Kant, 1781; Piaget, 1954). Specifically, the idea that people may rely on inference and "internal" feedback in addition to external information is in line with the well-documented knowledge that memory is not a mere recollection of externally presented information stored in memory but a (re)constructive process. numerous studies on top-down perception, schemas, false memories, and witness psychology have demonstrated how people do not simply perceive or remember external information. Rather, they interpret, infer, and supplement information that is lacking or that they expect to be there (e.g., Alba & Hasher, 1983; Dawes, 1994; Hawkins & Hastie, 1990; Loftus & Hoffman, 1989; Roediger, Watson, McDermott, & Gallo, 2001; Roediger & McDermott, 1995; Schacter, Norman, & Koutstaal, 1998).

Most research on inferential influences on information processing and judgement has been performed without specifying whether (or how) the inferences are actually assumed to be coded into memory. Other research concerns inferences that take place at the time of judgement, or retrieval, in which case the processes may be referred to as "*re*constructive" (e.g., Loftus & Hoffman, 1989). In contrast, as will be presented in a later section, several reports found evidence that specifically suggests the coding of inferences into memory. In these cases, the stored memory traces of the individual will contain externally presented information together with inferred information. In this thesis, "constructive processes" are a general expression referring to the broad range of phenomena that involves inferences and top-down processes, among which are reconstructive processes (it may be emphasised, however, that the process of *constructivist coding* is assumed to involve the actual coding of inferences into memory).

Information may be absent not only for certain outcomes but also for one or more features (cues). Judgements involving missing cue information may be illustrated by a physician making a diagnosis based on laboratory test results, of which some results are unavailable. Another example would be an investor predicting the outcomes of companies in which information of the characteristics of each company is incomplete. Research on the effects on judgements of missing cue information demonstrates that, rather than ignoring missing values, people infer and "fill in" lacking information based on previously presented information and the task context. Moreover, these inferences have been shown to influence their judgements (Garcia-Retamero & Rieskamp, in press; Jaccard & Wood, 1988; Johnson, 1987; White & Koehler, 2004). Clearly, people infer information based on their knowledge and beliefs, and such inferences affect current decisions. We will now turn to findings suggesting the coding of such inferences into memory and the effects of these inferences on later decisions.

Reality Monitoring and Source Monitoring

The source monitoring framework describes how people judge the origin of a particular memory; whether it was read in a book or told by a friend or by someone on the radio (source monitoring), or whether it was something that actually happened or only imagined (reality monitoring). The source of the memory is not evident in the memory trace as such. Instead, the source is attributed by the individual based on available cues, such as the level of sensory detail in the memory trace. Such cues may be insufficient to judge the actual origin of the memory trace, which is why misattribution occurs. Research on the ability of people to monitor the sources of their memories indicates that when the contexts of the sources are similar, the ability to distinguish remembered information from inferred or imagined information is poor (Johnson, 2006; Johnson, Hashtroudi, & Lindsay, 1993; Johnson & Raye, 1981). In fact, people have been found to remember events that never occurred and thus create "false memories" that they report with high confidence (Roediger & McDermott, 1995; Roediger, McDermott, & Gallo, 2001; Walther, Fiedler, & Nickel, 2003).

Such failures in source and reality monitoring imply that once an individual makes an inference, she will often be unable to distinguish inference from external reality. When the recruitment officer described above recalls the outcome of a particular applicant, she may easily confuse external feedback with inference and will therefore rarely know if this particular applicant was presented with external feedback or if the outcome was inferred (and considered highly likely). An inferred outcome, even if somehow coded as an inference rather than as an external experience, may be indistinguishable from actually presented information when the stored representations are retrieved and relied upon in later decisions.

Coding of Inferences

The overview of constructive cognitive processes presented above demonstrates how an individual easily make inferences when external information is lacking. Limitations in source monitoring ability indicate that once an inference has been made, the individual may be unable to tell the inference from externally presented information, especially if they share many characteristics. These processes may explain why an inference made in association with a decision is later confused with an actually presented outcome.

Moreover, research on the processes associated with false memories suggests that inferential processes may actually be coded into memory in a similar manner as externally presented feedback. The neural mechanisms underlying false memories indicate that the cognitive processes associated with a decision without feedback are highly similar to those associated with the encoding of externally presented feedback. Such research emphasises the importance of encoding as well as retrieval processes in the forming of false memories. Elaborate semantic and visual processing at the encoding stage is associated with later reports of a false memory, such as the memory of an associated word or the memory of having seen a picture rather than merely the word (Geng, Qi, Li, Fan, Wu, & Zhu, 2006; Gonsalves & Paller, 2000; Gonsalves, Reber, Gitelman, Parrish, Mesulam, & Paller, 2004; Kim & Cabeza, 2006; Nessler, Mecklinger, & Penney, 2001; Okado & Stark, 2009).

An example is provided by Gonsalves et al. (2004) in a study in which participants observed words (object names), some of which were followed by a photo of the corresponding object. The authors reported activation in areas that reflect visual imagery for words that were later falsely remembered as having been presented pictorially. These results suggest how constructive processes may actually take place at encoding, where a memory trace is formed based on elaboration or inference and is later retrieved as a false memory of externally presented information.

The notion that learning may take place in response to inferred or expected information, in addition to external feedback, has also been suggested in a clinical context. Lovibond (2001) described learning from "near miss" experiences in anxiety disorders. When experiencing a "near miss", the individual believes that a harmful outcome would have occurred if a specific action (such as escaping the situation) had not been performed. In this manner, learning takes place from the inference of a negative event that actually does not occur. Salkovskis, Clark and Gelder (1995) reported how panic disorder is maintained by avoidance, escape, or safety-seeking behaviour in anxiety-provoking situations. Such behaviour prevents the disconfirmation of the hypothesis that harm will come and may even serve as a confirmation of this expectation (e.g., "if I had not sat down then I would have fainted"), which would serve to maintain the disorder. Naturally, it is an adaptive ability to be able to learn about potential harm without having to actually experience it. However, in cases where the expected negative outcome and the perceived confirmation of a harmful outcome are not real, avoidance may result in the development or the maintenance of a clinical anxiety disorder.

Eiser et al. (2003) noted that when feedback is received only for positive decisions, attempts to avoid negative experiences will result in negative attitudes being more weakly based on actual experience than positive attitudes. Yet, the authors note that negative attitudes are often more resistant to extinction. As outlined above, one explanation is that avoidance, though not followed by actual information on the outcome, may be reinforced by a reduction in fear (or by the absence of the expected outcome, see Lovibond, Saunders, Weidemann, & Mitchell, 2008), which strengthens the avoidance response.

Eiser et al. (2003) implemented this idea of internally generated feedback, which confirms the expectation of a negative outcome, in a reinforcement learning network. In standard reinforcement learning, outcome information is contingent on positive action. Eiser et al. (2003) modified a reinforcement learning system so that avoidance (decision "no") was followed by feedback indicating that the outcome was actually negative. Consequently, the feedback received by the system in negative decisions always confirmed that the avoidance decision was correct. In this simulation, the expectation-confirming feedback was multiplied by an attenuation parameter of 0.1 (this was done to weaken the learning effect in response to these outcomes). Decision behaviour in the expectation-confirming network was compared with that in a common reinforcement learning network in which feedback was received only for approach responses (decision "yes"). Results in the two systems indicated a similar pattern of decision behaviour, possibly because of the attenuation parameter involved in expectation-confirming feedback, which made the systems quite similar. Although the results of the simulations are difficult to interpret, the study is noteworthy because it formalised the idea that people may learn from internal feedback, as generated from their expectations.

Briefly, another machine learning parallel to the coding of inferred information by humans is found in semi-supervised learning. Here, the learning of a network benefits from additional information provided in data without labels (feedback). In the example of *self-training*, the system is first trained with a small amount of labelled data and the information is used to classify unlabelled data, where the most confidently predicted labels are used in new training (Chapelle, 2006; Zhu, 2005). The process of using the outcomes of the learner's own classifications as actual labels is similar to the idea of constructivist coding of unseen outcomes. However, semi-supervised learning has not been related to human decision processes.

Statistical Imputation of Missing Data

In scientific practice, systematic use of the coding of an inference when information is lacking is found in the statistical analyses of multivariate data sets. In research studies (as in everyday life), data are likely not to be missing at random. Some participants may be unwilling to answer particular questions or people may be difficult to locate at later occasions in a longitudinal design. These participants may be systematically different from the rest of the sample on some characteristics. Accordingly, when data are absent according to some systematic process, the sample will not only be reduced in size but it will probably also be biased. When analysing the data set, simply deleting cases that contain missing data would result in an unrepresentative sample. Instead, imputed values may take bias into account by exploiting the information in existing data.

Various methods of statistical imputation thus benefit from information in existing data to infer the likely outcome for the missing piece of information. In this manner, missing values are replaced by "best guesses" in the form of imputed values (Rubin, 1976; Horton & Lipsitz, 2001; Sinharay, Stern, & Russell, 2001). Statistical imputing is a systematically applied practice performed by means of computerised statistical packages that benefit from large computational power to exploit the existing information. An everyday corollary to statistical imputing is suggested in constructivist coding, where the individual uses her knowledge in the intuitive and frugal coding of inferences and beliefs when external information is lacking. In this way, the individual benefits from the information received in externally presented information.

In summary, people have the ability to use inferential processes as a means to supplement external information. These inferences may be coded into memory along with externally presented information, and when combined with limitations in source monitoring, the conclusion is that learning may be based on inferences in the role of “internal feedback” in addition to external experience. Actually, when external information is absent, using experience to make “informed guesses” may be an adequate strategy – or the only possible – for dealing with an environment that offers some regularity but rarely clear-cut, immediate or unambiguous feedback. Despite the risks of inaccurate inferences, the ability of people to “fill in the gaps” and act upon knowledge that they believe to be true may be necessary to allow action, as well as the formation and maintenance of a coherent representation of the world. Studies I and II in this thesis investigated such learning from internal feedback in the cognitive process referred to as constructivist coding. In the presentation of Studies I and II, constructivist coding and its alternatives will be further outlined.

Purpose and Overview of the Studies

The three studies that are presented in this thesis investigated the effects of selective feedback on judgements and performance. As outlined earlier, the aim of the studies is the investigation of (a) the effects of selective decision-contingent feedback on performance, (b) explanations of such effects, and (c) coding in memory of selective feedback. Specifically, Study I examined the effects of selective feedback on decision behaviour and the coding of selective feedback. The study also introduced the hypothesis of constructivist coding. In Study II, constructivist coding of selective feedback was evaluated further against possible alternative principles for the coding of unknown outcomes. Finally, Study III pursued the investigation of the effects of selective feedback on decisions and examined to what extent these effects may be explained by the specific selective sample of experiences of the individual.

Method

The common features of the methodological aspects of the studies will be outlined and discussed in this section before each study is presented with its specific methodology, analyses, and results.

Task Structure and Overall Procedure

Participants were recruited primarily among university students who were compensated with a cinema cheque or course credits (performance-based compensation was added in some experiments). The experiments presented in this thesis were all based on the same basic computerised judgement task. Two task framings, or cover stories, were used. In Studies I and III, participants were instructed to decide whether to invest in each of a number of presented companies based on each company's values (between 0 and 10) on four features (e.g., amount of staff). In Studies II and III, participants were given the role of a recruitment officer who encountered various job applicants and whose task was to decide which applicants to employ based on the values on four characteristics of each applicant. In all tasks, the features of the companies or applicants were deterministically and linearly related to the outcome. The values of each applicant (or company) were presented on the computer screen as depicted in Figure 2. The participants were instructed that a positive judgement indicated a decision to recruit (or invest) while a negative judgement implied a rejection decision.

In a training phase, feedback was received after each decision (complete feedback) or only after some decisions (selective feedback contingent on decisions or on outcomes), depending on the learning condition of the participant. In a test phase, judgements were made without feedback. After each test decision, participants reported their confidence that the decision was correct (from 50%, indicating a mere guess, to 100%, indicating certainty). In some of the experiments confidence was reported also for training decisions.

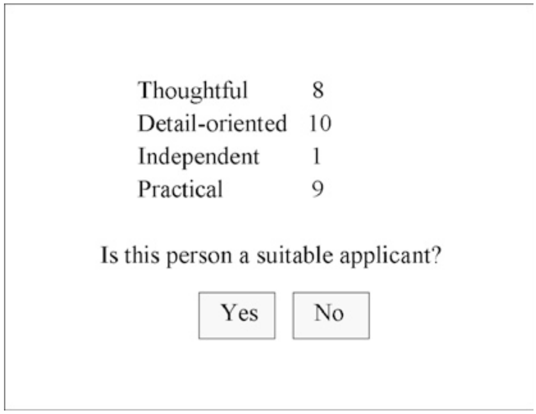

Figure 2. The presentation of an exemplar (job applicant).

Analyses and Dependent Measures

When people learn from their experiences, they gain knowledge about the structure of the environment as well as about their own judgemental ability. Such knowledge may be operationalised in different ways. Several measures of different aspects of knowledge were used in this thesis. A straightforward measure of performance is the proportion of correct decisions. In addition, we investigate participants' ability to distinguish between good and bad objects (e.g., suitable and unsuitable applicants). Another aspect of knowledge, or preference, is expressed in the response criterion of the individual: the tendency to favour positive or negative decisions. Further, knowledge concerning the distribution as a whole may also be abstracted by the individual, such as the overall proportion of positive outcomes. As mentioned, people also learn of their own abilities and knowledge and may express more or less appropriate confidence in their decisions as compared with their actual level of performance. These various measures will now be presented in more detail.

Signal Detection Theory

Signal Detection Theory (*SDT*; Swets, 1964) provides a framework and tools for analysing decisions under uncertainty, in terms of the knowledge and response behaviour of the individual. The individual's level of sensitivity or detectability is a measure of her ability to perceptually or cognitively discriminate between two outcome categories, such as signal vs. noise or suitable vs. unsuitable applicants. The detectability of the individual is calculated from the proportion of hits and false alarms ($d' = z_{\text{hit rate}} - z_{\text{false alarm rate}}$).

Unlike the proportion of correct decisions, the measure of detectability is not affected by the participant's tendency to favour one response over the other and thus it may be regarded as a more "pure" measure of knowledge.

The tendency to favour one response ("yes" or "no") over the other is measured separately in the response criterion of the individual, which indicates to what extent one response is preferred ($C = -.5 * [z_{\text{hit rate}} + z_{\text{false alarm rate}}]$). The response behaviour may be affected by the payoffs or consequences of the decisions: for example, severe negative consequences for false alarms are likely to lead to a general reluctance to make positive decisions. A preference for negative decisions (i.e. to reject job applicants) or conservative response behaviour is indicated by a positive response criterion, whereas liberal response behaviour or a tendency to accept applicants is indicated by a negative response criterion.

In addition to the parametric measures of detectability and response criterion, non-parametric alternatives are available when the distribution of outcomes fails to adhere to parametric assumptions (Donaldson, 1992; Macmillan & Creelman, 1991). Because of a skewed distribution of outcomes, Study I presents detectability as measured by the nonparametric A' and response criterion by B''_D. In Studies II and III, the more widely recognised parametric measures d' and C are used.

Overconfidence

When people learn from their experiences they also gain knowledge about their own judgemental competence. This meta-cognitive knowledge may be assessed by the degree of confidence expressed by the participant in relation to her decisions. The confidence of the individual may be more or less on a level with the actual proportion of correct answers, where underconfidence indicates that the level of confidence is lower than the proportion of correct decisions and overconfidence suggests that the participant is more confident than correct. When someone is well calibrated, her confidence is at level with the proportion of correct decisions (for other estimates of overconfidence, see, e.g., Moore & Healy, 2008).

Overconfidence, generally regarded as an indicator of bias, has been discussed in relation to selective feedback as an effect of confidence increasing with experience without a corresponding improvement in performance (Brehmer, 1980; Einhorn & Hogarth, 1978; Dawes, 1994; Hall, Ariss, & Todorov, 2007). Overconfidence has been demonstrated repeatedly in research, mainly in general-knowledge tasks, however, explanations of overconfidence effects are debated. Various accounts have been offered, including unrepresentative items and statistical effects (see e.g., Ayton & McClelland, 1997; Harvey, 1997; Juslin, Winman, & Olsson, 2000; Klayman, Soll, González-Vallejo, & Barlas, 1999). Overconfidence as a combined measure derived from the proportion of correct decisions and the level of confidence in decisions requires that care be taken not to confuse distributional properties of each measure with effects of a cognitive content. For example, ceiling effects in proportion of correct decisions may result in underconfidence simply because of a skewed distribution of measurement error (for a more extensive discussion of overconfidence effects, see e.g., Juslin et al., 2000).

Judgements of Distributional Properties

In addition to knowledge relied upon in individual decisions, people may record or infer general characteristics of the experienced environment, including the proportions of different outcomes, means, or variances. Hasher and Zacks (1979, 1984; Zacks and Hasher, 2002) have demonstrated how people automatically store information, such as the experienced frequencies, while interacting with an environment. In this thesis, participants' estimations of proportions of positive and negative outcomes in the material experienced with selective feedback were investigated in Study III.

Model Fit

The dependent measures presented so far concern the products of learning. These measures, which are various indicators of the knowledge of participants, however, usually provide no information on the specific processes and representations that underlie the decisions. As described in the section on cognitive processes, models of categorisation or judgement learning have been developed to describe the cognitive processes and representations involved in judgements. Essentially, these models attempt to describe the processes by which the individual updates her knowledge with experience and how she applies this knowledge in judgements of new exemplars.

A computational model enables the formulation and testing of more precise theoretical statements than is possible with mere verbal descriptions of cognitive processes. By fitting the model to empirical data, the free parameters of the model are set to capture the processes manifested in the decisions of the participants. The assumptions of the model can therefore be tested, preferably by means of cross-validation on another set of data. The ability of the model to capture the variance observed in data is indicated by the Root Mean Square Deviation (*RMSD)* or the coefficient of determination (r^2) between predicted and observed decisions. As with all descriptions of complex processes, the model seeks to capture and describe the common cognitive operations in different sets of empirical data rather than idiosyncratic (error) variance in a specific data set. The model should be flexible enough to capture the complexity of the cognitive process, but not so flexible as to produce overfit and low generalisability (Cutting, 2000; Pitt & Myung, 2002; Pitt, Myung, & Zhang, 2002).

To investigate the information involved in decisions with selective feedback, a cognitive process model was fitted to experimental data. In Study II of this thesis, the GCM (Nosofsky, 1986) was used as a tool to test the hypothesis of constructivist coding against alternative schemes of coding. Empirical data were compared with the predictions specified in the model regarding coding and information used in judgements. A slightly more detailed summary of the fitting of GCM to data in Study II is described in the presentation of the study (a complete account is found in the actual study).

Empirical Studies

Study I. Constructivist Coding: Learning From Selective Feedback

The aim of the study was to investigate the effects of learning from selective feedback on decisions and the coding of experiences without feedback. Two alternative hypotheses regarding the coding of selective feedback were suggested, namely positivist and constructivist coding. As outlined earlier in this thesis, *positivist coding* is coding of externally presented information when present. If feedback is absent, nothing is coded. In Figure 3A, when no feedback is received for misses and correct rejections, positivist coding entails that nothing is coded for these exemplars. Panel B illustrates the positivist coding of outcomes only for positive decisions.

As noted, positivist coding is the implicit assumption associated with the ignoring of selective feedback in learning (e.g., Ashby & Maddox, 2005; Kruschke, 2005; Minda & Smith, 2001) and is made explicit when assuming that learning takes place only in response to external feedback (Eiser et al., 2003; Sutton & Barto, 1998). However, naively applied positivist coding will produce extreme effects if the decision maker is assumed to entirely ignore non-feedback experiences in her decisions. If performance is good, most experiences with feedback will be positive (i.e. more hits than false alarms). Consequently, if the experienced base rate of positive outcomes is taken as the actual base rate for all exemplars, this should lead to the perception of an unproportionally high base rate of positive outcomes among all exemplars.

Positivist coding may therefore appear more plausible if complemented with an inference mechanism at the time of a later decision. Positivist coding always implies that only external feedback is coded into memory. However, a "sophisticated" positivist not only retrieves this coded material in later decisions but also assumes an equal proportion of correct decisions for negative decisions (where feedback was absent) as for positive decisions (where the level of performance may be evaluated against external feedback). We call this modified version *sophisticated positivist coding* with a symmetry assumption. The base rate of successful companies may thus be inferred from the experienced level of accuracy for positive investment decisions. If the level of accuracy for feedback decisions, *s*, is assumed to be the same as the accuracy for non-feedback decisions, the proportion of positive decisions, p, may be used to infer the base rate of positive outcomes as

$$p \cdot s + (1 - p) \cdot (1 - s)\,. \tag{1}$$

Like the naive positivist, the sophisticated positivist codes only the observed outcomes into memory. The inference made by the sophisticated positivist concerning the symmetry of accuracy is made at the time of later decisions and any coding of this inference into memory should be mentally distinct from, and not confused with, external feedback. (The non-separation or confusion of these representations in memory is regarded as an expression of constructivist coding.)

Like positivist coding, *constructivist coding* assumes that external feedback is coded when present. However, in the absence of feedback, internal information is constructed based on the knowledge or expectations of the individual. As outlined earlier in this thesis, this internally generated outcome feedback is coded into memory along with external feedback. Figure 3C illustrates constructivist coding when no external feedback is received for misses and correct rejections. For negative decisions in which feedback is absent, the decision maker codes the expected negative outcome. The constructivist coder, unlike the sophisticated positivist, thus codes the inference into memory, later to be retrieved along with coded material from external feedback. Consequently, the constructivist decision maker will estimate the base rate of positive outcomes from the level of accuracy for positive decisions, s, and the proportion of positive decisions, p, as

$$p \cdot s\,. \tag{2}$$

To summarise, positivist coding involves coding only of externally presented feedback. When feedback is received only for positive decisions, only these experiences are coded, and nothing is stored for negative decisions. At the time of a later decision, this is the coded material that is retrieved by the positivist decision maker who, in the case of a sophisticated positivist, uses this retrieved memory material to estimate the level of accuracy for the experienced outcomes (the proportion of hits out of all employed applicants) and to infer the level of accuracy for all presented applicants (the proportion of hits and correct rejections out of all employed and rejected applicants). In contrast, constructivist coding involves the coding of the inferred outcome for each decision made without feedback. At the time of a later decision, the memory material retrieved by the constructivist decision maker involves all presented applicants and the level of accuracy may be estimated directly from the retrieved proportion of hits out of all employed and rejected applicants.

A

Task Structure

Decision		Outcome: Positive	Outcome: Negative
Decision	"Yes"	Hits	False alarms
	"No"	Misses	Correct rejections

B

Positivist Coding

Decision		Coded Outcome	
Decision	"Yes"	Positive	Negative
	"No"	-	-

C

Constructivist Coding

Decision		Coded Outcome	
Decision	"Yes"	Positive	Negative
	"No"	*Negative*	*Negative*

Figure 3. Panel A: A schematic illustration of decisions and outcomes in a binary task. With selective feedback, contingent on positive decisions, no feedback is received for misses or correct rejections. Panels B and C illustrate the coded outcomes for positivist and constructivist coding. For sophisticated positivist coding with a symmetry assumption, feedback is coded in the positivist manner depicted in Panel B. This coded material is retrieved in later decisions and an inference based on symmetry of accuracy is made (see Equation 1).

Method

The effects of selective feedback were investigated in an experiment involving an investment task with complete and selective feedback with the basic design outlined earlier. The experiment was a 2 x 2 factorial design, with the two independent variables being complete vs. selective feedback contingent on positive decisions, and continuous vs. binary decisions and outcomes. Participants in the continuous condition predicted the outcome of presented companies on a scale from -50 to 50, where they received (complete or selective) feedback on the correct continuous outcome. In the binary condition, participants simply predicted whether the outcome would be positive or negative and they received binary feedback. A positive prediction (of an outcome ≥ 0) implied a decision to invest, which resulted in feedback independently of the feedback condition. A negative decision was followed by feedback only in the complete feedback condition, whereas with selective feedback no outcome information was presented for non-investment decisions. To enable comparison of the binary and continuous conditions, results were analysed in terms of the binary decisions of participants of whether to invest in the companies.

Predictions

Learning with complete feedback is expected to result in neutral decision behaviour, reflecting the equal base rates of positive and negative outcomes in the task. With selective feedback contingent on positive decisions, positivist and constructivist codings yield different predictions. Assuming some learning, the positivist coding hypothesis predicts that most experienced outcomes are positive. Naive positivist coding will therefore result in the perception of a high base rate of positive outcomes, which should be associated with a greater willingness to invest. Accordingly, naive positivist coding predicts liberal decision behaviour.

As specified in Equation 1, the decision behaviour predicted by sophisticated positivist coding is affected not only by the content of experienced outcomes but also by the proportion of investment decisions by the individual. A large proportion of positive outcomes (i.e. correct decisions) in memory is complemented with an inference that a large proportion of the unseen outcomes was negative (i.e. correct decisions). The perceived base rate of positive outcomes will therefore be lower with sophisticated, in comparison with naive, positivist coding and response behaviour less liberal, or even conservative.

With constructivist coding, internal feedback for unseen outcomes is constructed from knowledge and expectations. A negative expectation reflected in a negative decision results in no feedback and hence an internally generated negative outcome is coded. If performance is less than perfect, the participant will end up with an experience that contains a large proportion of negative outcomes (some of which were in fact positive, i.e. misses). Thus, constructivist coding of selective binary feedback is predicted to produce conservative decision behaviour, because of the perception of a high base rate of negative outcomes.

Previous research has demonstrated that feedback that provides information on a continuous scale, indicating, for example, the level of toxin in a bug, rather than merely binary information of the category label (poisonous or non-poisonous), allows the development of a more correct model of the task (Juslin, Olsson, & Olsson, 2003). It is therefore expected that continuous feedback on outcomes of the companies will result in higher performance independently of the feedback condition. Specifically, in the selective feedback condition, improved performance with feedback on a continuous scale should enable more correct inferences coded in the absence of feedback as compared with the binary condition. Consequently, a more correct perception of base rates and more neutral response behaviour, similar to that with complete feedback, are expected.

Results and Discussion

Study I investigated the effects of learning with selective feedback on decisions, as well as the coding of experiences without feedback. Results demonstrated that although experience was biased in the selective feedback conditions, selective feedback did not lead to impaired performance. In fact, performance was higher with selective as compared with complete feedback in this study. Subsequent experiments, however, reported in Studies II and III demonstrated that the typical finding involves slightly lower performance with selective as compared with complete feedback.

Figure 4A depicts the decision behaviour of the four conditions of the experiment. The decisions of participants with complete feedback are in line with their experience of an equal amount of successful and unsuccessful companies. Participants with binary selective feedback have a distinct tendency to judge a smaller proportion of companies as successful (.33), a tendency that is less extreme in the continuous selective feedback condition (.46). Base rates of positive outcomes in training as implied by naive positivist coding are .69 in the continuous selective condition and .76 in the binary selective condition, which are clearly inconsistent with the empirical data in Figure 4A.

Panels B and C of Figure 4 display participants' decisions as predicted by sophisticated positivist coding and constructivist coding. The experience of participants with selective feedback has been coded in line with each coding alternative and their test decisions are predicted from experienced and inferred base rates in line with the coding alternatives. For sophisticated positivist coding of selective feedback in Panel B, the predicted proportion of positive decisions at test are .48 in the continuous condition and .53 in the binary condition, which correspond poorly to the pattern revealed in the empirical data. In Panel C, however, constructivist coding of experience with selective feedback corresponds well to the qualitative pattern of outcomes in Panel A, with predicted proportions of investment decisions of .40 in the continuous conditions and .34 in the binary condition.

Further, at an individual level, a participant's experience as coded by naive positivist coding is uncorrelated with decision behaviour at test, and when adding the symmetry assumption, the correlation is negative. Constructivist coding of the participants' experience, on the other hand, is strongly correlated with decision behaviour at test (*r* = .85). In total, the support for constructivist coding against the suggested alternatives is strong in this study. In addition, an unpublished similar experiment with base rates of positive outcomes of .80 and .20, demonstrated the same pattern of results. Overall, decisions were responsive to the actual base rates: however, as predicted by constructivist coding, decisions after training with selective feedback reflected a consistent underestimation of positive outcomes.

In sum, the results of Study I demonstrated that participants' performance was not impaired by feedback being unrepresentative and contingent on their decisions, when compared with learning from the same number of feedback trials when received on every trial. Further, participants' judgements at test accurately expressed the frequencies in training as experienced with constructivist coding. In Study III, the effects of selective feedback on performance were investigated further and in Study II the constructivist coding hypothesis was additionally examined.

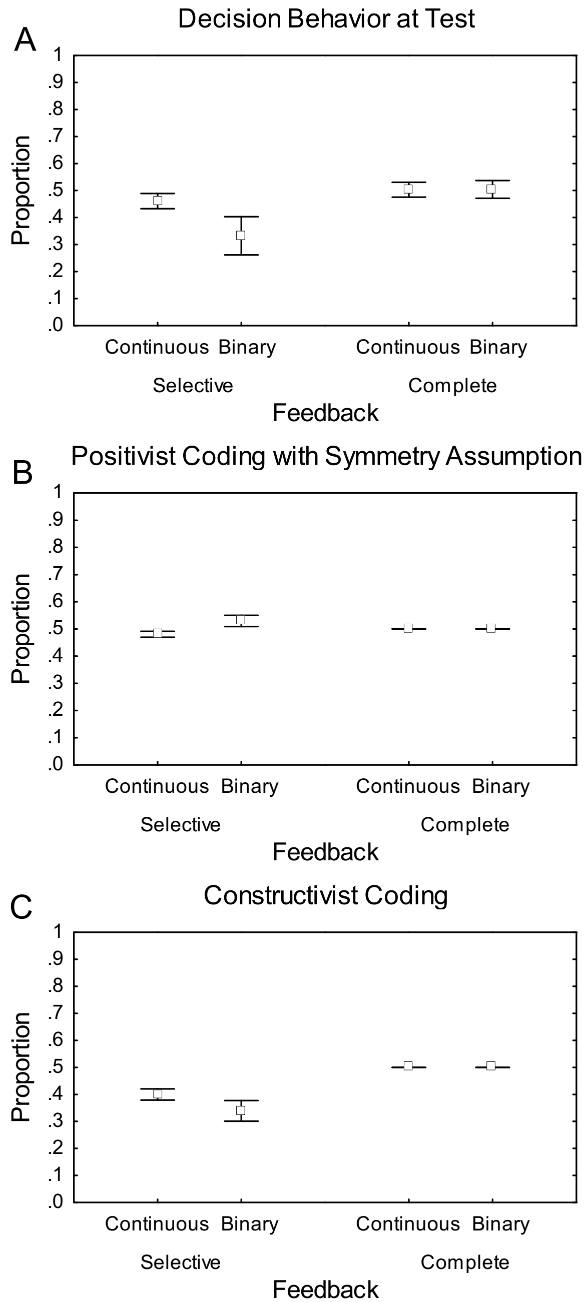

Figure 4. Panel A illustrates the observed decision behaviour in terms of the proportion of positive decisions with continuous and binary feedback as received selectively or completely for every decision. Panel B depicts participants' decision proportions as predicted by sophisticated positivist coding. In Panel C, the decision proportions as predicted with constructivist coding are given. Vertical bars denote .95 confidence intervals.

Study II. What is Coded Into Memory in the Absence of Outcome Feedback?

In Study II, the investigation of constructivist coding was extended in several ways. In addition to the investigation of participants' decision proportions as was done in Study I, the trial-by-trial judgements were analysed, including potential effects of perceived similarity relations between exemplars. The constructivist coding hypothesis was contrasted with two alternative coding schemes (positivist and agnostic coding) in three experiments involving selective feedback contingent on decisions or on outcomes. Study II thus allows a broader test of constructivist coding in relation to different sources of selectivity, where decision-contingent feedback depends on the actions of the individual while outcome-contingent feedback is governed by the structure of the environment.

Figure 5 shows the three coding alternatives in relation to selective feedback contingent on positive decisions, where Panels A to C are identical to Figure 3 in Study I. The possible outcomes in the task are outlined in Panel A, and with feedback selectively contingent on positive decisions, no feedback is received for misses or correct rejections. When outcome feedback is received (for hits and false alarms), all coding alternatives assume that the individual codes the outcome as presented. Positivist coding (Panel B) means that nothing is coded in the absence of feedback and constructivist coding (Panel C) means coding the expected outcome in the absence of external feedback.

The results in Study I suggested that when feedback is absent, people are likely to code an inference rather than code nothing at all. In Study I, this alternative was outlined in constructivist coding. Possibly, however, what is coded is not the expected outcome, but a randomly generated outcome, reflecting the ignorance of the participant regarding the unseen outcome. Like constructivist coding, *agnostic coding* involves some coding in the absence of feedback though without any assumption concerning the likely outcome. With agnostic coding, as illustrated in Panel D, a randomly generated outcome is coded when feedback is absent, i.e. each outcome (positive or negative) is coded with a .5 probability.

As outlined in the presentation of Study I, positivist coding of experienced outcomes (illustrated in Panel B of Figure 5) may be relied upon by the participant naively if later decisions are based solely on the retrieved material. The sophisticated positivist retrieves the same material, which is used with a symmetry assumption in which the level of overall accuracy is inferred from the proportion of correct decisions as experienced in outcome feedback (see Equation 1).

A

Task Structure

Decision		Outcome: Positive	Outcome: Negative
Decision	“Yes”	Hits	False alarms
	“No”	Misses	Correct rejections

B

Positivist Coding

Decision		Coded Outcome	
Decision	“Yes”	Positive	Negative
	“No”	-	-

C

Constructivist Coding

Decision		Coded Outcome	
Decision	“Yes”	Positive	Negative
	“No”	*Negative*	*Negative*

D

Agnostic Coding

Decision		Coded Outcome	
Decision	“Yes”	Positive	. Negative
	“No”	*.5 negative, .5 positive*	*.5 negative, .5 positive*

Figure 5. Panel A: A schematic illustration of decisions and outcomes in a binary task. With selective feedback contingent on positive decisions, no feedback is received for misses or correct rejections. Panels B to D illustrate the coded outcomes for the three coding schemes. For sophisticated positivist coding with a symmetry assumption, feedback is coded in the positivist manner depicted in Panel B. This coded material is retrieved in later decisions in which an inference is made about the overall level of accuracy (see Equation 1).

Figure 6 illustrates the three coding alternatives when feedback is selectively contingent on positive outcomes (Panel A presents the different outcomes of the task). Feedback is received contingent on positive outcomes if an investor consistently receives outcome information for any successful projects (whether or not she decided to support them) whereas no attention is given to projects with unsuccessful outcomes. In other words, when feedback is selectively contingent on positive outcomes, it is received only for hits and misses. All coding alternatives assume that this externally provided feedback is coded. With positivist coding of outcome-contingent feedback, nothing is coded when feedback is absent for false alarms or correct rejections (Panel B of Figure 6). Constructivist coding is illustrated in Panel C, where the absence of feedback for negative outcomes implies that outcomes are coded as positive or negative according to the actual decisions. Agnostic coding of a randomly generated outcome when feedback is absent is indicated by a probability of .5 for the coding of each outcome (positive or negative, see Panel D).

Again, positivist coding may be naive if later decisions are based solely on the retrieved material or it may be sophisticated if adding a symmetry assumption at the time of later decisions. In sophisticated positivist coding of outcome-contingent feedback, the base rate of positive outcomes is inferred from the level of accuracy experienced in outcomes with feedback (i.e. hits and misses), as specified in Equation 1.

In this task, participants are explicitly informed of the feedback contingency, and consequently feedback that is received contingent on positive outcomes is, in essence, highly similar to complete feedback. Because the participant is instructed that the absence of feedback always implies a negative outcome, feedback is either explicitly received or may be easily deduced. As we will see, however, results indicate that outcome-contingent feedback and complete feedback actually have quite different effects on performance.

Three alternative coding schemes were thus offered that describe what may be coded, if anything, in the absence of feedback that is received contingent on decisions or on outcomes. Like in Study I, each coding alternative (including the two applications of positivist coding) was evaluated against the actual proportion of positive decisions made by participants. This analysis concerns the overall judgements of participants. To evaluate the processes that produce the specific judgements made by the participants, the analysis is supplemented using a process model. Thus, each coding alternative was incorporated in the GCM (Nosofsky, 1986; Nosofsky & Johansen, 2000) and fitted to data from experiments involving selective decision-contingent or outcome-contingent feedback.

A

Task Structure

Decision		Outcome: Positive	Outcome: Negative
	"Yes"	Hits	False alarms
	"No"	Misses	Correct rejections

B

Positivist Coding

Decision		Coded Outcome	
	"Yes"	Positive	-
	"No"	Positive	-

C

Constructivist Coding

Decision		Coded Outcome	
	"Yes"	Positive	*Positive*
	"No"	Positive	*Negative*

D

Agnostic Coding

Decision		Coded Outcome	
	"Yes"	Positive	*.5 negative, .5 positive*
	"No"	Positive	*.5 negative, .5 positive*

Figure 6. Panel A: A schematic illustration of decisions and outcomes in a binary task. With selective feedback contingent on positive outcomes, no feedback is received for false alarms or correct rejections. Panels B to D depict the coded outcomes for the three coding schemes. For sophisticated positivist coding with a symmetry assumption, feedback is coded in the positivist manner depicted in Panel B. This coded material is retrieved in later decisions in which an inference is made about the overall level of accuracy (see Equation 1).

As outlined previously in this thesis, the GCM specifies the processes involved in categorisation decisions, describing categorisation as a function of both the frequency and the similarity of the exemplars in memory. Categorisation is depicted as a process of learning from outcome feedback by coding memory traces of concrete exemplars (e.g., job applicants) together with the category label (outcome) into memory. In later categorisation decisions, the coded exemplars are retrieved and the categorisation is made based on the similarity between the new exemplar and the previously experienced exemplars. The probability that a new exemplar will be classified as suitable increases with its similarity to previously encountered suitable exemplars as compared with its similarity to unsuitable exemplars (a slightly more elaborate description of the model is provided in Study II and further details are available in Nosofsky, 1986). In the categorisation of a new job applicant in the recruitment task, previously encountered suitable and unsuitable applicants will be retrieved from memory, and the new applicant is more likely to be categorised as suitable if the applicant is more similar to the previously encountered suitable applicants than unsuitable applicants.

As noted previously in the thesis, the GCM, as all models of categorisation or judgement, makes no assumptions about the coding of exemplars encountered without feedback. However, the model may be used as a tool in the evaluation of the predictions of the proposed coding schemes regarding the coding of non-feedback trials. This evaluation is achieved by incorporating each scheme in the GCM and comparing the predictions with the actual judgements of the participants. When outcome feedback is available, all three coding alternatives outlined above assume that the outcome is stored with the exemplar, as described in the standard GCM. However, the coding alternatives involve different assumptions about the coding of non-feedback trials. When the coding schemes are incorporated in the GCM, exemplars presented without external outcome feedback will be coded according to the particular coding scheme. When the participant encounters applicants for which no feedback is received, positivist coding means that no exemplars are coded into memory for these applicants. For constructivist coding, these exemplars will be coded with the most likely outcome, which is the outcome implied in the judgement of the participant. Finally, agnostic coding implies that for applicants encountered without feedback, a randomly generated outcome is coded (with a .5 probability on each trial). As described by the GCM, in later categorisation decisions all coded exemplars, whether coded with external feedback or inferred, are retrieved for similarity assessment.

The judgements of new exemplars based on their similarity to previously encountered exemplars, as coded by each scheme, are then evaluated against data. By fitting the parameters of the model to capture the similarity relations as expressed in the decisions of the participants, conditional on the experience specified by each coding scheme, the predictions of test decisions may be evaluated against the empirical results.

Method

Empirical data from three experiments with selective feedback were analysed to investigate which coding scheme best described the actual outcomes. The experiments were all based on the previously presented job recruitment task with binary judgements and feedback. Feedback in training was complete or selectively contingent on decisions, or on outcomes. The experiments involved some other manipulations additional to selectivity of feedback. *Feedback contingency* was positive or negative, involving feedback either for positive or for negative decisions (decision-contingent feedback), or for either positive or negative outcomes (outcome-contingent feedback). The *task focus* of participants was either to identify suitable applicants or unsuitable applicants. *Feedback attribution* involved feedback either attributed to the judge (in terms of correct vs. incorrect decision) or to the environment in terms of the category label (suitable vs. unsuitable applicant). None of these additional manipulations demonstrated any significant effect on decision behaviour or performance, which is why the experiments are presented and analysed together.

To test the constructivist coding hypothesis, the data in all conditions were coded so that a positive response criterion corresponds to conservative decisions with respect to the category overrepresented in feedback, and a negative criterion indicates liberal decision behaviour with respect to the category overrepresented in feedback. This means that a positive response criterion denotes the tendency to prefer to answer "unsuitable" if feedback involves only "suitable" (decisions or outcomes) and the tendency to answer "suitable" if feedback involves "unsuitable". For example, as seen in Study I, when the recruitment task involves feedback only for positive decisions, constructivist coding implies that the participants prefer to make negative decisions. In contrast, if the recruitment task involves feedback only for negative decisions, constructivist coding predicts that positive decisions will be preferred. To facilitate the presentation of the task, it will consistently be framed as a positive task focus on the recruitment of suitable applicants and feedback as received contingent on positive decisions or positive outcomes.

Predictions

The predictions from positivist, constructivist, and agnostic coding were compared using several analyses. First, the proportion of positive decisions in the test phase by each participant was predicted from her unique training history of suitable exemplars stored in memory according to each coding scheme. In this analysis, decisions are investigated as determined entirely by the coded proportion of outcomes, expecting that the coding of a large proportion of unsuitable applicants, for example, will lead to a decision behaviour that favours rejection decisions. This is the analysis performed in

Study I. As evident in Figures 5 and 6, constructivist coding predicts that the perceived base rate of positive outcomes with selective feedback contingent on positive decisions will be low (see Figure 5C), whereas the perceived base rate of positive outcomes with selective feedback contingent on positive outcomes will be high (Figure 6C). The predicted response criterion is therefore high (i.e., conservative) with decision-contingent feedback and low (liberal) with outcome-contingent feedback. The qualitative predictions from the alternative coding schemes are readily inferred from Figures 5 and 6 (for a detailed presentation of predictions, the interested reader is referred to Study II).

The examination of decision proportions has the advantage of not involving the fitting of any free parameters to data, which minimises the risk of overfit. To analyse decisions based on similarity as well as on frequency, the coding schemes are compared by incorporating each scheme in the GCM and fitting each version to data at the group and individual level. Constructivist coding is expected to demonstrate the smallest deviations between model predictions and empirical data.

Analyses and Results

Decision Behaviour

To investigate decisions based on the coded proportions of suitable and unsuitable applicants in training, the base rates of suitable exemplars stored in memory conditional on each type of coding were compared with the base rate of positive decisions at test. The constructivist coding hypothesis predicts that with selective decision-contingent feedback, the participants will become conservative regarding the category that is overrepresented in feedback; with selective outcome-contingent feedback, they will become liberal with respect to the category overrepresented in feedback. Consequently, in a recruitment task the perceived base rate of positive outcomes when feedback is contingent on positive decisions should be low (see Figure 5C), whereas the perceived base rate of positive outcomes when feedback is contingent on positive outcomes will be high (Figure 6C).

As predicted by the constructivist coding hypothesis, the present results consistently demonstrated conservative response behaviour relative to the category that was overrepresented in decision-contingent feedback and liberal response behaviour relative to the category overrepresented in outcome-contingent feedback. These results are illustrated in a recruitment task in Figure 7 as conservative response behaviour when feedback was contingent on positive decisions and liberal response behaviour when feedback was contingent on positive outcomes. As outlined earlier, the results summarised in Figure 7 involve different patterns of feedback contingencies, which suggest that the perceived base rates in the same task depend upon, and may be changed with, feedback contingency. For example, when feedback is received only for positive decisions, the participants prefer to reject most applicants, whereas when feedback is received only for negative decisions, they prefer to recruit the majority of applicants. This effect will be discussed later in this thesis. Moreover, Figure 7 illustrates the profound difference between the two conditions with selective feedback, also supporting the predictions of a constructivist coding.

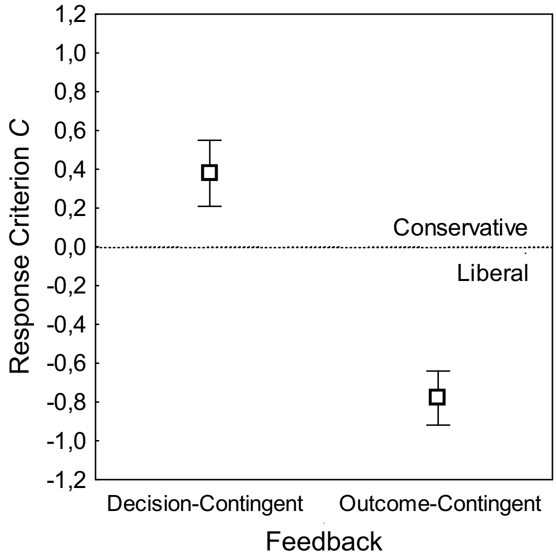

Figure 7. The response criterion C for all experiments in Study II framed as a recruitment task with feedback for positive decisions or positive outcomes. Vertical bars denote .95 confidence intervals

Further, the training experience of each participant, as coded according to each of the three coding schemes, was used to predict the overall proportion of positive decisions at test. Figure 8 depicts the mean *RMSD* between the predictions by each coding alternative and the empirical data. As evident in Figure 8, constructivist coding provided the best predictions of data for both decision-contingent and outcome-contingent feedback. In addition to the best predictions of the specific decision proportions, the constructivist coding alternative demonstrated the highest correlations between the training experience of the participants and their decision proportions at test in both feedback conditions.

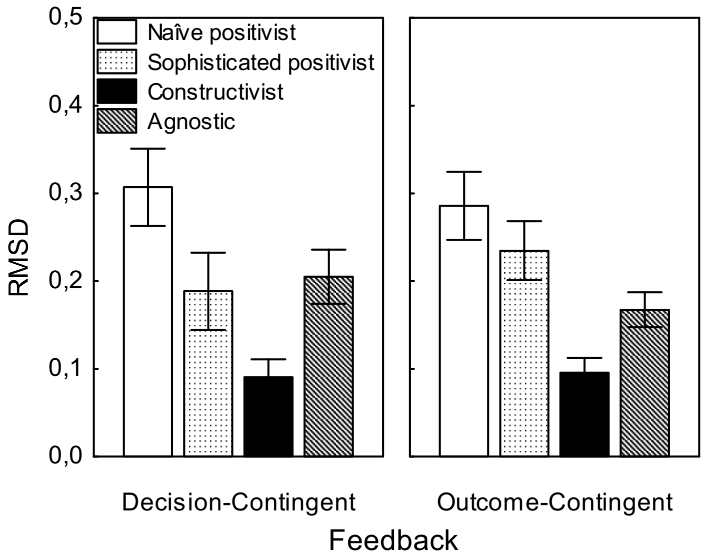

Figure 8. Prediction error as Root Mean Square Deviation (*RMSD*) between the predicted and the observed proportion of positive recruitment decisions for each individual participant in the conditions with decision-contingent and outcome-contingent feedback, for each of the three coding alternatives. Whiskers denote .95 confidence intervals. The lower the *RMSD*, the better the prediction of the participants' proportion of affirmative decisions at test.

Similarity-based Decisions

To evaluate the predictions by the three coding schemes regarding the similarity-based judgements of the participants, the GCM with each coding alternative was fitted to group test phase data. This analysis concerns the predictions of the item-by-item responses made by participants in the test phase, where the fitted parameters for similarity in the GCM take into account, and thus allow for, the similarity relations between the exemplars as the participants subjectively perceive them. The training exemplars provide the memory base of stored exemplars as contingent on each of the three coding schemes. The parameters of the GCM were fitted to minimise *RMSD* between the decision proportions predicted by the GCM with each coding scheme and the observed decision proportions for each test exemplar. In this way the fitted parameters capture the perceived similarity relations between exemplars.

Because this analysis involves the predictions of responses at the level of individual items, the sophisticated application of positivist coding is not included. This is because the sophisticated positivist makes a *general* symmetry assumption (involving symmetric accuracy for non-feedback exemplars as for exemplars experienced with feedback), and does not specify predictions at the level of individual responses.

As shown in Figure 9, constructivist coding demonstrated the best model fit for decision-contingent feedback and for outcome-contingent feedback (though *RMSD* was not significantly lower as compared with agnostic coding in the latter condition). Thus, at the group level, constructivist coding was the best alternative for predicting participants' test judgements as based on the similarity to previously experienced exemplars.

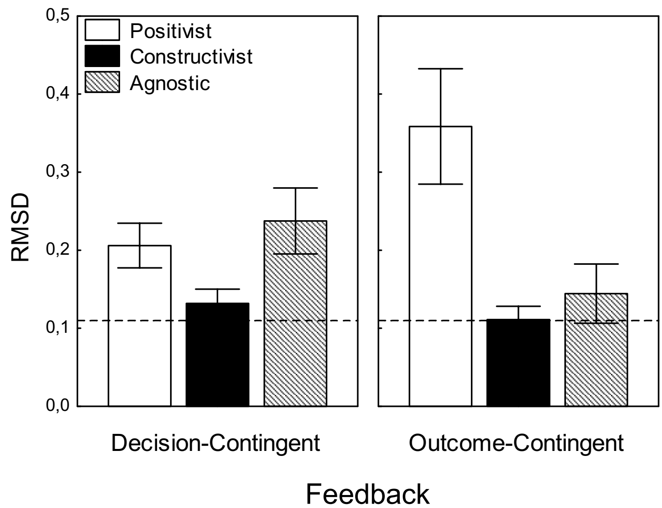

Figure 9. Model fit in terms of *RMSD* of the GCM implementing each of the three coding alternatives as applied to the empirical data. Whiskers denote .95 confidence intervals. The dotted horizontal line is the average standard error of measurement for the predicted data points, indicating a ceiling on the variance that can be validly predicted. Comparisons of the standard error and the obtained *RMSD* suggest that the variance not accounted for in model fit is due to unsystematic variance or noise in the data and that constructivist coding can actually explain most of the systematic variance in the data.

Further, to verify that the same rank order between the models is observed also in the analysis of individual participants, the GCM with each type of coding scheme was fitted separately to the data from each participant. Notably, the GCM is not ideal for application to individual participant data in the current data set. This is because it involves binary classifications of exemplars presented only once rather than response probabilities on a continuous scale from 0 to 1. Even if the model correctly describes the response probabilities, these will only materialise with repeated decisions. The analysis on individual data was nevertheless included to verify the pattern of results in the analysis at the group level. Indeed, model fit analyses on individual data revealed the same pattern as the analysis at the group level, with best fit for constructivist coding in both feedback conditions.

Performance

Finally, performance in the different feedback conditions was analysed. Some learning occurred independently of feedback condition. Performance was highest with complete outcome feedback (mean d′ 2.07), followed by decision-contingent feedback (d′ 1.59), and outcome-contingent feedback (d′ 1.19). The difference in performance for complete and decision-contingent feedback was non-significant. Again, we see how decision-contingent feedback does not appear to hamper learning to any large extent. However, outcome-contingent feedback was associated with impaired performance.

Brief Summary and Discussion

In summary, the results in Study II clearly suggest that participants code non-feedback trials constructively, according to their beliefs about the likely outcome. Constructivist coding predicted the decision behaviour of participants with selective feedback contingent on decisions and on outcomes, as well as the judgements of the participants as based on the similarity between new and previously encountered and coded exemplars. Thus, constructivist coding systematically proved the best alternative in Study II.

Most of the analyses in this study relied on the GCM in the evaluation of the coding schemes. As noted earlier in the thesis, cognitive models of categorisation often involve exemplar processes, and the GCM is one of the most widely used exemplar models. Alternative models may rely on abstracted information in terms of rules (e.g., Juslin, Olsson, & Olsson, 2003) or prototypes (e.g., Smith & Minda, 2000) rather than concrete experiences, and such models could also capture the processes involved in categorisation learning. As noted, however, a model that relies on the coding of exemplars naturally preserves information of the experienced frequencies. This is an interesting aspect of learning from selective feedback, where the expressed proportions of outcomes may be used to investigate the coding of non-feedback exemplars. For instance, in the abstraction of prototypes it is less evident how information of the proportion of suitable applicants is retained in memory. A model that relies on the storing of exemplars, such as the GCM, may explain straightforwardly why participants with decision-contingent feedback act as if the proportion of suitable candidates is low and participants with outcome-contingent feedback act as if the proportion of suitable candidates is high. The present results demonstrate that these decision tendencies are best predicted by the proportion of encountered exemplars in training as coded by constructivist coding. Moreover, as shown in

Figure 9, the GCM with constructivist coding actually accounts for most of the true variance in these data, suggesting that the GCM is an appropriate model for the analyses.

Study III. Living and Learning: The Interplay Between Beliefs, Sampling Behaviour, and Experience

In Study III, the effects of learning with selective decision-contingent feedback on decisions were investigated in three experiments. From a sampling perspective, impaired performance with selective feedback may be explained as the effect of a selective sample of experiences that is correctly, but naively, interpreted by the decision maker. In this perspective, when the aim of the decision maker is to produce good outcomes, restricted sampling behaviour may lead to an objective selective experience that actually reproduces and supports an incorrect prior assumption. Because the decision maker correctly interprets the specific sample of experiences but fails to acknowledge that the content of this experience is biased, selective experience may lead to the persistence of inaccurate beliefs.

The predictions from a sampling approach were investigated by offering an incorrect initial assumption to the participants, before their experience of the task environment with complete or selective feedback on decisions. Because the initial assumption was inaccurate in relation to the actual outcomes in the environment, the persisting influence of the instruction on the beliefs of the participants after extensive experience could be evaluated. The effects of the specific sample of experiences on the beliefs of participants were assessed by analysing the objective experience of participants with selective feedback. A sampling explanation of biased beliefs suggests that the persistence of an incorrect assumption should be larger with selective feedback to the extent that the instruction was reproduced in experience. Further, by manipulating the sampling behaviour of the participants, the causal effects of the characteristics of the sample on the beliefs of the participants could be investigated. A sampling approach suggests that less conservative decision behaviour with selective feedback and a sample of experiences that is more representative of the environment should enable the revision of an incorrect prior belief.

In addition to the investigation of the confirmation of beliefs in objective and selective experience, Study III offered the opportunity to test a prediction from constructivist coding. With constructivist coding, the manipulation of the restrictiveness of decision behaviour of participants should lead to differences in the perceived proportion of suitable applicants. Results reported in Studies I and II demonstrated conservative decision behaviour with selective feedback contingent on positive decisions, indicative of a low perceived proportion of suitable applicants. In Study III, these analyses were supplemented with the investigation of the base rates as explicitly assessed by the participants. This additional analysis was done to verify that the decision behaviour reflects the beliefs held by the participants concerning the actual proportion of positive outcomes.

Method and Predictions

Effects on decisions of selective, in comparison with complete outcome feedback, were investigated in the basic design outlined earlier. As in previous studies, in the first experiment of Study III the participants did not have any previous knowledge on which to base their decisions, i.e. they had to learn from experience alone. Under these circumstances, selective feedback was predicted to lead to a slight impairment in performance.

Predictions from a sampling approach were tested in Experiments 2 and 3. The predictions propose that the influence of an initial incorrect assumption may persist because participants with selective feedback act conservatively on the expectation, whereby their actual experience confirms the incorrect assumption. Participants were provided with either of two opposite assumptions that were both incorrect in relation to the environment that the participants experienced through complete or selective feedback. Both assumptions specified two of the characteristics of the job applicants as the only relevant features for the outcome, making known that the values on these characteristics should be high (in fact all four features were of equal relevance).

In Experiment 3, the sampling behaviour of participants was manipulated by inducing a focus on short-term performance vs. on long-term learning. This manipulation enabled the investigation of causal effects of the characteristics of the selective sample on the persistence of incorrect beliefs. With a task focus on "Performance" objectives, the incorrect assumption was stated as a fact and participants were instructed to perform to the best of their ability. This condition was expected to be associated with conservative decision behaviour aiming to avoid negative outcomes. In contrast, a "Learning" focus was induced by offering the assumption as a hypothesis to be tested in order to make correct judgements, thereby generating less conservative decision behaviour and a greater willingness to employ.

The predictions from the sampling approach were investigated in two steps. First, the persisting influence of the incorrect instruction after extensive experience with complete or selective feedback was analysed. The beliefs of participants were investigated by analysing the characteristics of the applicants recruited after 200 trials of experience. The persisting effect of an incorrect instruction was indicated by a larger emphasis on the cues that were specified as important in the instruction as compared with the other cues. The second step involved the actual experience of participants with selective feedback. The characteristics of experienced suitable and unsuitable applicants with selective feedback were investigated in an effort to evaluate whether the objective selective experience of participants confirmed the instruction.

The sampling approach predicts that the persisting influence of an incorrect prior assumption on the beliefs of the individual should be larger when experience is made with selective feedback, as compared with complete feedback, to the extent that the objective experience of the participants confirms the initial assumption. Conservative decision making on the part of the participants with a Performance focus should result in an objective experience with selective feedback that confirms the initial assumption. In contrast, in the Learning condition a more liberally collected experience is likely to be more representative of the actual outcomes and should thereby lead to the revision of the initial hypothesis and beliefs that correspond more to the actual structure of the environment.

Finally, the prediction from constructivist coding was evaluated by asking the participants in all conditions to estimate the actual proportion of suitable applicants. Constructivist coding suggests that because non-feedback trials are constructively coded with a negative outcome, conservative decision behaviour associated with the experience of few positive outcomes in selective feedback will lead to the estimation that the actual proportion of suitable outcomes in the environment is low. In contrast, when feedback is selective and the decision behaviour is neutral, participants experience a larger number of positive outcomes and fewer non-feedback exemplars which are constructively coded with a negative outcome. Therefore, the estimated proportions of positive outcomes should be lower in the Performance condition with selective feedback as compared with the Learning condition with selective feedback, the latter generating estimations closer to those given by participants with complete feedback who perceive all outcomes independently of their decisions.

Analyses and Results

As in Study II, the negative effects of selective decision-contingent feedback on performance and overconfidence were rather small. In the first experiment of Study III, in which participants learned from experience alone, performance was high independently of feedback condition, with a mean detectability, d', of 3.09 for complete feedback and 2.54 for selective feedback.

Experiment 2 served to investigate whether the negative effects of selective feedback on performance, though small in themselves, could be explained by an incorrect prior assumption being confirmed in the actual objective experience of the individual. The main results of Experiment 2 were replicated in the Performance condition of Experiment 3, which will be outlined in more detail. Experiment 3 further investigated the effects of incorrect instructions by a manipulation of decision behaviour via task focus. As noted, the predictions from the sampling approach were investigated in two steps. First, the persisting influence of the incorrect instruction after extensive experience was analysed in terms of a larger emphasis on the cues that were specified as important in the instruction. As expected, results demonstrated a larger persisting influence of the inaccurate instruction on the beliefs of the participants who received selective feedback as compared with participants who received complete feedback. Also, the influence of instructions was larger when focus was on performance than on learning.

The second step of analysis involved the actual experience of participants with selective feedback. Specifically, the analysis evaluated whether the objective selective experience of participants confirmed the incorrect instruction. The characteristics of experienced suitable and unsuitable applicants with selective feedback are presented in Figure 10 for the Performance and Learning conditions of Experiment 3. As evident in Panel A, the characteristics of the experienced applicants in the Performance condition were clearly consistent with the instruction specifying that the values for the first and third cue should be high. In this condition, the values for these cues were consistently very high, especially for the experienced suitable applicants, indicating that the experience of these participants confirmed the instruction. The characteristics of the experienced suitable and unsuitable applicants indicate that the participants in the Performance condition systematically relied on the instruction in their decisions, avoiding applicants with lower (or even medium) values on the told important cues. This condition produced an experience that supported the employing of applicants with very high values on the told important cues and the avoiding of applicants with lower values. As clearly observed, the decision behaviour of participants with selective feedback and a performance focus was highly conservative (mean C was 0.57, as compared with -0.01 in the Learning condition). As can be seen in Panel B, also the Learning condition produced an experience in line with the instruction though to a much lower extent. Suitable applicants were characterised by higher cue values for the cues specified in the instruction as compared with the other cues, although these values were obviously less extreme than in the Performance condition.

Predictions from a sampling approach were thus supported by a larger persisting influence of the incorrect instruction with selective feedback (as compared with complete feedback) when the objective experience of participants with selective feedback confirmed the instruction. Therefore, the persistence of an inaccurate instruction may be linked to the confirmation of an incorrect assumption in objective (however selective and restrictive) experience. When the decision maker has a strong expectation and focuses on producing good outcomes, the sample of experiences may result in the confirmation and persistence of the incorrect assumption. In contrast, when the participants were motivated to sample more liberally to test a weaker expectation or hypothesis, their samples of experiences became more representative of the environment. The beliefs of participants in the Learning condition accordingly came to reflect the structure of the environment more than the instruction.

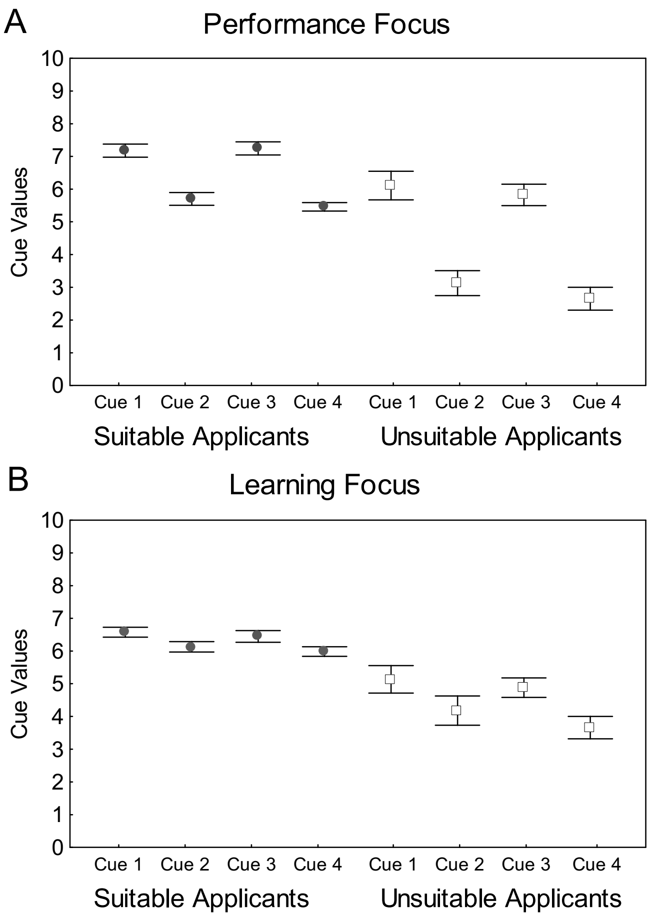

Figure 10. Actual experience with selective feedback in Experiment 3 of Study III. The graphs present the mean cue values of all suitable and unsuitable applicants experienced with feedback, when the participants are instructed that the first and third cue should be high in the Performance condition (Panel A) and the Learning condition (Panel B). Vertical bars indicate .95 confidence intervals.

Estimated Proportions and Constructivist Coding

The explicitly estimated proportions of suitable applicants in Experiment 3 are illustrated in Figure 11 for participants with complete or selective feedback and with a task focus on Performance or Learning. As seen in the graph, the highly conservative decision behaviour of participants with selective feedback in the Performance condition was associated with an estimated actual proportion of suitable applicants that was lower than in the other conditions. The neutral decision behaviour of the Learning condition with selective feedback was associated with estimations of the proportion of suitable applicants at the level with that of participants with complete and representative feedback.

This pattern of results was predicted by the constructivist coding hypothesis, which suggests that when no external feedback was received, internal feedback of the expected negative outcome was generated and coded. When the participants reject a high proportion of the applicants, they come to believe that the proportion of suitable applicants is low. By the same logic, a liberally collected experience with selective feedback in the Learning condition was associated with estimated proportions of suitable applicants at the level with complete feedback.

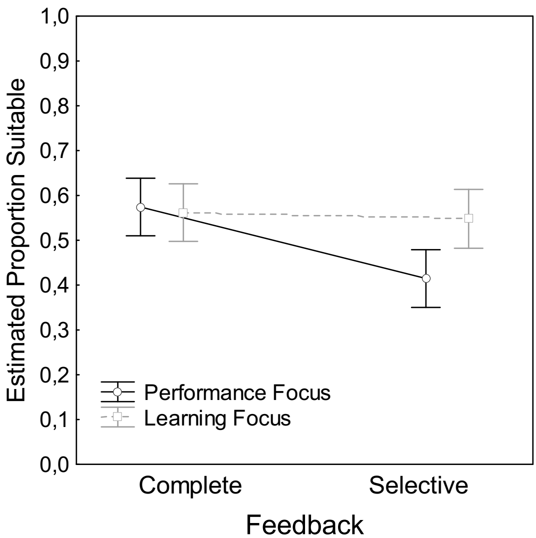

Figure 11. Estimated proportion of suitable applicants in Experiment 3 of Study III. Vertical bars indicate .95 confidence intervals. The "baseline" estimation of the proportion of suitable applicants with complete feedback and with selective feedback and a Learning focus is higher than .5. In itself, this effect is not essential for the questions addressed in this thesis. The reader may turn to Study III for a brief discussion of this effect.

Brief Summary

Results in Study III revealed how individuals with strong, inaccurate expectations and the aim to perform well may produce an objective (however restrictive and selective) experience of the environment that actually confirms the expectations, which results in the persistence of inaccurate beliefs. Consequently, even with extensive experience of the same environment, participants with opposite prior assumptions consistently decided to employ applicants with very different characteristics. The effect is explained by the reliance on the objective sample of experiences, correctly but naively interpreted "at face value", as predicted by a sampling approach.

The sampling approach makes no predictions regarding the perceived proportion of suitable applicants. By the coding also of non-feedback exemplars, the constructivist coding hypothesis may supplement the sampling approach in explaining the perceived proportions. In this manner, sampling effects on learning may be enhanced by the participants constructively coding unseen outcomes in accord with their expectations, which would further augment the dissimilarities in knowledge formed by participants with different expectations.

General Discussion

Three major findings were observed in relation to the aims of this thesis. First, the effects on performance of selective feedback contingent on decisions were small. Second, negative effects of selective feedback may, to a significant extent, be explained by the naive interpretation of a selective and biased sample of experiences. The extent of decision bias thus depends on the representativity of the experienced sample. Third, when feedback is selective, participants' learning appears to be based not only on externally presented information, but also on the coding of internally generated feedback by means of constructivist coding. These findings will now be discussed further.

Effects of Selective Feedback on Performance

A growing amount of empirical data, including those presented in this thesis, demonstrate that selective feedback contingent on decisions does not automatically lead to highly impaired learning and poor performance. In fact, reported negative effects of selective feedback tend to be small (Fazio et al., 2004; Griffiths & Newell, 2009; Grosskopf et al., 2006; Yechiam & Busemeyer, 2006). Importantly however, the effects of selective feedback are likely to depend on the structure of the specific environment and the interaction with other factors. More complex nonlinear and probabilistic tasks offer a poorer basis for learning from outcome feedback (e.g., Brehmer, 1980; Hammond et al., 1973; Jacoby et al., 1984; Kluger & DeNisi, 1996; Remus et al., 1996). Thus, perhaps selective feedback could result in an even larger impairment in performance than does complete outcome feedback in such tasks. On the other hand, under these circumstances selective feedback could possibly mitigate negative effects of complete outcome feedback when participants' performance and consistency in decisions benefit from their not being exposed to frequent negative feedback.

The effects of selective feedback on performance were investigated in the present series of experiments that involved different amounts of training trials. In Study I, the training phase consisted of 120 or 240 training trials for complete and selective feedback, respectively. The aim was to compare learning with complete and selective feedback with approximately the same number of feedback trials. Selective decision-contingent feedback in Study I was associated with slightly higher performance than complete feedback. In Study II, all participants received 240 training trials independently of condition, with the effect that participants with selective feedback experienced about half the amount of feedback trials as participants with complete feedback. Performance with selective decision-contingent feedback in Study II was slightly lower than with complete feedback. From this pattern, a plausible conclusion appears to be that the same number of feedback trials leads to performance with selective feedback at a level with, or higher than, performance with complete feedback. However, the same slight advantage for complete feedback seen in Study II was also demonstrated in Study III (Experiment 1) in which the number of training trials was 160 (complete feedback) and 320 (selective feedback). It therefore appears that the slightly higher performance with complete feedback in Study II is not necessarily the result of additional feedback trials. Because a large part of the learning appears to take place in the first part of training, participants may actually not benefit much from the additional feedback trials.

Importantly, the reason for these variations in design (i.e. the number of feedback trials was matched in Studies I and III and the total number of trials was matched in Study II) was the findings of good performance with selective feedback. The high performance with selective decision-contingent feedback may be an effect of constructivist coding and the use of information in feedback to generate inferred outcomes on non-feedback trials, which benefits learning. From this perspective, there appears to be insufficient reasons to administer additional training trials to participants in the selective feedback conditions.

The results in Study II indicate that learning is hampered when feedback is selectively contingent on outcomes. This finding may appear surprising considering that participants are instructed that the absence of feedback always signals the outcome opposite to the one received in explicit feedback. If feedback is received only for suitable applicants, the absence of feedback always implies that the applicant was unsuitable. Under these circumstances, outcome-contingent feedback is highly similar to complete feedback. However, feedback received only for suitable applicants requires the appreciation of the contingency and an active inference on non-feedback trials in order to generate valid representations of unsuitable applicants. Therefore, if the individual does not attend to the contingency on every trial or fails to apply the insight to her decisions, the highly skewed explicit feedback offers a poor basis for learning. To the extent that people are limited in this regard, the fact that decision-contingent but not outcome-contingent feedback provides explicit feedback for both suitable and unsuitable applicants may explain the counter-intuitive finding that performance was poorer with outcome-contingent feedback.

The differences in performance and response behaviour between decision-contingent and outcome-contingent feedback may thus be understood in terms of differences in the skew of external feedback received by the participants. When feedback is received contingent on positive decisions, the participant receives explicit feedback for both outcome categories (hits and false alarms). When feedback involves both suitable and unsuitable applicants, the participant may more easily develop the ability to differentiate between the categories, even without an appreciation of the feedback contingency. The resulting representations may be used as the basis for valid rejections to which no external feedback is received and the stored "internal feedback" on non-feedback trials will often be correct.

As will be discussed further in the following section, the results presented in this thesis suggest that negative effects of selective feedback are present to the extent that the specific sample of experiences is restricted and biased. Accordingly, negative effects of selective feedback are likely to be larger in circumstances where the sample of the individual is restrictive and biased as compared with complete feedback. When the sample of experiences contains information more useful for generalisation, the development of a more accurate mental model of the task is possible. Such a model allows the generation of more accurate "internal feedback" for non-feedback experiences.

Sampling Effects in Learning From Selective Feedback

A sampling perspective on learning emphasises the role of the actual experience of the individual in the development of knowledge (Fiedler, 2000; Fiedler & Juslin, 2006). Biased knowledge and decisions are the result of a naively generalised biased sample of experiences. Study III demonstrated how biased decisions may be explained by an incorrect assumption being reproduced in actual experience with selective feedback. When sampling is more liberal and the sample therefore more representative of the actual outcomes, an incorrect assumption is corrected by selective feedback. Bias is thus a combined effect of the validity of a previous assumption, the restrictiveness of the decision behaviour, and the selectivity of feedback.

Studies II and III demonstrate how participants who encounter the same basic recruitment task and act in essentially the same environment may develop different perceptions as to the proportions of outcomes and the relationships that govern the environment. The specific experiences of individuals differ because of the selectivity of feedback. The information in the sample is generalised as valid in future decisions, most likely complemented with constructivist coding of outcomes for non-feedback exemplars.

An implication of the results in Study II is that the perceived base rate of suitable applicants may be changed in the same task just by altering the category for which feedback is received. When participants receive outcome feedback only after decisions to employ an applicant, they develop a bias to believe that most applicants are unsuitable. However, if the feedback contingency is changed so that they only receive feedback after decisions to reject an applicant, they develop a bias to believe that most applicants are suitable, even though they operate in the same environment. Likewise, if feedback is received for only one category of outcomes (e.g., suitable applicants), the participant will perceive the base rate of that category as high. Clearly, the knowledge of the individual is determined not only by the specific characteristics of the environment but also by the principles that govern the availability of feedback.

The sampling approach used in Study III may explain how two personnel officers with different expectations regarding the importance of various characteristics of the job applicants may each have their assumptions confirmed by objective experience, even though they are both incorrect. In their efforts to make good decisions their experience may confirm prejudices, superstition or simply mistaken assumptions. Similar reasoning may be applied to the absence of improvement in performance after extensive experience that has been repeatedly demonstrated in clinical judgement (e.g., Camerer & Johnsson, 1997; Dawes, 1994; Garb, 1989; Grove, Zald, Lebow, Snitz, & Nelson, 2000). The focus of the clinician is likely to involve the outcomes for individual patients rather than long-term learning. Therefore, if the clinician relies on inaccurate theories or assumptions, experience may be inadequate for adjusting mistaken beliefs. Instead, erroneous beliefs or speculations on the part of clinicians, or any individual who wants to achieve good outcomes, are likely to be reproduced in actual experiences that confirm the inaccurate assumptions.

The results in Study III demonstrate the inevitable trade-off, or balance, between immediate performance and learning when information is contingent on decisions. Individuals must rely on the outcomes of their decisions for *current* achievement and survival, as well as for learning in order to benefit *future* achievement and survival. Decisions may often serve one of these objectives at the expense of the other.

Moreover, the sampling effects in learning from selective feedback appear to be enhanced by the individual constructively coding unseen outcomes in line with her expectations, which further augments the potential dissimilarities in knowledge formed by different individuals acting upon the same environment. Note that these effects may occur entirely without any distorted perceptions or deforming cognitions. They are the effects of a selectively experienced sample of information, correctly but naively interpreted by the judge, possibly complemented with constructivist coding of inferences when external information is absent. When the decisions of the individuals determine their experience, it is obvious that people actively (though most likely unknowingly) take part in the creation of the world they expect.

Constructivist Coding of Unseen Outcomes

As outlined in the introduction, constructivist coding is consistent with a large and well-established body of results in cognitive psychology that describe constructive processes in a wide array of tasks (e.g., Bartlett, 1932; Dawes, 1994; Gregory, 1970; Hawkins & Hastie, 1990; Loftus & Hoffman, 1989; Roediger et al., 2001; Samuel, 1981). Further, research that demonstrates people's limited abilities for source (or reality) monitoring (e.g., Johnson, 2006; Johnson et al., 1993; Johnson & Raye, 1981; Roediger & McDermott, 1995; Roediger et al., 2001; Walther et al., 2003) suggests that once an inference or decision has been made, the individual is often unable to mentally separate inference from externally presented feedback.

In the hypothesis of constructivist coding, the constructive processes known from memory research are applied to understand the coding in memory of selective feedback. Constructivist coding is therefore offered as a plausible, or even necessary, consequence of the constructivist nature of the human mind. Actually, in order not to be susceptible to constructivist processes, people would need to display an outstanding ability for source-monitoring, distinguishing applicants in memory that were only believed to be (un)suitable from applicants for which the belief was reinforced by external feedback.

The coding of inferences when external information is lacking offers a straightforward explanation of the base rates as perceived by participants with selective feedback. When feedback is contingent on decisions, the individual perceives the base rate of a category as opposite to the actually experienced proportion. That is, with feedback only for positive decisions, the perceived proportion of positive outcomes is low, even though the objective experience contains a large proportion of positive outcomes. With some level of performance and restrictive decision behaviour, participants with feedback contingent on positive decisions will encounter a high proportion of suitable applicants in external feedback. With constructivist coding of unseen outcomes, this externally presented information is complemented with internally constructed negative outcomes, which are coded into memory with external feedback. Consequently, the perceived base rate of suitable applicants will be low. More liberal sampling on the part of the participants offers fewer non-feedback experiences for which internally constructed negative outcomes are generated. Under these circumstances, the estimated base rate of suitable applicants is higher.

The sampling effects in learning outlined in the previous section may thus be supplemented by constructivist coding of non-feedback exemplars. When feedback is contingent on positive decisions, as in Study III, all non-feedback exemplars are coded with the inferred and expected negative outcome. Accordingly, the validity of negative decisions is confirmed by internal feedback. When inferences are imperfect, the confirmation of invalid inferences may result in overconfidence. Such effects illustrate the perils of constructivist coding and dealing with the absence of feedback by assuming that one is correct. Feedback, in fact, can be expected to be absent for a large part of people's decisions in everyday environments (Hogarth, 2006). Assuming correct knowledge each time one is not informed otherwise may lead to insufficient updating of false beliefs and overconfidence. Such effects may well contribute to the confirmatory behaviour, biases and prejudice often reported in research on intuitive judgement and in social psychology (Allport, 1954; Brehmer, 1980; Dawes, 1994; Denrell, 2003, 2005; Denrell & Le Mens, 2007; Einhorn & Hogarth, 1978; Fazio et al., 2004; Fiedler, 1991, 1996, 2000; Hogarth, 1981; Klayman, 1988).

On the other hand, using experience to make "informed guesses" when external feedback is absent may be an adaptive ability in an environment in which feedback is scarce. Like the statistical imputation procedures used by scientists to exploit information in existing data to infer (or impute) missing data, the everyday decision maker may use existing knowledge in constructivist coding to support the formation and maintenance of a coherent representation of the environment. In fact, it appears that the relevant question is not *whether* people rely on constructivist coding. Instead, we should investigate the specific characteristics, effects, and extent of the coding of inferences in the absence of external information.

Specifications of Constructivist Coding

So far, constructivist coding has been defined in general terms as the coding of expectations or inferences in the absence of feedback, without much elaboration of details. Several questions concerning the specific aspects of such a process await further investigation in which different types of constructivist coding may be outlined and compared. In this thesis, constructivist coding has been discussed and investigated in relation to binary outcomes, where an unseen outcome is constructed as a negative (or positive) outcome. With richer knowledge concerning the relationships between features and outcomes (e.g., between the characteristics of applicants and suitability), internal feedback may possibly be constructed on a continuous scale. Based on their knowledge, people may thus infer and code not only that a rejected applicant with no external feedback was unsuitable, but the coded inference may also specify the level of unsuitability.

Further, there are possible varieties of, or alternatives to, constructivist coding as outlined in this thesis. Constructivist coding has been presented as involving the belief of the decision maker that she makes a correct decision each time no feedback is received. Possibly, the level of confidence in a decision made without feedback impacts on the coding of the outcome. Decisions associated with strong beliefs may thus be constructively coded in the same manner as actual outcomes, whereas mere guesses are less likely to be coded with internal feedback.

On the other hand, the act of making a decision may be the crucial element in the process, where the additional attention or cognitive processing associated with the performing of an action results in a more vivid memory trace, independently of the level of confidence in the decision. The experiments presented in this thesis include well defined decision trials on which feedback is present or absent. In contrast, everyday decisions may often be made without any attention to the fact that alternative options were rejected. For instance, many products in a grocery store are never even considered. Presumably, an alternative has to be attended to in some manner for constructivist coding of internal feedback to occur.

A "sophisticated" version of constructivist coding may involve a later inference similar to the one described for the sophisticated positivist. The expected outcome is coded on non-feedback trials; however, later decisions are adjusted according to the externally experienced proportion of correct decisions. If a large proportion of externally received feedback reveals incorrect decisions (e.g., false alarms), later decisions may accommodate this knowledge and result in less reliance on internal feedback. To enable this later adjustment, however, the decision maker would need to somehow differentiate in memory the stored experiences from external and internal feedback, which may require an implausibly high level of source monitoring ability.

Constructivist coding, as outlined in this thesis, consists of the assumption that external feedback, when present, "overrules" the inference that leads to a judgement. For instance, if the participant infers a positive outcome and makes a positive judgement on which she receives negative feedback, the memory trace is assumed to reflect only the actually experienced outcome. However, it is possible that the judgement is coded as "internal feedback" whether or not external feedback is received. The stored memory traces thus include all inferences and external information.

In the straightforward version of constructivist coding presented in this thesis, inferred outcomes are retrieved as indistinguishable from externally presented information (in fact, inference and objective experience are likely to be aspects of the same representation). Possibly, however, the representations of inferences and those of actual experiences, though confused at the time of retrieval, may be governed by different processes. Internally generated (components of) representations may be distinct from those with an external origin in terms of processes involved in storage, associations, retrieval, etc. Research has demonstrated differences in processing of presented and generated words, where self-generated items were better, and more reliably, recalled than externally presented information (Fiedler, Lachnit, Fay, & Krug, 1992; Raye, Johnson, & Taylor, 1980). Such findings may be relevant for the specification of a process such as constructivist coding.

In summary, questions regarding the cognitive processes associated with incomplete experiences have been surprisingly neglected in research on categorisation learning and decision making. Today, constructivist coding is the only effort to specify the learning processes that deal with selective feedback. Constructivist coding is still a tentative hypothesis that should benefit from further specification and investigation. Possibly, other models may be developed that provide alternative accounts for learning with selective feedback. However, constructivist coding is not only one potential answer, but also the only empirically supported answer so far suggested to the question of what is coded into memory in the absence of feedback.

Further Considerations

In the introduction of this thesis, the characteristics of the typical laboratory task were contrasted with everyday decision situations. Selective feedback was identified as an aspect of everyday decision making and experiential learning that has been neglected in laboratory research. The research presented in this thesis investigated the effects of learning from selective feedback. The isolation and manipulation of one factor at a time is the method of the experimenter who acknowledges the need for control to enable causal inferences. Other variables are left to be manipulated in future designs in order to investigate their effects on decisions and their possible interaction effects with selectivity of feedback. Results in Study III demonstrate such interaction effects of prior beliefs and selective feedback on the estimated proportions of suitable applicants.

There are several aspects of the experimental task used in the studies presented in this thesis that makes it quite different from an everyday decision situation. The benefits of this artificially produced task lie in the analysis and interpretation of results which may more readily be related to relevant research in similar areas. The disadvantages involve the risk that the results are less applicable to situations that are very different from the ones specified in the experiment.

Some characteristics of an everyday environment may be particularly relevant for further research involving selective feedback. For example, the experiments presented in this thesis all have correct answers, deterministically predicted by the presented features. This is rarely the case in any natural decision environment, where the outcome may be only partly predictable even in the presence of a large amount of relevant information. Possibly, in a more unpredictable decision environment that involves unreliable and ambiguous feedback, people have to rely even more on inferential processes and constructivist coding to achieve a coherent representation of the environment.

The significance of the environment in human cognition is difficult to exaggerate. To understand the human mind, we must study it in relation to the world in which it operates. Important research is performed with the aim of representative design and random sampling of cognitive tasks in a natural environment (e.g., Dhami, Hertwig & Hoffrage, 2004; Hogarth, 2005, 2006; Hurlburt, 1997). Further, the effects of computationally specified aspects of cognitive processes may be investigated in simulations under a broader range of environmental conditions (e.g., McKenzie, 1994). The systematic investigation of human cognition under varying circumstances that people naturally encounter not only benefits the understanding of human cognition. Practical implications involve any human activity (personal or professional) in which people use the knowledge they gain from experience; in education, in science, and in social, medical, economic or legal decision making, or any area in which people's knowledge and beliefs form the basis for important decisions.

Acknowledgements

I am very lucky to have had the opportunity to benefit from the exceptional talent of my supervisor Peter Juslin. Not only is he a highly intelligent researcher; he is also a most knowledgeable, inspiring, and patient supervisor. The admirable competence of the other (past and present) members of the research group has further enriched my learning experience and I greatly appreciate the insight, comments, and contributions from Tommy Enkvist, Göran Hansson, Patrik Hansson, Maria Henriksson, Linnea Karlsson, Marcus Lindskog, Håkan Nilsson, Anna-Carin Olsson, Henrik Olsson, and Anders Winman.

The thesis was further improved with the help of Claes von Hofsten, Hedvig Söderlund, and Leslie Shaps. Finally, the quality of my working days owes much to my eminent colleague Erik Rautalinko and the other talented people at the department.

References

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin, 93,* 203-231.

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

Allport, G. (1954). *The nature of prejudice.* Cambridge, MA: Addison-Wesley.

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review, 91,* 112-149.

Ashby, F., & Maddox, W. (2005). Human category learning. *Annual Review of Psychology, 56*, 149-178.

Ayton, P., & McClelland, A. G. (1997). How real is overconfidence? *Journal of Behavioral Decision Making, 10,* 279-285.

Balzer, W. K., Doherty, M. E., & O'Connor, R. (1989). Effects of cognitive feedback on performance. *Psychological Bulletin, 106,* 410-433.

Barlow, H. (1989). Unsupervised learning. *Neural Computation, 1,* 295-311.

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

Bartlett, F. C. (1932). *Remembering.* Cambridge, England: Cambridge University Press.

Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex, 10,* 295-307.

Brehmer, B. (1980). In one word: Not from experience. *Acta Psychologica, 45*, 223-241.

Brunswik, E. (1943). Organismic achievement and environmental probability. *Psychological Review, 50*, 255-272.

Camerer, C.F. & Johnson, E.J. (1997). The process-performance paradox in expert judgment: How can experts know so much and predict so badly? In W. M. Goldstein & R. M. Hogarth (Eds.) *Research on judgment and decision making: Currents, connections, and controversies.* (pp. 342-364). New York: Cambridge University Press.

Chapelle, O., Scölkopf, B., & Zien, A. (2006). *Semi-supervised learning.* London: MIT Press.

Cooksey, R. W. (1996). The methodology of social judgement theory. *Thinking & Reasoning, 2,* 141-173.

Cosmides, L., & Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition, 58,* 1-73.

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

Cutting, J. E. (2000). Accuracy, scope, and flexibility of models. *Journal of Mathematical Psychology, 44,* 3-19.

Dawes, R. M. (1994). House of Cards: Psychology and Psychotherapy Built on Myth**.** New York: The Free Press.

Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science, 243,* 1668-1674.

Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review, 112*, 951-978.

Denrell, J., & Le Mens, G. (2007). Interdependent sampling and social influence. *Psychological Review, 114,* 398-422.

Doherty, M. E. & Brehmer, B. (1997). Paramorphic models versus process models. In W. M. Goldstein & R. M. Hogarth (Eds.) *Research on judgment and decision making: Currents, connections, and controversies.* (pp. 537-551). New York: Cambridge University Press.

Donaldson, W. (1992). Measuring recognition memory. *Journal of Experimental Psychology: General, 121,* 275-277.

Edwards, W., Lindman, H., & Savage, L. J. (1963). Bayesian statistical inference for psychological research. *Psychological Review, 70,* 193-242.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review, 85*, 395-416.

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology, 32*, 53-88.

Einhorn, H. J., Kleinmuntz, D. N., & Kleinmuntz, B. (1979). Linear regression and process-tracing models of judgment. *Psychological Review, 86*, 465-485.

Eiser, J., Fazio, R. H., Stafford, T., & Prescott, T. J. (2003). Connectionist Simulation of Attitude Learning: Asymmetries in the Acquisition of Positive and Negative Evaluations. *Personality and Social Psychology Bulletin, 29,* 1221-1235.

Elwin, E., Juslin, P., Olsson, H., & Enkvist, T. (2007). Constructivist coding: Learning from selective feedback. *Psychological Science, 18,* 105-110.

Estes, W. (1976). The cognitive side of probability learning. *Psychological Review, 83, 37-64*.

Evans, J. S. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59,* 255-278.

Evans, J. St. B. T, Clibbens, J., Cattani, A., Harris, A., & Dennis, I. (2003). Explicit and implicit processes in multicue judgment. *Memory & Cognition, 31,* 608-618.

Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valance asymmetries. *Journal of Personality and Social Psychology, 87*, 293- 311.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*, 659-676.

Fiedler, K. (2008). The ultimate sampling dilemma in experience-based decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 186-203.

Fiedler, K., & Juslin, P. (2006). *Information sampling and adaptive cognition*. New York, NY: Cambridge University Press.

Fiedler, K., Lachnit, H., Fay, D., & Krug, C. (1992). Mobilization of cognitive resources and the generation effect. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 45A,* 149-171.

Fischer, I., & Budescu, D. V. (2005). When do those who know more also know more about how much they know? The development of confidence and performance in categorical decision tasks. *Organizational Behavior and Human Decision Processes, 98,* 39-53.

Fischhoff, B., & Beyth-Marom, R. (1983). Hypothesis evaluation from a Bayesian perspective. *Psychological Review, 90,* 239-260.

Fritzsche, B. A., & Brannick, M. T. (2002). The importance of representative design in judgment tasks: The case of resume screening. *Journal of Occupational and Organizational Psychology, 75,* 163-169.

Fugelsang, J. A., Stein, C. B., Green, A. E., & Dunbar, K. N. (2004). Theory and Data Interactions of the Scientific Mind: Evidence From the Molecular and the Cognitive Laboratory. *Canadian Journal of Experimental Psychology, 58,* 86-95.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The Role of Representative Design in an Ecological Approach to Cognition. *Psychological Bulletin, 130*, 959-988.

Garb, H. N. (1989). Clinical judgment, clinical training, and professional experience. *Psychological Bulletin, 105,* 387-396.

Garcia- Retamero, R., Rieskamp, J. (in press). Adaptive mechanisms for treating missing information: A simulation study. The Psychological Record.

Geng, H., Qi, Y., Li, Y., Fan, S., Wu, Y., & Zhu, Y. (2007). Neurophysiological correlates of memory illusion in both encoding and retrieval phases. *Brain Research, 1136,* 154-168.

Gigerenzer, G. (2007). Fast and Frugal Heuristics: The Tools of Bounded Rationality. In Koehler, D.K., Harvey, N. (Eds.) Oxford: Blackwell.

Gigerenzer, G., Czerlinksy, J., & Martingnon, L. (2002). How Good Are Fast and Frugal Heuristics? In Gilovich, T., Griffin, D., & Kahneman, D. (Eds.) Cambridge University Press.

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

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

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

Gluck, M. A., Shohamy, D., & Myers, C. (2002). How do people solve the "weather prediction" task?: Individual variability in strategies for probabilistic category learning. *Learning & Memory, 9,* 408-418.

Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review, 109*, 75-90.

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

Gonsalves, B., & Paller, K. A. (2000). Neural events that underlie remembering something that never happened. *Nature Neuroscience, 3,* 1316-1321.

Gonsalves, B., Reber, P. J., Gitelman, D. R., Parrish, T. B., Mesulam, M., & Paller, K. A. (2004). Neural evidence that vivid imagining can lead to false remembering. *Psychological Science, 15,* 655-660.

Gregory, R.L. (1980). Perceptions as Hypotheses. In Noë, A. & Thompson, E. (Eds; 2002). Vision and mind: Selected readings in the philosophy of perception. (pp 111-133). x, 627 pp. Cambridge, MA, US: MIT Press.

Griffiths, O. & Newell, B. R. (2009). *The impact of complete and selective feedback in static and dynamic multiple-cue judgment tasks.* Paper submitted to the meeting of the Cognitive Science Society in Amsterdam, the Netherlands.

Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: indirect payoff information and its implications for choice. *International Journal of Game Theory, 34,* 285-302.

Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment, 12,* 19-30.

Hall, C. C., Ariss, L., & Todorov, A. (2007). The illusion of knowledge: When more information reduces accuracy and increases confidence. *Organizational Behavior and Human Decision Processes, 103,* 277-290.

Hammond, K. R. (1955). Probabilistic functioning and the clinical method. *Psychological Review, 62,* 255-262.

Hammond, K. R. (1996). Upon reflection. *Thinking & Reasoning, 2,* 239-248.

Hammond, K. R., Hamm, R. M., Grassia, J. L., & Pearson, T. (1997). Direct comparison of the efficacy of intuitive and analytical cognition in expert judgment. In W. M. Goldstein & R. M. Hogarth (Eds.) *Research on judgment and decision making: Currents, connections, and controversies.* (pp. 144-180). New York: Cambridge University Press.

Hammond, K. R., & Summers, D. A. (1972). Cognitive control. *Psychological Review, 79*, 58-67.

Hammond, K.R., Summers, D.A., & Deane, D.H. (1973). Negative effects of outcome-feedback in multiple-cue probability learning. *Organizational Behavior and Human Performance, 9*, 30-34.

Harvey, N. (1997). Confidence in judgment. *Trends in Cognitive Sciences, 1,* 78-82.

Harvey, N. (2001). Studying judgement: General issues. *Thinking & Reasoning, 7,* 103-118.

Hasher, L., & Zacks, R. T. (1979). Automatic and effortful processes in memory. *Journal of Experimental Psychology: General, 108*, 356-388.

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

Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology, 52,* 653-683.

Hawkins, S. A., & Hastie, R. (1990). Hindsight: Biased judgments of past events after the outcomes are known. *Psychological Bulletin, 117*, 311- 327.

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

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2006).

Hogarth, R. M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin, 90*, 197-217.

Hogarth, R. M. (2005). The challenge of representative design in psychology and economics *Journal of Economic Methodology, 12,* 253-263.

Hogarth, R. M. (2006). Is confidence in decisions related to feedback? Evidence from random samples of real-world behavior. In K. Fiedler & P. Juslin (Eds.) *Information sampling and adaptive cognition*, (pp. 456-484). New York: Cambridge University Press.

Horton, N. J. & Lipsitz, S. R. (2001). Multiple imputation in practice: Comparison of software packages for regression models with missing data. *American Statistician, 55*, 244- 254.

Hurlburt, R. T. (1997). Randomly sampling thinking in the natural environment. *Journal of Consulting and Clinical Psychology, 65,* 941-949.

Jaccard, J., & Wood, G. (1988). The effects of incomplete information on the formation of attitudes toward behavioral alternatives. *Journal of Personality and Social Psychology, 54,* 580-591.

Jacoby, J., Mazursky, D., Troutman, T., & Kuss, A. (1984). When feedback is ignored: Disutility of outcome feedback. *Journal of Applied Psychology, 69,* 531-545.

Johnson, J. E., & Bruce, A. C. (2001). Calibration of subjective probability judgments in a naturalistic setting. *Organizational Behavior and Human Decision Processes, 85,* 265-290.

Johnson, M. K. (2006). Memory and reality. *American Psychologist, 61,* 760-771

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3-28.

Johnson, M. K. & Raye, C. L. (1981). Reality monitoring. *Psychological Review, 88*, 67-85.

Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 924-941.

Juslin, P., Olsson, H., & Olsson, A-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General, 132*, 133-156.

Juslin, P., Winman, A., & Hansson, P. (2007). The naive intuitive statistician: A naive sampling model of intuitive confidence intervals. *Psychological Review, 114*, 678-703.

Juslin, P., Winman, A., & Olsson, H. (2000). Naive 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, US: Cambridge University Press.

Kant, I. (1952). The Critique of Pure Reason. (Translation by J.M.D. Meiklejohn). Chicago: Encyclopædia Britannica.

Kao, S.-F., & Wasserman, E. A. (1993). Assessment of an information integration account of contingency judgment with examination of subjective cell importance and method of information presentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1363-1386.

Karelaia, N., & Hogarth, R. M. (2008). Determinants of linear judgment: A meta-analysis of lens model studies. *Psychological Bulletin, 134*, 404-426.

Kim, H., & Cabeza, R. (2007). Differential contributions of prefrontal, medial temporal, and sensory-perceptual regions to true and false memory formation. *Cerebral Cortex, 17,* 2143-2150.

Klayman, J. (1988). On the how and why (not) of learning from outcomes. *Brehmer, Berndt (Ed); Joyce, C, R. B (Ed). (1988). Human judgment: The SJT view. (pp. 115-162). viii*, England: North-Holland.

Klayman, J. (1995). Varieties of confirmation bias. In J. Busemeyer, R. Hastie, & D. L. Medin (Eds.), *Decision making from a cognitive perspective*. New York: Academic Press (*Psychology of Learning and Motivation*, vol. 32), pp. 365-418.

Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review, 94,* 211-228.

Klayman, J., & Ha, Y.-W. (1989). Hypothesis testing in rule discovery: Strategy, structure, and content. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 596-604.

Klayman, J., Soll, J. B., Gonzalez-Vallejo, C., & Barlas, S. (1999). Overconfidence: It depends on how, what, and whom you ask. *Organizational Behavior and Human Decision Processes, 79,* 216-247.

Kluger, A. N., & DeNisi, A. (1996). Effects of feedback intervention on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin, 119,* 254-284.

Kluger, A. N., & DeNisi, A. (1998). Feedback interventions: Toward the understanding of a double-edged sword. *Current Directions in Psychological Science, 7,* 67-72.

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

Kruschke, J. K. (2005). Category learning. In K. Lamberts and R. L. Goldstone (Eds.), *The Handbook of Cognition*, Ch. 7, pp. 183-201. London: Sage.

Lagnado, D. A., Newell, B. R., Kahan, S., & Shanks, D. R. (2006). Insight and strategy in multiple-cue learning. *Journal of Experimental Psychology: General, 135,* 162-183.

Lipe, M.G. (1990) A lens-model analysis of covariation research. *Journal of Behavioral Decision Making, 3,* 47–59.

Loftus, E.F. & Hoffman, H. G. (1989) Misinformation and memory: The creation of new memories. *Journal of Experimental Psychology: General, 117,* 100-104.

Lovibond, P. F. (2001). The "near miss" as a fourth pathway to anxiety. *Behavioural and Cognitive Psychotherapy, 29,* 35-43.

Lovibond, P. F., Saunders, J., Weidemann, G., & Mitchell, C. J. (2008). Evidence for expectancy as a mediator of avoidance and anxiety in a laboratory model of human avoidance learning. *The Quarterly Journal of Experimental Psychology, 61,* 1199-1216.

Macmillan, N. A, & Creelman, C. D. (2005). *Detection theory. A user's guide.* 2nd ed., New Jersey: Lawrence Erlbaum Associates.

Maddox, W., Love, B. C., Glass, B. D., & Filoteo, J. (2008). When more is less: Feedback effects in perceptual category learning. *Cognition, 108,* 578-589.

March, J. G. (1996). Learning to be risk averse. *Psychological Review, 103,* 309-319.

Markman, A. B. (1989). LMS rules and the inverse base-rate effect: Comment on Gluck and Bower (1988). *Journal of Experimental Psychology: General, 118,* 417-421.

McKenzie, C. R., & Mikkelsen, L. A. (2007). A Bayesian view of covariation assessment. *Cognitive Psychology, 54,* 33-61.

Minda, J. P., & 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.

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review, 115,* 502-517.

Nessler, D., Mecklinger, A., & Penney, T. B. (2001). Event related brain potentials and illusory memories: The effects of differential encoding. *Cognitive Brain Research, 10,* 283-301.

Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). Straight choices: The psychology of decision making. *(2007), Straight choices: The psychology of decision making. xii*, NY, US: Psychology Press.

Newman, J. P., Wolff, W. T., & Hearst, E. (1980). The feature-positive effect in adult human subjects. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 630-650.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2,* 175-220.

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 and Review, 7,* 375-402.

Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18*, 211-233.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101,* 608-631.

Okado, Y., & Stark, C. E. (2005). Neural activity during encoding predicts false memories created by misinformation. *Learning & Memory, 12,* 3-11.

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

Peterson, C. R., Hammond, K. R., & Summers, D. A. (1965). Optimal responding in multiple-cue probability learning. *Journal of Experimental Psychology, 70,* 270-276.

Pitt, M.A., & Myung, I.J. (2002). When a good fit can be bad. *Trends in Cognitive Science, 6*, 421-425.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472-491.

Raye, C. L., Johnson, M. K., & Taylor, T. H. (1980). Is there something special about memory for internally generated information? *Memory & Cognition, 8,* 141-148.

Remus, W., O'Connor, M., & Griggs, K. (1996). Does feedback improve the accuracy of recurrent judgmental forecasts? *Organizational Behavior and Human Decision Processes, 66,* 22-30.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General, 135,* 207-236.

Roediger, H. L. & McDermott, K. B. (1995). Creating false memories – remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory and Cognition, 21*, 803-814.

Roediger, H. L., III, Watson, J. M., McDermott, K. B., & Gallo, D. A. (2001). Factors that determine false recall: A multiple regression analysis. *Psychonomic Bulletin & Review, 8,* 385- 407.

Ross, L., Lepper, M. R., & Hubbard, M. (1975). Perseverance in self-perception and social perception: Biased attributional processes in the debriefing paradigm. *Journal of Personality and Social Psychology, 32,* 880-892.

Rubin, D. B. (1976) Inference and missing data. *Biometrika, 63,* 581- 592.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926-1928.

Salkovskis, P. M., Clark, D. M., & Gelder, M. G. (1996). Cognition-behaviour links in the persistence of panic. *Behaviour Research and Therapy, 34,* 453-458.

Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General, 110,* 474- 494.

Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology, 49,* 289-318.

Schroth, M. L. (1997). Effects of frequency of feedback on transfer in concept identification. *American Journal of Psychology, 110,* 71-79.

Schwarz, N. (2002). Feelings as information: Moods influence judgments and processing strategies. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.) *Heuristics and biases: The psychology of intuitive judgment.* (pp. 534-547). New York, US: Cambridge University Press.

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

Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review, 63*, 129-138.

Sinharay, S., Stern, H. S., & Russell, D. (2001). The use of multiple imputation for the analysis of missing data. *Psychological methods. 6*, 317- 329.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3-22.
Slovic, P., Fischhoff, B., & Lichtenstein, S. (1977). Behavioral decision theory. *Annual Review of Psychology, 28,* 1-39.
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, E.E., Patalano, A.L., & Jonides, J. Alternative mechanisms of categorization. *Cognition*, *65,* 167-196.
Sutton, R. S., & Barto, A.G. (1998). *Reinforcement Learning: An Introduction.* London: MIT Press.
Swets, J.A. (1964). *Signal detection and recognition by human observers*. New York: Wiley.
Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science, 16*, 167-171.
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.
Unkelbach, C., Fiedler, K., & Freytag, P. (2007). Information repetition in evaluative judgments: Easy to monitor, hard to control. *Organizational Behavior and Human Decision Processes, 103*, 37-52.
Vallabha, G. K., McClelland, J. L., Pons, F., Werker, J. F., & Amano, S. (2007). Unsupervised learning of vowel categories from infant-directed speech. *PNAS Proceedings of the National Academy of Sciences of the United States of America, 104,* 13273-13278.
von Helversen, B., & Rieskamp, J. (2008). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General, 137,* 73-96.
Walther, E., Fiedler, K., & Nickel, S. (2003). The more we know the better? Influences of prior knowledge on constructive biases in social judgment. *Swiss Journal of Psychology, 4*, 219-233.
Weber, E. U., & Hsee, C. K. (2000). Culture and individual judgment and decision making. *Applied Psychology: An International Review: 49,* 32-61.
Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology. 60,* 53-85.
White, C. M., & Koehler, D. J. (2004). Missing information in multiple-cue probability learning. *Memory & Cognition, 32,* 1007-1018.
Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making, 19*, 1-16.
Zacks, R. T., & Hasher, L. (2002). *Frequency processing: A twenty-five year perspective*. New York, NY: Oxford University Press.
Zhu, X. (2005). *Semi-supervised learning with graphs.* Doctoral Dissertation, Carnegie Mellon University.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences 51

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-106880

ACTA
UNIVERSITATIS
UPSALIENSIS
UPPSALA
2009