THE ROLE OF INFORMATION SYSTEMS IN DECISION-MAKING BIASES

Mini dissertation by

SIBELLA MARGARETHA TURPIN
(87062471)

submitted in partial fulfilment of the requirements for the degree

MAGISTER PHILOSPHIAE IN INFORMATICS

in the

DEPARTMENT OF INFORMATICS
SCHOOL OF INFORMATION TECHNOLOGY

of the

FACULTY OF ENGINEERING, BUILT ENVIRONMENT AND INFORMATION TECHNOLOGY

Supervisor: Prof NF du Plooy **November 2003**

ACKNOWLEDGEMENTS

I hereby wish to thank everyone who has assisted me with this study, directly and indirectly. I especially want to acknowledge the following parties for their support:

- The CSIR, and in particular Hans Ittmann;
- My colleagues Johan, Mario, Isabel and Renée;
- My supervisor, prof Niek du Plooy;
- My family: Andrew, Marguerite and Laura; and
- The numerous family members and friends who have helped to look after my children during the course of my studies.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

SUMMARY

Information systems and in particular decision support systems have been developed to supplement human information processing and to assist with decision-making. Human decision-making is facilitated by the often unconscious use of heuristics or rules of thumb in situations where it may not be possible or feasible to search for the best decision. Judgemental heuristics have previously been found to lead to biases in decision-making. When information systems are used as decision aids, they may have an influence on biases.

This study investigates the possible role of information systems in introducing, reinforcing or reducing biases of decision-making.

It has been found that information systems have the ability to introduce new biases and to reinforce biases. Information systems can also reduce biases, but this requires innovate thinking on the way information is represented and the way human decision-making processes are supported. It has also been found that in the real world, other than the laboratories where biases are usually measured, other constraints on rational decision-making, such as politics or data errors, can overshadow the effects of biases.

1. INTRODUCTION AND OVERVIEW

1.1. Introduction

Information systems and in particular decision support systems have been developed to supplement human information processing (George *et al*, 2000). The ability to perform complex calculations, to store large amounts of data and to perform operations on the data, has extended people's capability to process and analyse information. Computers were initially used for data processing, where automation resulted in vast improvements in efficiency (Introna, 1997). They also entered into the domain of providing management information by performing queries on increasingly large databases. Systems for supporting managers, such as management information systems, decision support systems and executive information systems all claim to assist with or improve managerial decision-making.

What are the decision-making processes that need to be supported by information systems? Different models of decision-making exist; of these, the rational model is believed to be the norm or ideal (Keen and Scott Morton, 1978). In classic rationality, the subjective expected utility is calculated and it is believed to provide an optimal decision. Behavioural decision researchers observed that in the real world, people's decision-making processes do not follow the rational model. A widely recognised alternative for the rational model is Simon's model of bounded rationality (Simon, 1979). Simon claimed that if people are not able to choose an optimal outcome under conditions of uncertainty, complexity and time pressure, they 'satisfice'. People learn short cuts or rules of thumb that guide them to a good enough rather than optimal decision. These short cuts or heuristics are extremely useful, but they may also lead people to make biased decisions. With the introduction of Bayesian statistics in psychological research in the 1960s (Kahneman *et al*, 1982), it became possible for Tversky and Kahneman (1974) to measure the biases contained in heuristics. Subsequent to Tversky and Kahneman's work, a number of other psychologists have researched the behavioural phenomenon of heuristics as well as its consequences. Hogarth (1980) describes a large number of heuristics and associated biases that are found in the way people deal with information, namely regarding information acquisition, processing, output and feedback. Information systems perform all the latter functions, and also in general encode rules of behaviour in organisations. Thus, it can be expected to find biases in the 'rules' that information systems comprise of and prompt people to use. On the other hand,

computers are built to be more logical and consistent than human decision-makers; they are the artefacts of the mechanistic world view (Dahlbom and Mathiassen, 1993). Perhaps they can assist in reducing decision-making biases?

A number of studies have investigated the effect of information systems on the occurrence of biases in decision-making, for example Rai *et al* (1994), Skitka *et al* (1999), George *et al* (2000), Shore (1996) and Lim, Benbasat and Ward (2000). However, no general framework has been found for investigating the role of information systems in propagating or reducing biases. This study will attempt to create a better understanding of this role.

What follows in the rest of the chapter is some background on heuristics and biases in decision-making as well as the role they play when decision-making is supported by information systems. The problem statement for the research project is given, as well as an outline of the research undertaken. Finally, an outline for the remainder of the document will be provided.

1.2. Decision-making: a historical perspective

Herbert Simon, in the lecture he delivered when receiving the 1978 Nobel Prize in Economic Science, gave an overview of the dominant views on decision-making in the twentieth century up to then (Simon, 1979). According to Simon, the decision-making theory of perfect rationality was inherited from the classical and neo-classical economists. Whether aiming towards full employment of resources, efficient allocation of resources, or maximisation of profit, these economists strived towards the optimisation of mathematical models that described the economy at a macro-level. During World War II, mathematical modelling to support decision-making in the military environment became popular. These techniques were also applied in the business environment, and termed ‘operations research’ or ‘operational analysis’ by mathematicians and statisticians, and ‘management science’ by economists. “Optimisation techniques were transported into management science from economics” (1979, p 498). According to Simon, mathematical modellers had to make recommendations while operating under the limitations of the available data and processing capability of computers. Already at that stage, operations researchers were forced to ‘satisfice’ by either simplifying the model that would be optimised, or to sacrifice optimality and attempt to find good enough solutions by simplifying the searches through the problem space. Thus, they could either attempt to find optimal models of a simplified world or satisfactory solutions for a more realistic world.

Simon formulated his theory of bounded rationality while studying behaviour and decision-making of firms from the 1930s onwards. He observed that behaviour in organisations did not follow the classical models of rational choice. One version of the classical model is discussed in Simon (1977). The four-step decision model consists of intelligence (finding occasions for making a decision), design (inventing, developing and analysing possible courses of action), choice (choosing one of the possible courses of action) and review (assessing past choices). When applying this model, neo-classical economists and operations researchers strive to maximise the Subjective Expected Utility (SEU). It means that the alternative options are *quantified* and the *best possible one* (with the highest utility) chosen. When using the four-step model and simultaneously striving to maximise the SEU, it is assumed that managers:

- "have knowledge of all possible alternatives;
- have complete knowledge about the consequences that follow each alternative;
- have a well-organised and stable set of preferences for these consequences, and
- have the computational ability to compare consequences and to determine which one is preferred" (Kreitner and Kinicki, 2001).

In reality, decision-makers face uncertainties and have incomplete knowledge (Simon, 1979). They also have limited computational powers. Simon defines the concept of bounded rationality by the activities of *searching* and *satisficing*. A person has certain aspirations of what they seek. Since all alternatives are not available upfront, a search is undertaken for alternatives that would satisfy the aspirations. Alternatives are found and evaluated in sequence, until a satisfactory alternative is found. The process of sequential searching and evaluating against aspirations until satisfied is called 'satisficing'.

Bounded rationality is regarded as a major improvement on the classic rationality of optimisation, and Simon has devoted much research effort in order to establish the theory of bounded rationality and gain acceptance of it. Yet, behaviour under bounded rationality is also rational (or reasonable or logical). It approximates perfect rationality under the constraints or limitations of reality.

Not all decision researchers agree that decision-makers behave rationally or even boundedly rational. Evidence of this can be found in Keen and Scott Morton's (1978) and Huber's (1981)

categorisations of decision-making models. Apart from the views of rational and boundedly rational behaviour (which Huber believes to be so similar compared to the other models that he groups them together), are views such as the political model of decision-making (Pfeffer, 1981), the garbage can model (Cohen *et al*, 1972) and the avoidance model (March, 1988). A recent addition is the Recognition-Primed Decision Model of Gary Klein (1999), an example of naturalistic decision-making. These and other models will be discussed in more detail in Chapter 3.

1.3. Heuristics and biases in decision-making

One of the streams of research that followed Simon's conceptualisation of bounded rationality is the study of heuristics and biases in human judgement. Daniel Kahneman and Amos Tversky are regarded as significant contributors to this field (Simon, 1979). It was mentioned in the previous section that operations researchers faced with computationally complex problems 'satisfice' and attempt to find good enough solutions by simplifying the searches through a problem space. These simplified searches are called *heuristics*. Decision researchers have adopted the term to describe boundedly rational behaviour by humans. Heuristics in this context are rules of thumb or informal reasoning strategies to cope with the complexity inherent in most decisions (Klein, 1999; Hammond *et al*, 1998). Heuristics are useful tools: they are relatively simple and intuitive (Harvey, 1998). However, Tversky and Kahneman have found that the use of heuristics can lead to severe and systematic errors in judgement (Tversky and Kahneman, 1974).

Tversky and Kahneman have built on the work of Simon that showed how people's decision-making did not follow the process of maximising subjective expected utility (SEU). When Bayesian statistics became used in psychological research in the 1960s, Tversky and Kahneman had a means to quantitatively measure the deviation of people's choices from the optimum. People in laboratory settings were given questions on probabilities. These people had different levels of training in probability theory. Tversky and Kahneman showed that under one set of circumstances, decision makers give far too little weight to prior knowledge and based their choices almost entirely on new evidence, while in other circumstances new evidence has little influence on opinions already formed (Simon, 1979). Also, people weight gains and losses differently and in most circumstances tend to be risk averse (Dietz and Stern, 1995). Based on their findings, Tversky and Kahneman categorised people's judgement-related behaviour into three heuristics, termed representativeness, availability, and adjustment and anchoring

(Kahneman *et al*, 1982). Associated with each heuristic are a number of typical decision-making biases or judgement 'errors'. A number of other biases such as decision framing (Tversky and Kahneman, 1981) were added to the initial lists of biases.

Kahneman, Tversky and their collaborators' work were supplemented by a number of publications with similar findings regarding decision-making biases. Three volumes of publications edited by Kahneman and others have appeared to date (Kahneman *et al*, 1982; Kahneman and Tversky, 2000; and Gilovich *et al*, 2002). Following the reasoning of the cognitive school of psychology that regards humans as information processors, Hogarth (1980) drew up a comprehensive list of biases associated with information processing. The findings of Tversky, Kahneman and their collaborators gave decision researchers the opportunity to become management consultants, providing advice to managers and organisations to improve their decision-making by eliminating biases. Examples of these are Russo and Schoemaker (2002) and Hammond *et al* (1999).

Criticism of the heuristics and biases literature

The literature on heuristics and biases leaves the impression that humans are not competent decision-makers. "It appears that people lack the correct programs for many important judgemental tasks… it may be argued that we have not had the opportunity to evolve an intellect capable of dealing conceptually with uncertainty" (Slovic, Fischhof and Lichtenstein, quoted in Chase *et al*, 1998). Yet, a number of authors do not agree that perfect rationality or information processing is the norm that people should be measured against, nor do some agree with the manner in which people's rationality has been tested in the heuristics and biases literature.

Whereas Tversky and Kahneman tested people and found their judgement to be flawed, Gary Klein (1998) observed and analysed the decision-making of people in their natural settings and found them to be competent decision-makers. What is the difference between the two situations? It appears as if Klein did not interfere with the decision-making process, but tried to understand it. On the other hand, it is implied by Lopes (Klein, 1998) and Chase *et al* (1998) that Tversky and Kahneman have selected and framed their questions in such a manner that their subjects have been set up for failure. Dietz and Stern (1995) argue that people's analytical decision-making abilities are not very advanced when it comes to arithmetic and algebra, but that people are very sophisticated when it comes to pattern recognition and classification. Tversky and Kahneman test

the arithmetic kind of skills, whereas Klein's Recognition Primed Decision model relies strongly on pattern recognition and classification.

Dietz and Stern (1995) write about the social context of decision-making. They discuss a number of socially related decision-making behaviours that are valid in a social context but can lead to outcomes different from the maximum SEU. Many of the biases listed by eg. Kahneman *et al* (1982) and Hogarth (1980) show some irrationality that could be socially motivated. Dietz and Stern hold that human decision-making is possibly far more sophisticated and evolved than the rational mode of decision-making, and not less sophisticated as some researchers would have it.

Boland (2001) criticises some of the assumptions on which the literature on heuristics and biases are based. One of them is the view of humans as information processors. Drawing on the work of Bruner, he calls for an alternative mode of cognition, namely the narrative mode. In the information processing mode, people test hypotheses, deduce consequences and look for if-then rules. In the narrative mode, a story is constructed within which a series of events become plausible in the context of people's experience and culture. It is a form of sense-making that allows a richer context for decision-making. Klein (1998) also emphasises the role of story-telling in decision-making and particularly in sharing experience. He believes the context of the story is significant; it is something that is lacking in decision-making rules.

Despite the criticism on the heuristics and biases literature, rational decision-making is still held to be a norm by many. Under this assumption, biases will be found, whether these are referred to as psychological traps (Hammond *et al*, 1999), flaws in judgement (Tversky and Kahneman, 1974) or general mistakes in decision-making (Russo and Schoemaker, 2002). The position taken for this study is that biases do exist, although their occurrence is perhaps over-emphasised in some of the literature, or not viewed in context.

1.4. Heuristics and biases in Information Systems

Information systems are designed to assist people with information processing, information analysis and even with decision-making. Within information systems, information is typically analysed according to the rational model of decision-making (Stair and Reynolds, 2001). In this context, it is possible that information systems inherit or reinforce biases of information processing. Also, organisational decision-making is often a matter of implementing rules rather than exercising choice (March, 1988). Since these rules are shaped by experience, some of them could be viewed as heuristics. Information systems abound with such rules, specifying for example what data is displayed in management reports and how performance variables are calculated and compared. Many of these rules would have been specified by the owners of the information systems, but others could have been implemented on behalf of the organisation by business consultants or even systems analysts or programmers, whose assumptions are embedded into systems with the best of intentions. Some evidence has been found of biases studied in an information systems context, namely:

- A DSS for house appraisals (George *et al*, 2000): It was found that the anchoring and adjustment bias (based on initial values) persisted even after attempts to counter it.
- A flight management system (Skitka *et al*, 1999): When using autopilot software, people allowed the system's recommendations to override their own judgements and performed worse than people in a non-automated setting.
- A battle management and radar tracking system (Fisher and Kingma, 2001): The disaster of the USS Vincennes shooting down an Iranian passenger airliner in the Gulf War could be partially ascribed to the Aegis system introducing an expectancy bias and not assisting users to correct the impression formed.
- Executive information systems (Rai *et al*, 1994): For example, salient graphics and exception reporting increase biases associated with the heuristics of availability, regression and overconfidence, as discussed by Tversky & Kahneman and Hogarth.
- The development and use of an expert system (Shore, 1996). It is shown how the biases of the different parties involved with the expert system, namely the subject-matter experts, knowledge engineers, end users as well as people validating and maintaining the system, contribute to errors occurring when the system is used.

- A performance appraisal system designed by Lim, Benbasat and Ward (2000). Multimedia is used to reduce first impression bias when doing a performance appraisal.

Despite the above examples it appears as if the topic of heuristics and biases has not yet been thoroughly researched from an information systems perspective. A starting point for assessing the possible impact of biases in information systems design and use is Hogarth's (1980) discussion of the characteristics of information processing that makes it susceptible to biases. Do or can information systems also possess these characteristics?

1.5. Problem statement

The rational model of decision making (eg. Simon, 1977) pervades the decision-making literature as well as the field of decision support systems. Yet, in situations of uncertainty, complexity and time pressure, the conditions for the use of the rational model cannot be met. In reality, people 'satisfice' or make do with what information and resources they have (Kreitner and Kinicki, 2001). They also use heuristics or rules of thumb to guide them. These heuristics were found to contain biases, such as those measured and reported by Tversky and Kahneman (1974). The literature on heuristics and biases has been criticised, amongst others for assuming that people should behave clinically rational and act like information processors. Even if people do not behave like this themselves, they often de-contextualise or rationalise their own behaviour in order to delegate some of their functions to machines. It is possible that information systems can have an effect on biases of information processing and rational decision-making, depending on how they are designed and used. Some case studies were found to prove this assertion, for example Rai *et al* (1994), Skitka *et al* (1999), George *et al* (2000), Shore (1996) and Lim *et al* (2000). However, there does not appear to be an overall or recently updated framework that could assist IS specialists to understand and deal with biases occurring as a result of information systems design and use. Following this, the main research question for the proposed research project is as follows:

- How does information system – supported decision-making get affected by biases?

1.6. Research to be undertaken

A literature study will be done to provide background on the use of heuristics in decision-making, the biases associated with the use of heuristics as well as the occurrence of biases that result from the design and use of information systems. A framework will be drafted to investigate the role of information systems in decision-making biases. Case studies of the role of information systems in reducing or reinforcing biases, as found in the literature, will be discussed. Conclusions from the case studies will be compared to sections of previously discussed theory, as well as to the framework of the role of information systems in decision-making biases.

The empirical component of the research will consist of an investigation into a decision support project undertaken by a local organisation, where aspects of the information analysis and reporting provided the opportunity for biases to occur. As with the case studies from literature, the theory of heuristics and biases will be used to investigate the occurrence of biases on this project.

1.7. Summary of the remainder of the document

The remainder of the document is structured as follows:

- Chapter 2 describes the research methodology.
- Chapter 3 provides background information on decision-making, which forms the context for the use of heuristics. It also introduces the concept of heuristics.
- In chapter 4, the literature on the alleged biases associated with the use of heuristics is discussed. It includes a criticism of the heuristics and biases literature.
- Chapter 5 investigates the occurrence of decision-making biases in the manner that information systems are designed and used.
- Chapter 6 reports on the empirical work undertaken. Results of the empirical study are interpreted and compared to aspects of the literature study that was undertaken in chapters 3 to 5.
- Chapter 7 summarises and concludes the study.

2. RESEARCH METHODOLOGY AND RESEARCH QUESTIONS

2.1. Introduction

This chapter focuses on the process aspects of the research. Du Plooy, Introna and Roode's (1994) process framework is used as a basis for deriving research questions.

The majority of the study consists of a literature survey. A theoretical understanding of the role of heuristics and biases in decision-making is established, after which the possible role of information systems on decision-making biases is investigated. The literature study concludes with an analysis of a number of case studies where biases are investigated in an information systems context.

A small empirical study will be undertaken. The empirical study will investigate the occurrence of a particular type of bias in a decision support project where information systems are used.

2.2. The context of the information systems discipline

Banville and Landry (1989) classify scientific disciplines according to three variables, namely functional dependence, strategic dependence and task uncertainty. Functional dependence will be high when researchers are expected to use the specific ideas and procedures of fellow specialists in their research. Strategic dependence will be high if researchers have to convince colleagues of the strategic importance of their problem and approach in order to be accepted. Strategic task uncertainty is low when there is an agreed hierarchy of research problems in the field of research, and high when the schools of thought are only loosely coupled. The field of information systems are classified as having a low degree of functional dependence, a low degree of strategic dependence and a high degree of strategic task uncertainty. Banville and Landry refer to this classification as a 'fragmented adhocracy', and describe it as follows:

- "Research in information systems is rather personal and weakly co-ordinated in the field as a whole;
- A researcher can gain a reputation by contributing in a way that is largely specific to a group of colleagues or a research site;

- A field is largely open to an educated public and amateurs can affect the field's standards;
- Barriers to entry in the field are weak and going from one fragment to another is quite easy;
- Reputations are fairly fluid, control of resources is unstable, coalitions are likely to be ephemeral and leadership is often charismatic in nature; and
- Common-sense languages dominate the communication system" (Banville and Landry, p 80).

It can be inferred from the above that information systems researchers are allowed a fair amount of freedom in terms of research topics and research approaches.

The pluralistic nature and porous boundaries of the information systems discipline can be seen in the wide number of reference disciplines it draws on, namely behavioural science, computer science, decision theory, information theory, organisation theory, management theory, language theories, systems theory, research, social science, management science, artificial intelligence, economic theory, ergonomics, political science and psychology (Barki *et al*, in Du Plooy, 1998).

The disciplines that will be referenced in this study are behavioural science, decision theory and psychology.

2.3. Sociological paradigms

Burrel and Morgan (1979) describe four different paradigms according to which the world, and organisations, can be viewed. The paradigms were obtained by combining the extremes of two variables.

The first variable has to do with how we view the world (the nature of social science), and the extremes are *subjectivity* and *objectivity*. In the *objective* world view, the world is regarded as an object, remote from the observer. There is a single, correct third party's view on things, people and situations. According to the *subjective* world view, the world, and reality itself, are social constructs.

The second variable is about the way the world behaves, or the nature of society. The associated extremes are order and conflict, or, as preferred by Burrel & Morgan, *regulation* and *radical change*. Under the view of order or *regulation*, the world is seen as an inherently orderly system. Society operates best in this manner, and should thus be regulated at all costs. Change can only be allowed as far as it adheres to existing rules, processes and agendas. The conflict or *radical change* view assumes that conflict is inherent in society, because of existing injustices and oppression.

The different combinations of the extremes of the two variables result in four paradigms, namely the functionalist paradigm, the interpretive paradigm, the radical structuralist paradigm and the radical humanist paradigm (figure 2.1):

- The functionalist paradigm assumes an orderly world with an absolute frame of reference, as described by the language of the natural sciences.
- The interpretive paradigm assumes an orderly world that is defined and understood from the frame of reference of the observer. It is a passive view, aimed at understanding and reflecting rather than improving.
- The radical structuralist view promotes radical change with the assumption of objectivity. There is thus a single 'right' way towards which the world should be changed. Ideologies aimed at freeing the oppressed (at group or societal level) can be classified here.
- Radical humanism promotes radical change but allows for multiple or subjective views of the world. It strives towards true emancipation at the level of the individual.

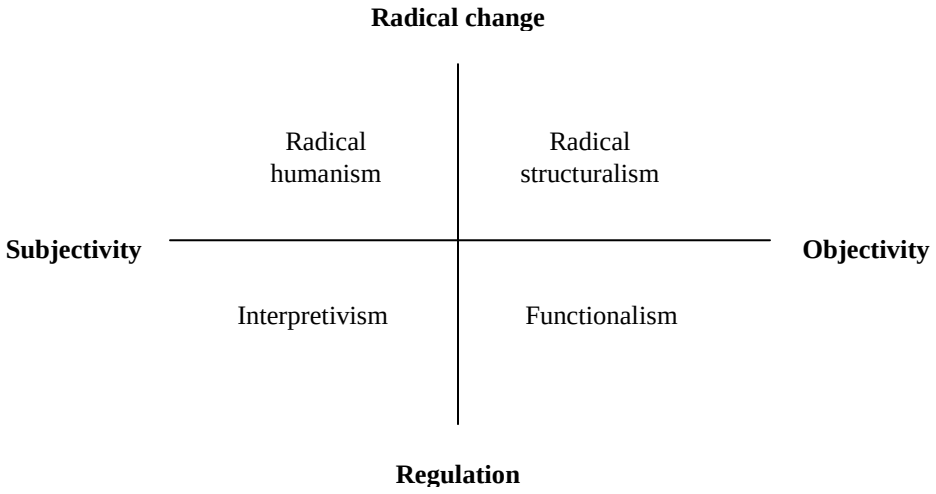

Figure 2.1: The four sociological paradigms

(Burrel and Morgan, 1979, p 22)

Positivist research methodologies can be classified under the functionalist paradigm. Likewise, studies with an interpretive approach can be regarded to fall into the interpretive paradigm. Critical social theory with its associated emancipatory goals address the order-conflict extremes, and such studies will touch upon the radical humanist paradigm. However, it is not necessary to limit oneself to one paradigm. Du Plooy (1998), based on arguments by Boland, calls for a multi-paradigmatic approach. It is argued that "objective fact is socially constructed, while subjective opinion is grounded in observable fact" (Du Plooy, *op cit*, p 52). Rather than having research labelled as 'objectivist' or 'subjectivist', it can be attempted to include views from more than one paradigm.

2.4. The Process Framework

Du Plooy, Introna and Roode (1994) developed a process framework for developing research questions. The framework encourages deliberate movement through the four paradigms of Burrel and Morgan. Explicit recognition of the different paradigms and accompanying assumptions, and thus a deliberate viewing of the problem from different perspectives, leads to an appreciation of the unique problem situation and a more holistic approach. The mentioned traversing of the problem space is also meant to lead to greater insight into the research problem.

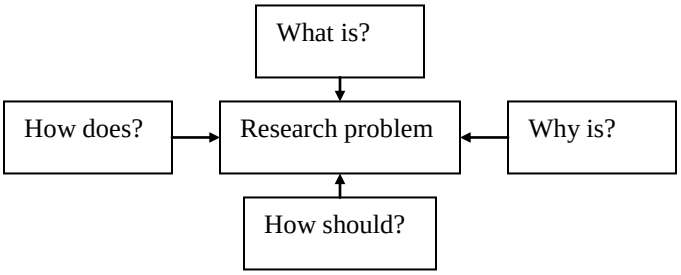

Figure 2.2: Traversing the research question space

(Du Plooy, 1998)

Although there is no one-to-one correlation between the different kinds of questions and the four paradigms, it can be noted that the 'Why is?' question relies heavily on the functionalist paradigm (Roode, 1994, p 13), and the 'What is?' and 'How does?' questions on the interpretive paradigm. The 'How should?' question advocates change and therefore a movement towards the radical structuralist or radical humanist paradigms. On the other hand, a 'How should?' question could imply a prescriptive view from the functionalist paradigm. Questions can be enriched with views from different perspectives, for example a 'How does?' question can be approached from the functionalist as well as interpretive paradigms.

2.5. Research questions

Research questions for this study are of the 'What is?' and 'How does?' types. 'What is?' questions explore the fundamental nature of the topic under consideration. 'How does?' questions are concerned with observing the phenomenon and describing the way it manifests itself in reality (Roode, 1994, p 12). Conclusions from the study do contain a few 'How should' remarks, although the study is not prescriptive in nature.

The main research question for the proposed research project is as follows:

- How does information system – supported decision-making get affected by biases?

Derived research questions

Heuristics and the associated biases occur in a context of decision-making. In particular, certain models of decision-making are assumed by the researchers that have documented the occurrence of heuristics and biases. Therefore, it is attempted to gain an understanding of the decision-making environment with the following questions:

- How do people make decisions?
- What is a heuristic?
- How are heuristics used in the process of decision-making?
- What is the effect of heuristics on the outcomes of decisions or judgements?

The term 'bias' enters the scene with the last mentioned question. However, the word 'bias' has not been used in any of the above questions. This is because the literature contains mixed opinions on the significance and role of biases in decision-making; they are not viewed by all researchers as a necessary consequence of the use of heuristics.

Despite the mentioned controversy, evidence of biases has been documented. The effect of a number of these biases has been studied in the context of information systems. The background information gained from addressing the four derived questions can equip one to study biases in an information systems setting, and thus assist to address the main research question.

Views from the different paradigms

Much of the literature on heuristics and biases, as well as the underlying assumptions about decision-making, is perceived to be extremely positivist in nature. The 'How does?' as well as the 'What is?' questions listed above will be viewed from the functionalist paradigm in order to adequately capture the thinking of this literature. On the other hand, the positivist assumptions made in this literature needs to be questioned. The same questions will also be viewed from the interpretive paradigm, and contrasting views on the topics will be discussed. Since the emphasis of the study is on understanding rather than on addressing conflict and advocating change, the radical structuralist and radical humanist paradigms will not receive similar attention.

2.6. Research planning

The research questions are addressed as follows in the remainder of the report:

Literature survey

Chapter 3 consists of a literature survey addressing the following research questions, placing the notion of heuristics in the context of decision-making:

- How do people make decisions?
- What is a heuristic?
- How are heuristics used in the process of decision-making?

Chapter 4 consists of a literature study addressing the following research questions, focusing on the biases associated with heuristics:

- How are heuristics used in the process of decision-making? (continued from Chapter 3)
- What is the effect of heuristics on the outcomes of decisions or judgements?

Chapter 5 addresses the main research question from the viewpoint of the literature:

- How does information system – supported decision-making get affected by biases?

Empirical work

Chapter 6 consists of a case study investigating the occurrence the bias of regression effects in exception reporting for decision-makers in a government department. The main research question is studied here from a practical point of view.

The majority of the empirical studies reported in the literature around the occurrence of biases, and in particular biases in information systems, are laboratory experiments. In these experiments, information systems are tailor-made to test the responses of experimental and control groups in an environment isolated from ‘noise’ factors that could cloud decision-making, such as distractions of attention, social or political influences or data errors.

Apart from the fact that the effort and time of setting up such an experiment would be practically undesirable given the scope of the present study, a case study from real life is believed to shed different light on the subject. As will be shown in chapter 4, the use of lab experiments to study biases has received a large amount of criticism. Several of these criticisms imply that the effect of biases is amplified in a lab environment. At least some of this criticism can be refuted by studying biases in a real life context.

Du Plooy (1998) gives an overview of the advantages and disadvantages ascribed to case study research. Case studies are viewed as useful when trying to understand a phenomenon within its context rather than seeking general laws about it. Multiple-case studies are preferred, although single-case studies are equally valid. Disadvantages of case studies include the lack of ability to make controlled observations, the non-replicability of the situation, the fact that it does not allow one to generalise, as well as the subjectivity of the observer/researcher. However, these criticisms are only valid under the assumption that the positivistic or natural science model is the norm for performing research. Du Plooy argues that "there is no need to legitimise social science research methods by modelling them on the methods of the natural sciences" (*op cit*, p 59). For example, subjectivity might be beneficial rather than undesirable.

2.7. Conclusion

The research methodology and research questions have been explained within the context of the information systems discipline, Burrel and Morgan's sociological paradigms as well as the process framework for deriving research questions. The empirical component of the research will be done by means of a case study. In contrast to most of the empirical work described in the literature, the case study is based on a real project rather than a laboratory experiment.

3. DECISION-MAKING AND HEURISTICS

3.1. Introduction

This chapter aims to introduce the notion of heuristics within the context of decision-making, of which it forms part. It contains a literature survey that will deal with the following research questions:

- How do people make decisions?
- What is a heuristic?
- How are heuristics used in the process of decision-making?

Much of the work on heuristics and biases implicitly assumes a rational model of decision-making similar to that of Simon (1977). Yet, this model is not the only means by which to understand decision-making. A comparison of Simon's rational model with other decision-making models enables one to see its limitations, which is relevant to the review of the theory on heuristics and biases.

3.2. Psychological approaches to decision-making

This section investigates the contribution of psychology to understanding the process of decision-making. Psychology is concerned with mental processes as well as human behaviour (Jennings and Wattam, 1994). Jennings and Wattam discusses four different schools of thinking in psychology. These schools will be described, as well as their implications for decision-making.

3.2.1. The Psychodynamic approach

The psychodynamic approach is based on the work of Sigmund Freud. Behaviour and psychological functioning are influenced by instinctive forces or 'drives'. These drives originate from the three facets of the psyche, namely the id (animal), ego (regulating self) and superego (conscience). These three facets have conflicting desires that lead to anxiety. Individuals behave in a way that reduces this anxiety, making use of certain defence mechanisms. Many of the

defence mechanisms work in a way that denies or distorts reality and interferes with objective perception. An added complication is that the mentioned mental processes (desires, anxiety, defence mechanisms) operate largely at a subconscious level. A consequence for decision-making is that a decision appearing irrational to an impartial observer could appear reasonable to the decision-maker, given their interpretation of reality.

3.2.2. The Behaviourist Approach

The behaviourist school differs from the psychodynamic approach in that it is only concerned with external behaviour that can be observed and measured, as opposed to internal mental processes. It views and treats psychology as a natural science. Behaviour is regarded as responses to stimuli, thus determined by environmental causes and experience rather than innate causes. For example, it is believed that behaviour will become repetitive if reinforced. Decision-making is regarded as the result of habitual and learned responses to environmental factors.

3.2.3. The Cognitive Approach

The cognitive approach recognises both internal processes and responses to the environment as influences on behaviour. The individual is perceived as an 'information processor'. Information is perceived from the environment and processed or made sense of by means of personal constructs. Information is perceived differently by different individuals and processed using different constructs. Two people can therefore come to totally different conclusions or decisions when given the same information.

3.2.4. The Humanistic Approach

Humanists differentiate themselves with their interpretive view: any opinion or fact is an interpretation of reality. Reality is thus constructed by the individual, going further (being more interpretive) than the concept of personal constructs described by the cognitive school. The humanist school emphasises a person's free will to act. This is in contrast with the views of behaviour as solely determined by internal mental processes or environmental stimuli. The influence of motivation on behaviour is also recognised. A decision-maker according to the

humanist school is more of a 'free agent' whose individual beliefs and interpretations are respected.

3.2.5. Conclusion

The two psychological approaches that appear to be favoured within decision-making literature are the behavioural (eg. Simon, 1956) and cognitive (eg. Hogarth, 1980) approaches. The psychodynamic approach, or Freud's work, being mainly applied in psychotherapy, appears to be of limited use. However, it is surprising that the humanist approach or the philosophy behind it (such as emphasising interpretation) is not more strongly represented, as it is for example in the work of Boland *et al* (1994).

3.3. Models of decision-making

3.3.1. Introduction

Various theories or views of decision-making can be found in the literature. The categorisation given below is based on the categories given by Keen and Scott Morton (1978), Huber (1981) and Das and Teng (1999). The last item, namely naturalistic decision-making, is fairly new and did not appear in the mentioned categorisations. Das and Teng's classification is by itself a meta-classification. Keen and Scott Morton considered the field of decision-making in general. Huber has written about decision-making in organisations, and Das and Teng's classification concerns strategic decision-making in organisations. The different views, theories or models are as follows:

- *The rational manager view* assumes a rational and completely informed decision-maker ('economic man') as described by neoclassical microeconomic theory around the middle of the previous century. Simon (1977) contributed to this stream of thought, but also pointed out some of its limitations.
- *The 'satisficing', process-oriented view* is based primarily on Simon's (1979) work on bounded rationality, admitting that the rational manager does not always have complete information, and that optimal choices are not always required. Das and Teng as well as Huber do not distinguish between the satisficing and the rational manager views.

- *The logical incrementalist view* involves a step-by-step process of incremental actions and keeps the strategy open to adjustment. Under Lindblom's (1959) disjointed incrementalism ('muddling through'), marginal, feasible changes are made, working from the status quo to solve existing problems rather than towards goals. Other researchers describe a process of 'muddling with a purpose' (Das and Teng, 1999).
- *The organisational procedures view* seeks to understand decisions as the output of standard operating procedures invoked by organisational subunits. March (1988) contributed to this theory. Huber names this view the 'program model', indicating that the decisions are pre-programmed in existing procedures as well as the routinised thinking of the people involved. Das and Teng refer to it as the 'avoidance mode' which views decision-making as a systematic process aimed at maintaining the status quo. On the other hand, Krabuanrat and Phelps (1998) regards this view in a positive light, namely as the use of codified organisational experience.
- *The political view* sees decision-making as a personalised bargaining process, driven by the agendas of participants rather than rational processes. People differ on the organisation's goals, values and the relevance of information. The decision-making process never ends, but it is a continuous battle between different coalitions. After one group wins a round of the battle, other parties might regroup or become even more determined to win the next round. Influence and power is wielded in a deliberate manner and to further self-interest. The goals of the coalitions are defined by self-interest rather than what is good for the organisation as a whole. Pfeffer (1981) is one of the major contributors on politics and power in decision-making.
- *The garbage can view* describes decision-making in an 'organised anarchy' and is based on the work of Cohen, March and Olsen (1972). Similar to the political model, it recognises a pluralist environment with multiple goals and agendas. However, it emphasises the fragmented and disjointed nature of decision-making rather than deliberate political behaviour.
- *The individual differences perspective* focuses the attention on the problem-solving behaviour of the individual manager, as influenced and differentiated by the manager's decision-making style, background and personality. It is only mentioned by Keen and Scott Morton.
- *Naturalistic decision-making* is concerned with investigating and understanding decision-making in its natural context. Real decision situations are analysed in detail in order to understand people's 'true' natural behaviour. Gary Klein's (1998) Recognition-Primed

Decision Model (RPD) describes such behaviour. The strong empirical foundation of naturalistic decision-making differentiates it from other descriptive models, such as the organisational procedures and political views.

What follows in the remainder of the section is a selection of decision-making models, providing examples for some of the categories listed above. Simon's (1977) rational model of decision-making is given in support of the rational view. Simon's (1956, 1979) work on bounded rationality is mentioned with regard to the process-oriented view. The Garbage Can model is described. As an example of naturalistic decision-making, the Recognition-Primed Decision Model (Klein, 1998) is described. The remainder of the categories are not further dealt with, for reasons stated as follows. In terms of the organisational procedures / program / avoidance views: it is believed the models for these views are given by the relevant organisational procedures themselves, and how they are implemented in practice. The individual differences view focuses on different perspectives and does not present a unitary model of decision-making. Regarding the political view, Huber mentions that despite a good many writings related to the view, no composite theory or model is yet available.

Hogarth's (1980) and Russo and Schoemaker's (2002) decision-making models are also discussed below since they will be used later when discussing heuristics and biases.

3.3.2. Simon's model of rational decision-making

First and foremost in terms of its continued influence on decision-making thought, is Simon's four-step decision model (Simon, 1977). These steps are closely related to the stages of problem solving described by John Dewey in 1910, namely:

What is the problem?
What are the alternatives?
Which alternative is best?

Simon performed his research on people who had to solve logical problems in a laboratory setting. Subjects were encouraged to think aloud while trying to solve problems, and their thinking patterns were analysed**.** This research confirmed his belief in the four-step model:

Intelligence: finding occasions for making a decision;
Design: inventing, developing and analysing possible courses of action;
Choice: selecting a particular course of action from those available; and
Review: assessing past choices.

According to Simon, executives and their staff "spend a large fraction of their time surveying the economic, technical, political and social environment to identify new conditions that call for new actions. They probably spend an even larger fraction of their time, individually or with their associates, seeking to invent, design, and develop possible courses of action for handling situations where a decision is needed. They spend a small fraction of their time in choosing among alternative courses already developed to meet an identified problem and already analysed in terms of their consequences. They spend a moderate portion of their time assessing the outcomes of past actions as part of a repeating cycle that leads again to new decisions. The four fractions, added together, account for most of what executives do." (1977, p 40)

Simon believed that even beneath terms such as intuition lay a logical problem-solving process. He claimed that if all the relevant information and enough processing power were available, a computer would be able to make decisions in the same way as a human being.

Simon's model can be viewed as a general model for rational problem solving. It also describes the steps that are taken for the calculation of the Subjective Expected Utility (SEU), a term used by decision analysts (Wright, 1984). Decision analysis aims at arriving at an optimum or best decision by maximising the SEU. The SEU is calculated by attaching to each alternative:

- an estimated probability of the likelihood that it will occur; and
- a weight or "utility" indicating the decision-maker's preference for that alternative.

The above two values are multiplied and the option with the highest value (highest SEU) is concluded to be the preferred alternative. Simon's model can and often is associated (eg. Kreitner and Kinicki, 2001) with the 'perfect optimality' kind of decision-making although the model is not restricted to such application.

3.3.3. Simon's contribution on bounded rationality

Simon (1956) compared psychological theories of adaptive behaviour with economic theories of rational behaviour. According to Simon, the theories of adaptive behaviour, such as learning theories, better describe observed or real behaviour. Furthermore, the adaptive behaviour falls short of the ideal of maximising given by the economic theories. In reality, decision-makers face uncertainties and have incomplete knowledge (Simon, 1979), and are thus not in a position to calculate the SEU. They also have limited computational powers.

He asserts that organisms adapt well enough to satisfice, while they do not optimise (Simon, 1956). On the one hand, an organism possesses limited information and limited computational capacity, thus limiting its behaviour. On the other hand, the task environment in which it operates allows for simplification of its behaviour. Simon explains the term 'satisficing' by means of a model of an organism searching for food. The organism only needs to maintain its energy levels by a regular intake of food. It has a limited capacity to consume and has no use for additional food. If it does not find food before its energy is depleted it will die. The organism searches the environment and stops at the first source of food it can find. Food is distributed randomly on a smooth surface. This means that the organism does not have to employ any specific search methods or means of travelling, it can just move randomly in search for food. In this case, the structure of the task environment influences the behaviour of the organism.

Simon (1979) more generally defines the concept of bounded rationality by the activities of searching and satisficing. A person has certain aspirations that need to be satisfied. Since all alternatives are not available upfront, a search is undertaken for alternatives that would satisfy the aspirations. Alternatives are found and evaluated in sequence, until a satisfactory alternative is found. The process of sequential searching and evaluating against aspirations until satisfied is called 'satisficing'.

Simon's (1956) organism is an artificial and highly simplified example of a decision-maker. Yet, it accurately reflects aspects of people's daily behaviour in decision-making: whether aiming towards meeting a sales target or keeping an adequate supply of milk in the fridge, having 'enough' is often more important than optimising.

3.3.4. The Garbage Can Model of Organisational Choice

In sharp contrast to the rational model is Cohen, March and Olsen's (1972) garbage can model of decision-making in an 'organised anarchy'. According to them, every organisation to some extent displays the characteristics of an organised anarchy. These characteristics are as follows:

- Ambiguity with regard to preferences and goals: "the organisation operates on the basis of a variety of inconsistent and ill-defined preferences" (Cohen *et al*, 1972, p 295). Preferences are discovered through actions rather than actions being governed by preferences. Also, decisions are made without consistent, clear goals.
- Unclear processes: the organisation's processes are not understood by its own members. "It operates on the basis of simple trial-and-error procedures, the residue of learning from the accidents of past experience, and pragmatic inventions of necessity" (*op cit*).
- Fluid participation: participation of the organisation's members varies for each activity and it changes with time.

Although organisations might be viewed as places where well-defined problems are solved, or where conflict is resolved through bargaining, they also provide procedures where participants arrive at an interpretation of what they are doing and have done while they are in the process of doing it. "From this point of view, an organisation is a collection of choices looking for problems, issues and feelings looking for decision situations in which they might be aired, solutions looking for issues to which they might be the answer, and decision-makers looking for work" (*op cit*). Thus, *"one can view a choice opportunity as a garbage can into which various kinds of problems and solutions are dumped by participants as they are generated"* (emphasis added). In the garbage can model, "a decision is an outcome or interpretation of several relatively independent streams within an organisation". The four streams are:

- Problems: all the various issues that occupy people's minds and require attention, whether work related or not.
- Solutions: a solution is a product that wants to be sold, an answer looking for a question.
- Participants: Participants come and go, and an entrance into one garbage can means an exit somewhere else. Attention varies due to other demands on the participants' time, rather than being dictated by the choice situation itself.
- Choice opportunities: The occasions during which decisions need to be made.

The four streams interact but have existences of their own. When a decision is made, the garbage can of the choice opportunity is removed. This might happen without having solved all or some of the associated problems in the garbage can. Since participants are the ones generating the garbage, or problems and solutions, the decision made is totally dependent on the make-up of the team of participants in the can.

According to Cohen *et al*, the garbage can does not resolve problems well. Yet, many problems are resolved despite living in an organised anarchy. They conclude that "the great advantage of trying to model garbage can phenomena is the possibility that the process can be understood, that organisational design and decision-making can take account of its existence and that, to some extent, it can be managed" (*op cit*).

3.3.5. Klein's Recognition-Primed Decision Model

Similar to Cohen, March and Olsen, Gary Klein (1998) is concerned with decision-making in its natural environment, rather than in a laboratory. He conducted numerous field studies and interviews with people who had to think and act under pressure, such as fire fighters, soldiers and nurses. He attempted to understand and describe the manner in which decisions really happened, and his models are thus descriptive and not prescriptive. Based on the analysis of over 600 real decision situations, Klein developed the recognition-primed decision model (RPD), shown in figure 3.1.

Central to the model is the decision-maker's ability to recognise a situation as being similar to a situation that was previously experienced. Part of what is recognised are the goals associated with such a situation, as well as important cues and what to expect. Decision-makers also recognise a course of action likely to succeed. The course of action is evaluated by means of a mental simulation, where the decision-maker visualises how the action is implemented.

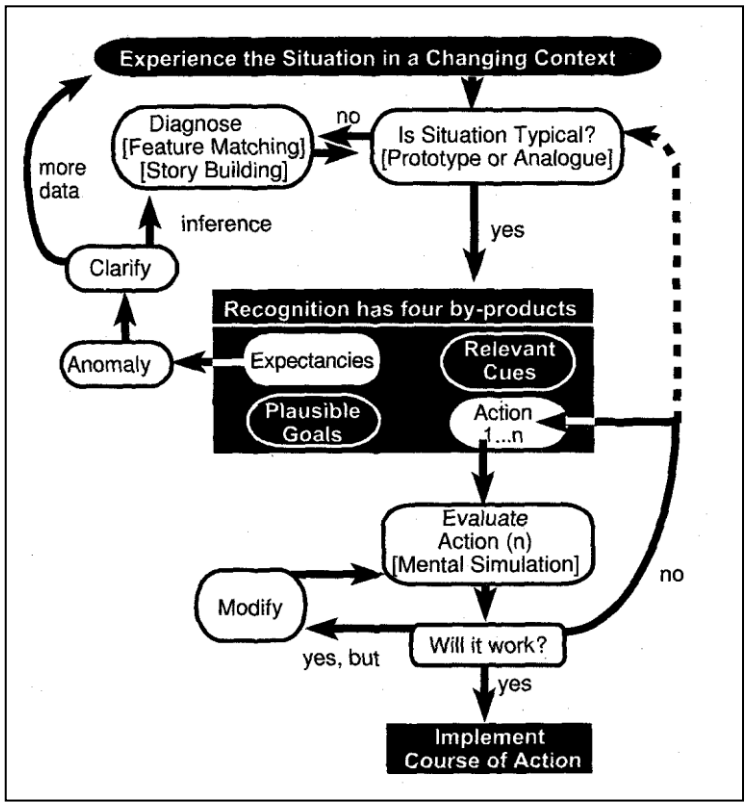

Figure 3.1: Integrated version of the recognition-primed decision model

(Klein, 1998)

The script is revised until the decision-maker is comfortable with it, after which it is implemented. All this can happen in a course of a few seconds. If a situation is not recognised as typical, more energy needs to be spent to diagnose the situation, and additional information will be collected. According to Klein, one of the key factors leading to good decisions is experience. The RPD model supports the notion that experience will increase the person's ability to recognise a situation.

Two primary strengths or 'sources of power' (a phrase used by Klein) in the decision-making process are intuition and mental simulation. Intuition is about recognising patterns, the first step in the RPD model. According to Klein, many people associate intuition with some supernatural or extrasensory perception. Yet, where he analysed instances of acting on 'unexplainable' moments of clarity, it was found that logical reasoning was made based on information from the

environment. The only difference was that it happened subconsciously, or so fast that the incumbent did not have time to consider the facts deliberately. Intuition will improve if one has more patterns in memory (from one's experience database) to match with the present situation. Klein states that one cannot transfer this experience by means of rules to a novice. A pattern needs to be understood in a particular context. A better way to build someone's pattern database is by means of storytelling or case studies.

The second important source of power is that of mental simulation, "the ability to imagine people and objects consciously and to transform those people and objects though several transitions, finally picturing them in a different way than at the start" (Klein, 1998, p 45). Mental simulation can be used to build a plausible sequence of events about either the past, the present or the future. A jury needs to put together the facts of a court case into a convincing story of the past. A decision-maker trying to diagnose a situation needs to find a story to explain the present situation. When deciding how to implement a course of action, the future needs to be visualised step by step. If a set of facts or actions is coherent, it will lead to a plausible story. Otherwise, it will lead to some anomaly or contradiction, which one will need to either explain away or change the story line.

3.3.6. Hogarth's model of judgement

Hogarth's (1980) model of judgement is a vehicle for showing where and how decision-making bias occurs. At this stage, only the model will be discussed. The model is a system comprising three components that mutually influence one another. These elements are: the decision-maker (or person performing the act of judgement); the actions resulting from choices of the decision-maker; and the decision-making environment or task environment. The 'schema' on the diagram (Figure 3.2) refers to the decision-maker's schema. It represents the person in terms of their beliefs concerning the task environment as well as their perception of the task. The schema includes the person's memory of previous experiences. From Giddens' structuration theory (Walsham, 2001), one could say that this is the person's 'structure' (their tacit beliefs, values and knowledge) as pertains to the task at hand. Within the person's schema, information is acquired, both from memory and the task environment. The information is processed and an output is obtained. The output is the person's decision. After the decision has been made, activity moves from structure to agency, and the person's decision is implemented within the task environment. The outcome is the effects of the action, which may have consequences beyond the immediate

task environment. From the outcome there is feedback to both the person's schema/structure and the environment.

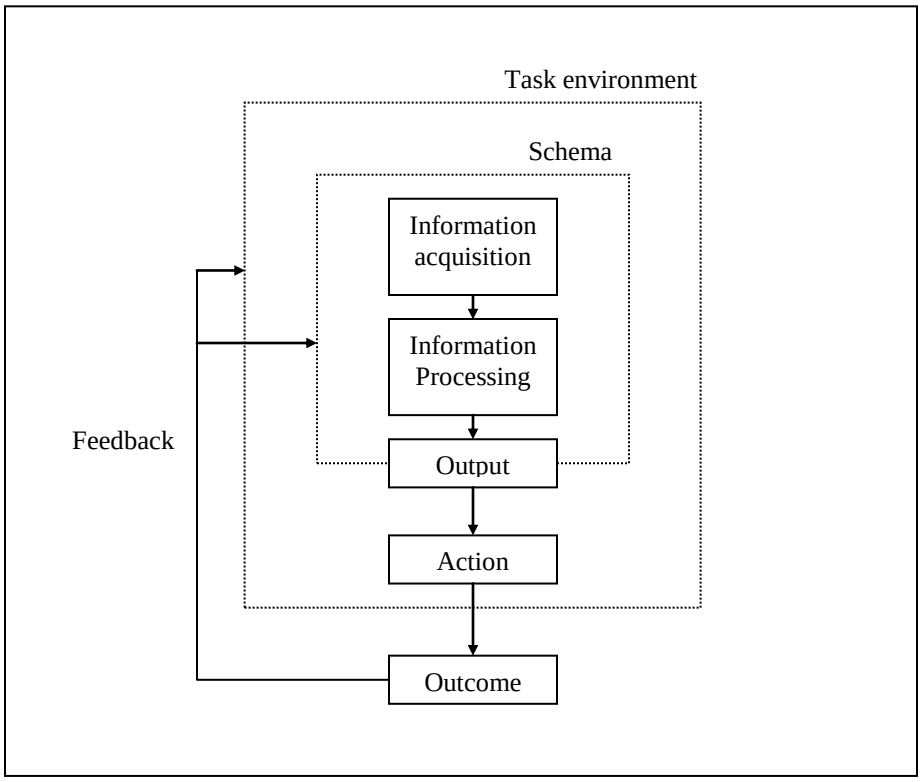

Figure 3.2: Hogarth's conceptual model of judgement

(Hogarth, 1980, p 157)

Hogarth's model falls within the cognitive school of psychological thinking: it recognises the role of internal mental processes as well as responses to the environment on human behaviour. The individual is viewed as an information processor. Information is perceived from the environment and processed or made sense of by means of personal constructs (similar to the concepts of a schema and a person's structure, as mentioned above). Since information is perceived and processed differently by different individuals, two people can come to totally different conclusions when given the same information to base a decision on.

According to Hogarth's model, people's information processing (and thus decision making) capacity is limited by the following factors (1980, *op cit*):

- Selective perception;
- Sequential information processing;
- The use of short cuts or heuristics; and
- Limited memory capacity.

3.3.7. Russo and Schoemaker's model

Russo and Schoemaker (2002) are researchers and management consultants who wrote a text specifically to assist people in making better decisions and reduce decision-making bias that occur in the process of decision-making. They believe that people go through the following steps when making decisions:

- *Decision-framing*: The way in which a problem is perceived and aspects of a situation are prioritised, is influenced by the decision-maker's background, world view, mental models or even the rut they are in. Russo and Schoemaker include as part of this step the generating of options or alternatives.
- *Gathering intelligence*: The gathering of information to assist the decision-maker to choose between options. This includes validation of knowledge.
- *Choice*: The use of an appropriate choice approach for the decision that needs to be taken. Methods can vary from intuition to detailed analysis.
- *Learning from experience*: Learning is more that just the accumulation of experience. Timely learning, right from when the decision is made, can lead to refinements to the decision and its implementation.

Similar to Russo and Schoemaker's contribution is that of Hammond, Keeny and Raiffa (1999). They provide guidelines for better decision-making, also following a number of steps:

- Frame the problem;
- Specify objectives;
- Create alternatives;
- Understand the consequences;
- Grapple with the tradeoffs;

- Clarify uncertainties;
- Take account of risk tolerance; and
- Consider linked decisions.

Both the models of Russo and Schoemaker and Hammond *et al* are extensions of and similar in nature to Simon's rational model.

3.3.8. Decisions vs. judgements

Two different terms are used in the above-mentioned literature to describe the process of decision-making, namely 'decision-making' and 'judgement'. The term 'choice' has also been observed (eg. Hogarth, 1980) although it is not used as frequently. The heuristics and biases literature uses the word 'judgement' (eg. Tversky and Kahneman, 1974). In the preface to their first collected readings on heuristics and biases, Kahneman *et al* state: "Although the boundary between judgement and decision-making is not always clear, we have focused here on judgement rather than on choice. The topic of decision-making is important enough to be the subject of a separate volume" (1982, p xii). However, no definition of 'judgement' as opposed to 'decision-making' was found in the mentioned literature. The following definitions have been sourced (Online Dictionary, 2003):

Judgement: "The cognitive process of reaching a decision or drawing conclusions; the capacity to assess situations or circumstances and draw sound conclusions." (*op cit*). The term judgement is further defined in a legal context. It appears to be implicitly assumed that values or principles are used in the drawing of a conclusion.

Decision: "A position or opinion or judgement reached after consideration" (*op cit*). A decision can also refer to a cutting. "Decision is a cutting short. It implies that several courses of action have been presented to the mind, and that a choice is now finally made" (*op cit*).

The above definitions do not provide a clear explanation of why some decision researchers would prefer the term 'judgement' above 'decision' or 'decision-making'. What does the context of the literature suggest? In the heuristics and biases literature, where the term 'judgement' prevails, decision-making is studied in a very specific context. People are being given multiple-choice questionnaires in a laboratory environment and have to select the options they believe are correct.

The term 'decision-making' is used in a much wider context in literature covering various disciplines. It might be that the early authors of the heuristics and biases literature appropriated the term 'judgement' to distinguish between the specific context of decision-making as they apply it, and the wider context of decision-making as it is used elsewhere, as implied above by Kahneman *et al* (1982).

In this study, the terms 'judgement' and 'decision-making' are used interchangeably. However, when authors are cited that use the term 'judgement', it will be attempted to use the term 'judgement' in reference to their work, and vice versa.

3.3.9. Concluding remarks

A number of striking differences between the various decision-making theories or models can be observed. Some of these are:

Decision-making as a choice between alternatives

Simon's rational model, Russo and Schoemaker's model as well as decision theory in general (eg. Wright, 1984), assume that a choice is made between simultaneously presented alternatives. Decision analysis goes further by making use of the concept of the SEU to calculate the best alternative. Theory on bounded rationality refers to a sequential evaluation of options that continues until a good enough option is found. In Klein's model, people under time pressure do not have the luxury to reflect upon options, and they intuitively decide on a course of action. The Garbage Can model displays decision-making as so haphazard that no systematic comparison of options occurs. The organisational procedures view describes decision-making as the implementation of a set of organisational rules, thus no choices are made between alternatives. March (1988) describes how options are evaluated sequentially in organisations, by being discussed and/or implemented one at a time. The political model implies that opposing parties are each 'pushing' their own preference(s) and are not interested in dealing with a wider array of alternatives. Hogarth's model also processes information in sequence, although information processing can mean the choice between alternatives. It appears as if the models or theories that support the simultaneous evaluation of alternatives, as does the rational model, are in the minority.

The role of intuition

Simon (1977) and Klein believe that intuition is a rational process that has not been made explicit. Simon, who associates the use of intuition with 'nonprogrammable' decisions, has attempted to rework nonprogrammable problems into programmable ones. Simon asserts that the portion of the iceberg of reasoning that is 'below the surface' is similar to what we see above the surface. Klein firmly believes in the power of intuition and describes how it should be nurtured. Others such as Russo and Schoemaker believe that use of the term intuition is often an excuse for a decision not well thought through, and that intuition is only applicable in cases where experienced people deal with repetitive tasks.

The rationality of the human decision-maker

The rational model of Simon promotes the ideal of a rational decision-maker. This rationality is taken to the extreme if the four steps are used as a basis for calculating the SEU. Hogarth (1980) and Russo and Schoemaker also believe in an ideal of a best decision that can be rationally arrived at. In the theory of bounded rationality, a decision-maker does not optimise but still behaves rationally. Klein's decision-maker does not deliberately analyse, yet it is possible to justify decisions made. March does not believe that the behaviour of the organisation or its decision-makers is rational. In the political model, people's behaviour and preferences are politically motivated rather than rational. Hogarth, Russo and Schoemaker and Tversky and Kahneman (1974) all believe in the ideal of Simon's rational model. Their work on heuristics and biases shows how humans fail to display rational or 'correct' behaviour.

3.4. Heuristics

The concept of heuristics has been referred to in the section on decision-making. Here, it will receive more explicit attention.

3.4.1. What is a heuristic?

It appears as if the term 'heuristic' is used in mainly two different contexts. The first context is a purely mathematical or analytical one, where heuristics refer to "nonexhaustive strategies for

searching through a problem space" (Klein, 1999, p 307). The heuristic is a formula or algorithm that does not guarantee an optimal solution yet satisfies some minimum criteria. It is used where exhaustive search strategies are not feasible because of their computational demands. People using heuristic methods in the mathematical context normally know exactly the error range (or possible deviation from optimum) that they are dealing with, and are more than willing to accept near-optimality because of the associated benefits.

The second context for the term heuristic is the one used by decision researchers and psychologists. It is also the one that will be used for the purpose of this study. Here, heuristics "refer to a rule of thumb, or informal reasoning strategy, as opposed to a mathematical formula that can be calculated" (Klein, 1999, *op cit*). Hammond *et al* (1998, p 47) refer to heuristics as "unconscious routines to cope with the complexity inherent in most decisions." According to Jennings and Wattam (1994, p 106), heuristics are "simple generalisations or guidelines that individuals use to reduce mental effort in processing information. They are broad categories which simplify and speed up the process of interpreting or classifying new information."

Marsh (2002) emphasises the often unconscious use of heuristics. According to him, "we can think of heuristics as being the rules by which knowledge (both internal and external) is filtered and selected before being passed to the decision-taking realm that is consciously sensed" (p 51).

According to Jennings and Wattam (1994), the concept of heuristics in decision-making has been developed in the cognitive school of psychology, as described under Hogarth's model above. Jennings and Wattam relate the ideas of "schema", "personal construct" and heuristics. They view all of these as "ready made frameworks" that individuals apply when processing information. These frameworks are means of classifying, organising and interpreting information, and hence present opportunities for introducing bias. Marsh (*op cit*) also views heuristics in an information-processing context, and refers to search heuristics, assessment heuristics and selection heuristics.

Simon's work on bounded rationality, in which heuristics play an important role, has been influenced by the behaviourist rather than the cognitive school of thought. The organism searching for food (Simon, 1956) will develop or adapt its behaviour until it has found a satisfactory or 'satisficing' strategy, given the limits of its own capability, the information that it

has as well as the conditions of the environment. It will then continue to follow that strategy, which can be labelled a heuristic.

3.4.2. Examples of heuristics in decision-making

To illustrate the use of heuristics in decision-making, Kahneman *et al*'s (1982) categorisation of heuristics will be presented, the first two items as discussed in Kreitner and Kinicki (2001):

- The availability heuristic refers to the tendency to base decisions on more readily available information. This might pertain to recent events (which are fresh in one's memory), salient or unusual information and information that evokes strong emotions. For example, information from an article read in the newspaper before going to work might influence a person's decision-making later in the morning more strongly than other relevant but less easily retrievable facts.
- The representativeness heuristic is used to estimate the probability of an event occurring. The estimate is influenced by outcomes of similar events in the past. For example, a manager decides to hire a graduate from a particular university because previous appointments from that university were successful.
- Adjustment and anchoring: this heuristic pertains to cases where a value needs to be estimated or predicted. A suggested initial value might be given or guessed, and then adjusted upwards or downwards. For example, if sales for the next financial year need to be forecasted, people will tend to base or anchor their estimate on the forecast for the current financial year.

Another application of heuristics in decision-making is given by Krabuanrat and Phelps (1998). They support the organisational procedures view of decision-making (Keen and Scott Morton, 1978), as discussed under the models of decision-making. They regard following organisational routines in a positive light and do not share the sentiments of, for example, the associated 'avoidance mode' (Das and Teng, 1999). According to them, organisational routines have been built up by the organisation in the course of experience and would have proven themselves to be robust. They refer to such routines as heuristics. Krabuanrat and Phelps believe that these heuristics are often applied unknowingly, or without decision-makers being able to articulate them.

The heuristics found by Krabuanrat and Phelps are categorised as follows:

- Simplification: intentionally ignoring aspects of the decision in order to reduce the cognitive load on the decision-maker.
- Reference to past cases: a situation is recognised to be similar to a past case, and the latter's decision outcomes are taken as guidelines.
- Imitation: the decision-making behaviour of other organisations is imitated.
- Risk aversion: where risk is identified, small-scale experiments are carried out in order not to repeat failures that have been observed in similar situations.
- Satisficing: an acceptable solution is sought rather than an optimal one, and the search for alternatives is stopped once an acceptable solution is found.
- Cooperation: knowledge and resources are pooled and risk shared with customers, suppliers or competitors.

Chase *et al* (1998) believe that the use of heuristics is an important component of boundedly rational behaviour. They refer to Simon's bounded rationality as behaviour being shaped with a scissors whose two blades are the structure of the task environment and the cognitive capabilities of the decision-maker. One class of heuristics that makes use of both Simon's blades is termed the 'fast and frugal' heuristics. 'Fast' means that the answer is found soon, and 'frugal' refers to the fact that it does not require a lot of information. As an example of a fast and frugal heuristic, Chase *et al* describe 'Take the Best'. It is based on the so-called recognition heuristic. If one has to choose which of two cities is the largest, then choose the one you have heard of before. If you have heard of both, choose the one that you know more about with respect to a second variable (for example, you might know the one has an airport but you are not sure whether the other one has). Continue until a city has been selected, then stop. Chase *et al* provide evidence that the 'Take the Best' heuristic can do as well as multiple regression when required to infer or predict a variable, while at the same time requiring less information or being more frugal. 'Take the Best' borrows from both the sets of heuristics given by Kahneman *et al* and Krabuanrat and Phelps above. The recognition heuristic on which 'Take the Best' is based is related to the availability as well as the representativeness heuristics from Kahneman *et al*'s list. From Krabuanrat and Phelps' list, the 'simplification' and 'satisficing' rules are used in the 'Take the Best' algorithm. Chase *et al* argue that 'Take the Best' works by capitalising on the environmental regularity that, in the mentioned example, big cities are more likely to be well known. With the latter argument, they also involve the second blade of Simon's scissors, namely that the structure of the environment

encourages boundedly rational behaviour. The presence of the first blade, namely the limits of people's cognitive capabilities, has been shown in the use of the heuristic. 'Take the Best' is an example of heuristic behaviour that has been formalised into an algorithm. In that sense, it differs from the belief that heuristics are often applied in an informal, unknowing and less specified way, being shaped by the context it is used in as well as by the traits of the decision-maker, as can be understood from Marsh (2002).

Marsh (2002) discusses the use of heuristics are used in a social context. Social behaviour often follows recognisable patterns without actors even being aware of the heuristics applied. Examples are the use of nonverbal cues in order to communicate intent, and social imitation. Social imitation includes a wide range of behaviour, from dressing the part when aspiring to join a group, to focusing on similarities when establishing a friendship, to a general social rule of 'when in doubt, follow others'. The heuristic of imitation has also been identified by Krabuanrat and Phelps, as shown above.

3.5. Conclusion

The focus of this chapter has been decision-making. Various views on decision-making have been presented and these have been illustrated by a variety of decision-making models. In doing so, the following research question has been addressed:

- How do people make decisions?

The rational model of decision-making has been observed to dominate the decision-making literature, despite the availability of equally valid alternative models, such as Cohen *et al*'s Garbage Can model, and Klein's Recognition-Primed Decision Model. The rational model takes preference especially in literature with prescriptive views on decision-making, eg. Hogarth, Russo and Schoemaker, and Hammond *et al* (1999).

Simon's model of bounded rationality suggests a modification to the rational model. Simon, influenced by the behaviourist school of psychology, has shown how *heuristics* are used when people 'satisfice' or go for good enough rather than optimal decisions. The role of heuristics within the cognitive school of psychology has been shown by Jennings and Wattam. It was

attempted to define the concept of heuristics from the literature, using definitions from Klein, Hammond *et al* (1998) and Jennings and Wattam. In doing so, the next research question:

- What is a heuristic?

has been addressed. The question

- How are heuristics used in the process of decision-making?

has implicitly been addressed by pointing out the reference to heuristics within the decision-making literature and by giving some examples of the use of heuristics within decision-making. In the examples quoted, the usefulness of heuristics was recognised by people supporting either the bounded rationality or organisational procedures view of decision-making, both of which are descriptive. In contrast, the prescriptive literature on decision-making (based on the rational model) argues that heuristics lead to decision-making bias. From this, one can see how assumptions on decision-making influence the approach towards heuristics, and whether heuristics are regarded in a positive or negative light (ie. as a source of biases). Biases have not been discussed to date: the topic of 'heuristics and biases' will be taken further in the next chapter.

4. HEURISTICS AND BIASES

4.1. Introduction

This chapter will continue the discussion on heuristics in decision-making, with an added emphasis on the biases that could be attributed to the use of heuristics. The research question:

- How are heuristics used in the process of decision-making?

that has been addressed in the previous chapter will receive further attention. In addition, the following research question will be addressed:

- What is the effect of heuristics on the outcomes of decisions or judgements?

The literature suggests that although heuristics are very useful tools in decision-making, they leave room for biases to occur. Tversky and Kahneman (1974) were the first psychologists to systematically measure the occurrence of bias in human judgement. Following their lead, numerous contributions have been made by decision researchers to what can be called the "heuristics and biases" literature (Klein, 1998). In response to the latter research question, the literature on heuristics and biases will be reviewed.

Heuristics and biases will be discussed with reference to human information processing and rational decision-making, since these are the two assumptions of the heuristics and biases literature: that humans are information processors, and that rational decision-making is the norm against which they should be measured. In the next chapter, it will be argued that these assumptions also hold for information systems. The occurrence of heuristics and biases within the social context of decision-making will also be discussed. The latter was found to be a substantial field of research in the psychology literature, but the social context of decision-making is not accounted for under the assumptions of information processing and perfect rationality. The sometimes un-rational behaviour in a social context can either be used as a source to describe even more human biases, or it can provide a basis for criticising the heuristics and biases literature. The discussion concludes with a review of the criticism against the heuristics and biases literature.

4.2. Heuristics and Biases

4.2.1. Tversky and Kahneman's work

Kahneman *et al* (1982) explain that there were a number of driving forces or influences behind Tversky and Kahneman's (1974) initial work on heuristics and biases. The first was Simon's (eg. 1956, 1979) contribution on bounded rationality, according to which decisions are based on limited rather than complete information, on the use of judgemental heuristics and outcomes that are satisficing rather than optimal. The second driving force behind Tversky and Kahneman's work was a publication by Paul Meehl in 1954, where he provided evidence for discrepancies in the judgement of experts. He showed that people's impressions of how well they reason were not in line with their actual performance. The third influence on Tversky and Kahneman's work was the introduction of Bayesian statistics into psychological research. It provided a way to quantify optimal performance under uncertainty, with which human judgement could be compared. With the mentioned three influences in place, it was known that people's judgements were not always totally rational, and there was a way to measure deviation from the optimum. Tversky and Kahneman did just that: they set up a variety of statistical experiments where people had to predict outcomes, with the aim of 'measuring' their rationality.

Tversky and Kahneman (1974) set up experiments where people had to assess the probabilities or likelihoods of uncertain events. For example, in considering tosses of a coin for heads and tails, people were asked whether they thought the sequence H-T-H-T-T-H more likely than the sequence H-H-H-H-T-H. Tversky and Kahneman performed numerous such experiments on people with different levels of training in probability theory. They came to the conclusion that "people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgemental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors" (1974, p 1124). They found people to use three general types of heuristics, namely representativeness, availability, and adjustment or anchoring. Each of these was found to lead to a number of decision-making biases. Later on, they added another heuristic, called problem framing (Tversky and Kahneman, 1981). A short description of the heuristics was given in the previous chapter, and will be repeated here with definitions as well as a list of the biases associated with each:

Representativeness heuristic

How well does an instance or subset represent a bigger class? Tversky and Kahneman found the following biases associated with representativeness (Harvey, 1998):

- Insensitivity to prior probability of outcomes (not applying Bayes' rule);
- Insensitivity to sample size (the law of small numbers): smaller sample sizes normally have larger deviations than people expect;
- Misperceptions of chance: see the heads and tails example in the previous paragraph. A random process does not guarantee a uniform distribution among a small subset; and
- Misconceptions about regression towards the mean, which results in undue attention given to outliers.

Availability heuristic

Ease of recollection of a past event makes people think that such an event is more likely to occur than a less easily recollected event. In particular, the following sources of bias arise from availability (Harvey, 1998):

- Familiarity: If someone is familiar with a certain class of events, a higher likelihood of occurrence will be attached to it, and vice versa;
- Salience: events evoking strong emotions are more easily recalled;
- Recentness: a more recent event is easier recalled;
- Effectiveness of the search set: the way in which we store information makes some items easier to retrieve from memory than others, and they are thus regarded as more likely;
- Illusory correlation: some classes of events may be strongly associated in a person's mind, even if they have little correlation; and
- Ease of scenario construction: events that are easy to imagine seem more likely to occur.

Adjustment and anchoring

An expected value is estimated in relation to a given initial value. The danger is that the initial value might not be justified (it might be a guess or suggestion) with the result that a slight upward

or downward adjustment of the initial value is still way off the mark. Overly narrow confidence intervals are also assumed in relation to the initial value.

Problem or decision framing

A 'decision frame' refers to the decision-maker's conceptions or beliefs about a choice. The frame is a result of the formulation of a problem as well as a person's norms, habits and personal characteristics (Tversky and Kahneman, 1981). Different frames would provide different views on the same problem. Rational decision-making requires that a change of frame does not alter the decision made. However, it was found that people's preferences differed significantly when problems were framed (or worded) differently.

4.2.2. Biases related to information processing

Building on Tversky and Kahneman's work, Hogarth (1980) wrote a text for decision makers not schooled in decision theory, where he discusses biases in dealing with information and how to overcome these. Hogarth subscribes to the cognitive school of psychology that views people as information processors. He developed an extensive list of biases and categorised them in terms of information acquisition, information processing, output and feedback. Tversky and Kahneman's biases are included in the list (Table 4.1), but phrased or grouped slightly differently.

In Hogarth's conceptual model of judgement, the influence of internal mental processes as well as inputs perceived and processed from the task environment is acknowledged to determine the behaviour of a decision-maker. The model was described and drawn in Chapter 3. There, it was also mentioned that information is perceived differently by different individuals and processed using different mental constructs. Two people can therefore come to totally different conclusions when given the same information to base a decision on (Jennings and Wattam, 1994). According to Hogarth's model, people's information processing capacity is limited by the following factors, leading to bias in decision-making:

- Selective perception;
- Sequential information processing;
- The use of short cuts or heuristics; and
- Limited memory capacity.

Stage of information processing	Bias/source of bias	Description
Information acquisition	Availability	Ease of recollection of a situation deems a similar situation more likely.
	Selective perception	Information is filtered by one's own experience, anticipations/expectations and preconceived hypotheses. Conflicting evidence disregarded.
	Frequency	Information on absolute number of occurrences used (eg. successes only), rather than a relative frequency (successes vs. losses)
	Concrete information	The experience of eg. a friend carries more weight than abstract information such as consumer reports.
	Illusory correlation	The belief that two variables co-occur (based on prior experience) when in actual fact they don't.
	Data presentation	The manner in which data is presented can be misleading (eg. what is shown and what is omitted, the sequence of display, mode of display, scales used for graphs etc).
Information processing	Inconsistency	Inability to judge a set of cases consistently (eg. in personnel selection, grading of papers).
	Conservatism	Failure to revise opinion after receipt of new information.
	Non-linear extrapolation	Inability to extrapolate eg. exponential growth; tendency to underestimate joint probabilities.
	Heuristics used to reduce mental effort:	
	Habit/rules of thumb	Choosing an alternative that has proved satisfactory before.
	Anchoring and adjustment	Prediction based on adjustment of initial value (a clue or a guess).
	Representativeness	Judging an event/person based on the characteristics of the class to which it belongs.
	Law of small numbers	A specific case of representativeness: small samples of data assumed to represent their population.
	Justifiability	A seemingly rational rule used to justify a decision.
	Regression bias	Undue attention given to outliers, ignoring law of regression towards the mean.
	"Best guess" strategy	Ignoring a number of uncertainties and basing guess on likely outcome of less uncertain variables.
	Bias associated with the decision environment:	
	Complexity	Complexity caused by time pressure, information overload or distractions reduces ability to process information.
	Emotional stress	Stress can lead to panic judgements instead of careful assessment of information.
	Social pressures	Social or group pressure can distort judgement.
	Bias associated with information sources:	
	Consistency of information sources	A large volume of consistent (including redundant) data increases confidence but not accuracy of predictions.
	Data presentation bias	As discussed under information acquisition.
Output	Question format (decision framing)	The wording or format of the problem posed influence the way that people respond to it.
	Scale effects	Eg. the scale chosen for a questionnaire affects responses.
	Wishful thinking	A desired outcome is rated to be more likely.
	Illusion of control	Planning or forecasting can induce a feeling of control over the future.

Stage of information processing	Bias/source of bias	Description
Feedback	Outcome irrelevant learning structures	Only a selection of the outcome of decision is fed back, eg. performance of rejected candidates is not known.
	Misconceptions of chance	Gambler's fallacy: eg. successive appearance of reds on roulette table makes people believe black is more likely with next round.
	Success/failure attributions	Tendency to attribute success to skill and failure to bad luck.
	Logical fallacies in recall	Inability to recall details makes people "reconstruct" events.
	Hindsight bias	In retrospect, people are not surprised about what happened in the past.

Table 4.1: Biases in information processing

(adapted from Hogarth, 1980, pp 166 – 170)

Hogarth's list of heuristics and biases was the most complete of what could be found in the literature. In addition, his categorisation of biases in terms of the stages of information processing provides a useful framework for investigating the occurrence of bias in information systems.

4.2.3. Biases related to the decision-making process

Tversky and Kahneman's work suggests that people do not aim at maximising utilities in their judgements. The assumption behind utility maximisation, or calculating the SEU as discussed in Chapter 3, is that the rational model of decision-making is followed. The literature that was studied concerning biases in decision-making, namely Das and Teng (1999), Russo and Schoemaker (2002) and Hammond *et al* (1999) all use variations of the rational model as a departure point. The biases described in this literature are all deviations from the 'correct' application of the rational model.

Hammond *et al* distinguish between two kinds of errors related to the rational decision-making process. The first set of mistakes is process mistakes, or not following the suggested steps of decision-making correctly, and the second is errors as a result of 'psychological traps'. Strictly speaking, the second kind is what Tversky and Kahneman as well as the psychological literature on heuristics and biases refer to. Yet, Das and Teng and Russo and Schoemaker write about both kinds of errors interchangeably. In the discussion that follows, an attempt will be made to focus on the psychological traps, although there is an argument for generally referring to typical mistakes in decision-making as biases.

For each step of the rational model, the literature mentioned above discusses possible biases that can occur while executing that step, in some cases with suggested remedies for improvement. A summary of the listed biases will be provided. The summary is based on Russo and Schoemaker's work, and supplemented by contributions from Das and Teng and Hammond *et al*. Most of these biases are already mentioned in the discussion of Tversky and Kahneman's and Hogarth's work above.

Decision framing

Russo and Schoemaker refer to the *limitation of frames* or the way people look at the world (referred to as problem framing bias by Tversky and Kahneman; selective perception by Hogarth). Information is filtered by one's frame or worldview and preconceived hypotheses about the problem are used. An incorrectly framed or stated problem can lead to the wrong problem being solved, and to limitations on the objectives considered as well as the number and variety of alternatives generated.

Intelligence gathering

The *overconfidence bias* refers to a situation where the uncertainty related to an alternative is underestimated and too narrow confidence intervals used. It is often accompanied by the anchoring bias (Tversky and Kahneman; Hogarth): the value of an unknown variable is guessed and the confidence interval obtained by slight upward and downward adjustments of the guessed value. Issues around the decision are considered with false confidence, and the SEU can be calculated incorrectly. Overconfidence can also result from having large amounts of data/information, even if the data are incorrect or contain redundancies.

The *confirmation bias* discussed above can result in limited or selective data (that confirms one's hypotheses) being collected in the intelligence gathering exercise.

The *availability bias* (Tversky and Kahneman) can lead to information being gathered based on the ease of it being collected or recollected, recentness, salience and plausibility.

When a number of variables are estimated over-conservatively (just to be safe), the cumulative effect of the conservative estimations may lead to an unrealistic joint estimate. This is referred to as the prudence bias (Hammond *et al*).

Hammond *et al* warn against the trap of seeing correlation in gathered data (or dependence between variables) where none exist. This is referred to as the *illusory correlation bias* (Tversky and Kahneman; Hogarth).

Another trap that Hammond *et al* warn against is the *surprised-by-surprises bias.* The fact that an event is unlikely does not mean that it cannot or will not occur. Also, the joint risk of a number of unlikely events equates to a higher likelihood that at least one of them will occur.

Choice

All the biases mentioned by Hogarth under the *'Information Processing'* heading in Table 4.1 can influence the outcome of a choice. Russo and Schoemaker discuss the *biases that are associated with methods of choice*. If only intuition is used, it can lead to inconsistency in judgement. If decision-making is governed by rules (such as whether a bank should approve a loan), only a selection of information, namely what is contained in the rules, are used to base a decision on. When circumstances change, the rules can become outdated. More sophisticated methods such as the numerical comparison of alternatives based on decision analysis, for example the calculation of the SEU, can create the false confidence that the future can be predicted from analysing historical data. Also, when alternatives are weighted, the weights can be (and often are) rigged in order to make the results conform to prior hypotheses. Buchanan and Corner (1997) provide evidence that the anchoring bias is present in multiple criteria decision analysis. The most sophisticated decision method discussed by Russo and Schoemaker is value analysis. In addition to the factors considered by decision analysis, it also weights variables against the tacit values (as opposed to facts) relevant to the decision. Total transparency in decision-making means that the inclusion and weighting of values are made public. Russo and Schoemaker mention that the explicit declaration of values can expose decision-makers, so that honesty can be compromised in the decision-making process or its communication.

A bias that can occur in a group decision-making process is the *groupthink* bias. Hogarth refers to it as 'social pressures'. It can be the result of people's desire to conform or be accepted and thus

claim allegiance to a group, or it can result from the mental laziness of members of a homogeneous group to question each other's claims.

Learning

Hogarth's list of biases associated with information feedback (Table 4.1) is relevant to the learning that Russo and Schoemaker advocate should follow choice and implementation. Of these biases, the ones also discussed by Russo and Schoemaker or elsewhere are given further mention.

After a choice is made, the decision is implemented. An implementation bias that does not acknowledge learning that has occurred during implementation is the *status quo bias*, also referred to as "escalation of commitment" (Kreitner and Kinicki, 2001). People realise that a wrong decision has been made, but have invested so much in the implementation that they do not want to acknowledge a mistake and cut their losses. They would persevere with the decision in the hope that they will recover the losses or regain control over the situation, just to accumulate further losses in the process. Nations participating in wars often land in this situation: Kreitner and Kinicki mention the Vietnam War as an example of escalation of commitment.

Russo and Schoemaker further refer to the *illusion of control* bias when reflecting on the success or failure of a decision. Hogarth refers to the same as "success/failure attributions". When people are successful, they like to attribute outcomes to their own skill. When they fail, they rationalise the outcome and blame it on, for example, chance or bad luck. Russo and Schoemaker stress that there are always uncertainties or 'chance' aspects to decision-making. No matter how thorough the process that was followed, a 'good' decision might fail or a 'bad' one saved by a fortunate turn of events. Thus, they believe that people should not over-rationalise or justify outcomes, but rather judge themselves by the process followed.

When evaluating a decision, it can happen that feedback or learning is not based on a complete set of the relevant information. Hogarth refers to *outcome irrelevant learning structures*. For example, no information is available on the financial performance of people whose loan applications were rejected. Russo and Schoemaker refer to the Challenger disaster as a case where learning was only done based on partial information. A number of engineers felt uncomfortable with the launch of the Challenger at very cold temperatures, since O-ring failures have previously occurred when launching at low temperatures. A graph plotting O-ring incidents

against temperature at launch did not convince decision-makers and the launch continued. However, a graph drawn afterwards that indicated the temperatures at which flights with incidents as well as incident-free flights took place showed that incident-free flights only occurred *above* a certain temperature. This more complete graph (with seemingly irrelevant information) would clearly have indicated the risk of launching at a temperature lower than the threshold value indicated on the graph.

Hindsight bias refers to the tendency to selectively recollect information that might have predicted an outcome and stating that the person "knew it all along", or "they should have seen it coming". Russo and Schoemaker go as far as arguing that the storage of historical data in the brain is altered after the outcome of a decision has been made available, such that the selection of information that could have predicted the outcome is linked to where the outcome of the decision is stored.

Another learning bias referred to by Russo and Schoemaker is the *self-fulfilling prophecy*. After a person has committed to a decision, their interests are at stake to prove themselves right. By committing additional effort in order to make their predicted success a real one, they make their own decision come true.

4.3. The social context of decision-making and the related occurrence of heuristics and biases

The majority of contributions on heuristics and biases in decision-making (eg. Russo and Schoemaker, Hammond *et al*) follow the lead of researchers such as Tversky and Kahneman and Hogarth in assuming a rational model of decision-making, and/or the cognitive view of humans as information processors. Yet, some of the biases mentioned in the previous section, such as groupthink and escalation of commitment, have a social component to them rather than being purely analytical. A stream of research was found to acknowledge the social context of decision-making. Some of these, for example Slovic *et al* (2001), Lipshitz *et al* (2002), Marsh (2002) and Chase *et al* (1998) investigate the use of heuristics in a social context, while researchers such as Dietz and Stern (1995) and Chase *et al* (1998) proceed to question the findings of earlier work on heuristics and biases that did not account for a social context. This section focuses on the contributions on heuristics and biases in a social context. Criticism of the heuristics and biases

literature for, among others, not acknowledging a social context is taken further in the next section.

Marsh's (2002) contribution on the use of heuristics in a social context was introduced in paragraph 3.4. One of the main heuristics discussed by Marsh is *social imitation*. Social imitation includes a wide range of behaviour, from dressing the part when aspiring to join a group, to focusing on similarities when establishing a friendship, to a general social rule of 'when in doubt, follow others'. Another social heuristic is to 'seek those who are already familiar' when looking for a partner or a person to employ. For example, one might be more comfortable going on a date with the family member of a friend, since one has knowledge about the person's background. Marsh stresses the usefulness of being able to solve socially complex problems with limited time, energy and computational resources. He acknowledges that speed and efficiency could lead to sacrifices in accuracy or errors in judgement. A bias that can occur when people rely on a group's preferences and socially transmitted stereotypes is unfair social prejudice. Marsh is however convinced that the benefits of social heuristics far outweigh the drawbacks.

Lipshitz *et al* (2001) describe the *one-of-us effect* that occurs when the outcome of a decision is evaluated. The one-of-us effect refers to the evaluator's identification (or not) with the decision-maker whose decision is evaluated. They describe a study where Arab and Jewish subjects had to judge the success of decisions made by either the Arab or Jewish underground in operations directed against the British authorities in Palestine. It was found that the Jewish subjects would, for example, judge the decision-making proficiency of the Arabs more harshly, even when outcomes of their decision-making were successful. On the other hand, they would be forgiving towards failures of the Jewish underground. Subjects were found to side with decision-makers they identified with, no matter what the outcome of decisions. Lipshitz *et al* conclude that the mentioned behaviour does not only violate rationality, but also hinders the ability to learn from experience, handicaps the distribution of rewards and punishments and weaken the authority of valid operating procedures. They also point to the fact that decision-making in the above circumstances cannot be researched without understanding of the socio-cultural context of decision-making.

Chase *et al* (1998) provide a number of illustrations of behaviour that is socially valid yet violate the rules of rational choice in decision theory. For example, in competitive situations, it is sometimes desirable to behave inconsistently in order to confuse an opponent. Or, at a dinner

party, one might out of politeness not take the last éclair – yet, if the supply is replenished, one might decide to take the same éclair and violate one of the basic principles of internal consistency, namely that a choice should be independent of other alternatives offered. In terms of social adaptability, it is more important to outwit an opponent or to be regarded as polite at a dinner party than to satisfy a decision theorist. Chase *et al* propose that the meeting of social expectations should be included in the definition of rationality, rather than to refer to the socially justified violations of classic rationality as 'biases'.

Slovic *et al* (2002) discuss how decision-making behaviour is mediated by the *affect heuristic*. According to them, images and experiences in memory have an 'affect' tag, indicating an emotion or feeling knowingly or unknowingly associated with it. The affect tag also accompanies anticipated events, either sounding an alarm or serving as an incentive. This is because affects related to past events are consulted in the process of making judgements. Affect serves as a cue, in the same way as imaginability or similarity does when the availability and representativeness heuristics are used. An affective impression is easier to come by than an analytical evaluation of a situation. Even if an analysis is performed, the affect tag influences decision-making: "…people base their judgements of an activity or technology not only on what they *think* about it but also on what they *feel* about it" (p 333). Slovic *et al* provide empirical support for the influence of the affect heuristic on decision-making. They also provide examples of how it can misguide people. One of the dangers of affect is that it is much more 'real' in the present than it is for the distant future. For example, people start smoking based on immediate pleasure, even though they can anticipate future problems. Once the problems of addiction and health risks are experienced, they cannot change their behaviour. A number of 'illogical' judgements ascribed to the affect heuristic are also shown through experiments, for example in the comparison of gains and losses, such as the saving of lives versus the loss of lives. These experiments are similar to the ones performed by Tversky and Kahneman (1981) in their description of problem framing, as discussed in section 4.2.1. It can be argued that the framing or wording of a question produces an affect on which decisions are based.

On the one hand, one can conclude that the affect heuristic is problematic in the sense that it leads people away from a purely rational decision-making process. On the other hand, it can be used to explain or justify the biases found in heuristics-and-bias experiments such as Tversky and Kahneman's (1981), and affect can be recognised as a valid component of reasoning.

4.4. Criticism of the heuristics and biases literature

The literature on heuristics and biases can leave the impression that humans are information processing machines that contain a few programming errors in their algorithms. Or perhaps, the code is too simple and leaves people unable to cope with all aspects of the real world. "It appears that people lack the correct programs for many important judgemental tasks… it may be argued that we have not had the opportunity to evolve an intellect capable of dealing conceptually with uncertainty" (Slovic, Fischhoff and Lichtenstein, quoted in Chase *et al*, 1998). Yet, a number of authors do not agree that perfect rationality or proficiency in information processing is the norm that people should be measured against, nor do some agree with the manner in which people's rationality has been tested in the heuristics and biases literature.

Contradictory findings regarding people's judgemental competence can be seen in the following comparison. Tversky and Kahneman tested hundreds of people in laboratory settings and found their judgement to be fundamentally flawed. Gary Klein (1998) observed and analysed the decision-making of numerous people in their natural settings, many of them by the nature of their work (firemen, soldiers and nurses) faced with life-or-death situations. In contrast to Tversky and Kahneman, he found his subjects to be competent decision-makers. What is the difference between the two situations?

In the one situation, Tversky and Kahneman gave people carefully formulated hypothetical problems in laboratory settings. It appears that their examples have been deliberately chosen: "researchers in the heuristics-and-biases program select problems in which reasoning by cognitive heuristics leads to violations of probability theory" (Chase *et al*, 1998, p 206). A study by Lopes (Klein, 1998) also implies that Tversky and Kahneman set people up for failure. An example she quotes is a question about whether the letter R appears more often in the first or third position of an English word. When the majority of people incorrectly respond that it appears more often in the first position, Tversky and Kahneman ascribe this behaviour to the availability heuristic. They argue that people find it easier to recall words that begin with an R than words with R in the third position. However, Lopes points out that of the twenty consonants, twelve are more common in the first position. Thus, R was carefully chosen to make people respond wrongly; there is in fact nothing wrong with the heuristic that the people used.

Gary Klein, on the other hand, adhered to the principles of naturalistic decision-making. He did not interfere with the decision-making process, but observed it in its natural settings and tried to understand it as it were. Klein highlights the irony that one of the primary biases referred to by for example Tversky and Kahneman is the confirmation bias (only seeking for information to confirm a hypothesis), and implies that Tversky and Kahneman have done the same.

Another difference between Tversky and Kahneman's studies and those of Gary Klein surfaces in Dietz and Stern (1995). They argue that people's analytical decision-making abilities are not very advanced when it comes to arithmetic and algebra, but that people are very sophisticated when it comes to pattern recognition and classification. Tversky and Kahneman test the arithmetic kind of skills, whereas Klein's Recognition Primed Decision model relies strongly on pattern recognition and classification.

Similar to the authors mentioned in the previous section, Dietz and Stern (1995) write about the social context of decision-making. According to them, it is more important for the human species' survival to be socially adept than to be good with arithmetic, or to be able to calculate the maximum Subjective Expected Utility (SEU). They discuss a number of socially related decision-making behaviours that are valid in a social context but can lead to outcomes different from the maximum SEU:

- Values, a concern for the welfare of others (altruism) or moral imperatives can dominate people's choices, but rational choice models do not adequately cater for the influence of values on decision-making.
- 'Hot' cognitions, such as overriding passions, drives or motivations are irrational or irrelevant within a SEU framework, yet they can be valid determinants of decision-making.
- Social influence and imitation are important aspects of learning. These are also recognised in theories on innovation diffusion (eg. Nicolau, 1999), persuasion, propaganda, mass media and social networks. The calculation of the SEU can take into account the sum of a number of individuals' sentiments, but not the multiplicative effect of social influence.

Many of the biases listed by eg. Kahneman *et al* (1982) and Hogarth (1980) show some irrationality that could be socially motivated. In the previous section, a number of such examples

have been shown in the work of Chase *et al* (1998) and Marsh (2002). Dietz and Stern hold that human decision-making is possibly far more sophisticated and evolved than the rational mode of decision-making, and not less sophisticated as the heuristics and biases literature would have it. In addition to Dietz and Stern, Lipshitz *et al* (2001) and Chase *et al* (op cit) support the idea that the social context should be taken into account when studying decision-making.

Boland (2001) criticises some of the assumptions on which the literature on heuristics and biases are based. One of them is the view of humans as information processors. Drawing on the work of Bruner, he calls for an alternative mode of cognition, namely the narrative mode. In the information-processing mode, people test hypotheses, deduce consequences and look for if-then rules. In the narrative mode, a story is constructed within which a series of events become plausible within the context of people's experience and culture. It is a form of sense making that allows a richer context for decision-making. Klein (1998) also emphasises the role of story telling in decision-making and particularly in sharing experience. He believes the context of the narrative is significant; it is something that is lacking in decision-making rules.

The heuristics and biases literature has indeed been attacked on many fronts. A comparison of some of the criticisms shows incoherent results. For example, Fox and Levav (2000) argue that when testing probability, Tversky and Kahneman should not ask, as they have done: "Which of the following is *more likely*?" Rather, if one asks: "Which of the following is *more probable*?" the bias is reduced. Chase *et al*, again, argue that information should not be presented as *probabilities*, but as *natural frequencies*, and people will make fewer mistakes. Lipshitz *et al* point to the importance of the semantic content of decision problems: they refer to studies where syntactically identical but semantically different decision problems produced opposite attitudes towards risk. The above arguments leaves the following question: do these arguments explain away the biases found by Tversky and Kahneman, or do they confirm Tversky and Kahneman's problem framing bias?

Another issue raised by both the contenders and the critics of the heuristics and biases literature is that of learning from mistakes. As Marsh (2002, p 55) states: "many demonstrations of 'failings' in human judgement are derived from experimental designs where subjects are given single choices with no feedback". Stanovich and West (1999) gave people questions on probability theory similar to those posed by Tversky and Kahneman. They then allowed people to revisit their choices after giving them additional information (but not the answers to the questions). The

majority of respondents whose choices were biased, corrected them after giving it a second thought. According to Stanovich and West, the results confirmed Slovic and Tversky's understanding/ acceptance principle: that the deeper the understanding of a normative principle, the greater the tendency to adhere to it. The fact that people respond more correctly after better understanding is used to defend the heuristics and biases literature from arguments that it does not judge people against appropriate normative models (Stanovich and West, *op cit*). Yet, Stanovich and West's findings can equally be used as criticism against Tversky and Kahneman's laboratory-based experiments.

Rationality implies internal consistency (Chase *et al*, *op cit*). Yet, it can be said that some of the biases or decision-making errors contradict each other. For example, the combined biases of inconsistency and conservatism can place a person in a no-win situation. If a person changes their mind about a situation after receiving new information halfway down a process and adapt their behaviour accordingly, their behaviour can be regarded as inconsistent. If they do not adapt their behaviour, they are conservative. In the same sense, the bias of 'misconceptions of chance' can, in certain situations, contradict the regression bias.

The findings of the heuristics and biases literature as well as the assumptions it is based on, namely that decision-making should be rational and humans are information processors, have been attacked from a large variety of angles. Although the criticism is not always coherent, a widespread unease can be sensed regarding the claim that human judgement is fundamentally flawed.

4.5. Conclusion

The focus of this chapter has been on the biases associated with the use of heuristics in decision-making or judgement. The origins of the heuristics and biases literature have been discussed. In the heuristics and biases literature (eg. Tversky and Kahneman, 1974; Hogarth, 1980) two major assumptions are made: that decision-making should be rational and that humans are information processors. For each of these assumptions, the biases associated with them have been discussed.

It was argued that the two assumptions mentioned above contain a major oversight – neither takes into account the social context of decision-making. Depending on one's departure point on whether perfect rationality is the proper norm to measure people against, the social context of

decision-making (and its associated heuristics) is used to either criticise or defend the argument of the heuristics and biases literature, namely that human judgement is flawed.

The research question that prompted this review of the literature, was:

- What is the effect of heuristics on the outcomes of decisions or judgements?

In the next chapter, the role of information systems in promoting or preventing biases in decision-making will be discussed.

5. BIASES AND INFORMATION SYSTEMS

5.1. Introduction

The research question to be addressed in this chapter, is:

- How does information system – supported decision-making get affected by biases?

The two major assumptions made in the heuristics and biases literature are those of humans as information processors and perfect rationality as the norm of behaviour. Against these assumptions, it was shown that a number of fundamental flaws or biases exist in human judgement or decision-making. The validity of the two assumptions were questioned in Chapters 3 and 4, where descriptive models of decision-making were contrasted to the rational model, alternative behavioural views to that of information processing were presented, and the social context of decision-making was discussed.

How do information systems fare when assisting humans with information processing and rational decision-making? Do they assist people to process information more consistently and rationally, thereby reducing decision-making biases? How are biases reduced in the first place? These issues will be explored by considering our understanding of the role of computer systems in human judgement. A number of case studies on the role of information systems in decision-making bias, as found in the literature, will also be discussed.

5.2. Debiasing

Debiasing methods are of interest when attempting to use information systems to reduce decision-making biases. Ways and means to reduce biases are discussed to different degrees in the heuristics and biases literature. Russo and Schoemaker (2002), for example, take people through the steps of decision-making and for each step discuss possible biases and how to overcome them. Critics of the heuristics and biases literature, such as Chase *et al* (1998) and Klein (1998) list similar debiasing strategies and argue that people make use of these strategies (for example feedback and learning) in real life but are not afforded opportunities to do the same in laboratory environments.

Roy and Lerch (1996) discuss three different debiasing strategies to overcome the bias of base-rate neglect, where prior probabilities are not taken into account as they should be. They suggest different presentations of the information, the training of subjects in information processing strategies and the replacement of human decision-makers with a model or system that follows the normative rules.

Fischhoff (1982) drew up a general framework of debiasing strategies, although his focus is on debiasing during psychological experiments in a laboratory environment. According to Fischhoff, biases occur either as a result of the decision-maker (faulty judge), the task environment (faulty task) or the mismatch between decision-maker and task. Faulty judges could benefit from the following four-level schedule, where each level suggests a stronger intervention than the previous: warnings about the possibility of bias; descriptions of the direction of bias; feedback about the subject's behaviour; and an extended programme of training. Faulty and unfair tasks could be debiased by: raising the stakes; clarification of instructions; discouragement of second-guessing; use of better response modes and the posing of fewer questions. Faulty and misunderstood tasks could benefit from the demonstration of: alternative goals; semantic disagreement; impossibility of task and overlooked distinctions. As can be seen, the majority of interventions suggested for faulty tasks are related to the posing of questions in a laboratory environment. More generally relevant is Fischhoff's suggestions for addressing the mismatch between judge and task, namely: making knowledge explicit; searching for discrepant information; elaborating on the problem; considering alternative situations; offering alternative problem formulations; and general education of the judge.

5.3. The Role of Information Systems in Human Judgement

The mechanistic view, computers and rationality

The notion that human judgement is flawed, that probability theory provides a norm against which to measure judgement, and that human judgement should be improved in order to arrive at the same answers as probability theory, is true to the mechanistic world view as described by Dahlbom and Mathiassen (1993). The mechanistic world view originates in the philosophy of modernity, which separates the 'true' external world from the way humans perceive it. Humans need to adapt their perceptions in order to better reflect the external world. The language of

mathematics (to which probability theory belongs) assists in describing the external world and provides us with the exact symbols we need to represent the world truthfully.

Computers are based on the use of this explicit and formal language of reasoning: "the mechanistic ideas of representation and formalisation are at the very heart of computing" (Dahlbom and Mathiassen, 1993, p 12). Perhaps, this solid foundation of rationality on which computers are built, can assist humans to be more precise in their thinking, more mathematically correct? Perhaps, it can assist them in adhering to the logic of probability theory and Bayesian statistics, so that their judgement will be more consistent and less flawed.

Information processing biases and information systems

Information systems are meant to assist humans with information processing. The possible role of information systems with regard to information processing biases will be investigated here.

According to Hogarth (1980), information processing bias can occur as a result of the decision-maker's mental *schema* or structure (as per structuration theory) as well as the *characteristics of the task* at hand. This is related to Fischhoff's (1982) classification of biases as caused by the judge, the task or the mismatch between the two. In what follows, Hogarth's explanation of information processing biases will be summarised and the role of information systems in this context will be investigated.

Hogarth states that a person's *schema* varies on three dimensions, namely *veridicality*, *stability* and *generality*.

Veridicality refers to the degree to which a person's schema represents reality. A mismatch between a person's perception and reality will cause a bias of veridicality. Examples of such biases are those related to availability, representativeness and illusory correlation, as discussed in Chapter 4.

Stability means that information should be collected and processed consistently. Inconsistent problem framing or aspects of the social context of decision-making could lead to non-consistent judgements.

Generality refers to how widely applicable the information processing rule is that is being applied. If the rule has been inferred from the observation of many cases, it is more general. Bias related to generality would occur if a rule were being used outside the scope of where it applies, or if false generalisations are made by a person with limited experience.

How would information systems fare on these dimensions, when processing information? In terms of the *veridicality* of information processing rules, it is believed that an information system will represent the veridicality of the people who have been involved with the system (designers, developers, users). Regarding *stability*, an information system is expected to behave more consistently than its human counterparts. However, a stable schema will not mean much if it is not veridical. As far as *generality* is concerned, a schema needs to be based on a large 'case base' and also be updated frequently with new experience and learning in order to remain as widely applicable and relevant as possible. The danger of the schema or rule/data base of an information system is that it will probably not be updated as frequently as human learning occurs. Also, the rule/data base of an information system will only contain a subset or a generalisation of the total set of experiences of the people who have designed it, and thus be less discerning. The conclusion from this is that in terms of generality, an information system is bound to fare worse than humans.

All in all, it appears as if an information system's *schema* will perform similarly or worse rather than better compared to humans when it comes to information processing biases.

How does the *task environment* affect biases of information processing? Hogarth mentions four characteristics of the task environment, of which only one actually refers to the task; the other three have to do with an individual's response to a particular task. The characteristics mentioned are: the *complexity* or difficulty of a task; *uncertainty* of the decision-maker as to how to deal with the task; the *ownership* or commitment of the decision-maker to the successful outcome of a situation; and the *stress* experienced by the decision-maker when performing a task. The potential for bias increases when a task is complex, when the decision-maker is uncertain how to execute it, when there is lack of ownership or when the decision-maker experiences high levels of stress related to the execution of the task

Where do information systems come into the picture of *task-related* biases in information processing? It is believed that an information system, when it is used, becomes an extension of a

person's *schema* for dealing with a decision task. The manner in which the system is used to deal with the *task*, is up to the user. First of all, an information system needs to be appropriate for assisting with the decision task. Furthermore, if the information system can assist the user to reduce or manage the *complexity* of the task, or if it can reduce the *uncertainty* of the decision-maker by providing information or suggesting means of solving the problem, it can assist in reducing potential information processing bias. If the information system is a key component of the decision-maker's dealing with the task, it is possible that an information system can influence the decision-maker's feeling of ownership as well as stress levels. Feelings of ownership can increase if the user has ownership of the system and is comfortable to use it, or vice versa. Similarly, a user's decision-making stress can increase or decrease depending on whether the use of the system increases or decreases their stress levels.

From the above, it can be said that the influence of an information system on *task-related biases* is indirect and depends on the way it is appropriated in the decision-making process. It depends on whether the system is an appropriate decision support tool, which is a function of its design. Also, the user's level of comfort when using the system (eg. does it increase or decrease task-related stress?) can impact decision-making bias. Thus, the role of an information system in reducing or increasing task-related biases of information processing is dependent on its design, use and perceived usefulness in relation to the task at hand.

Information systems and biases: a suggested framework

It has been shown how information systems could play a role with respect to biases as a result of the schema of the judge as well as biases related to the task environment.

An information system can be an extension of a decision-maker's mental schema when it assists with information processing that needs to be done by the decision-maker. In this regard, it could also have been influenced by the mental schemas of people involved in the design of the system.

The information system can be a part of the task environment, and can include interpretations of the task environment. An example of the latter will be when the organisation's financial manager uses a financial system as a representation of the real financial situation of the organisation, and makes decisions based on the reports of the financial system. As with the mental schema, the way

in which the task environment is represented by the system is influenced by other people, be it the designers or other users that have populated the system with information.

Where an information system is used in a decision-making situation, it plays a mediating role between the user (decision-maker) and the task at hand. It can play a role in judgement, in the way that the decision task is perceived as well as in the match between decision-maker and task, covering all three areas of possible bias in decision-making mentioned by Fischhoff (1982). In addition, the information system allows other peoples' (designers, developers, users) perceptions and biases to enter the decision-making situation. The following diagram illustrates the mentioned influences:

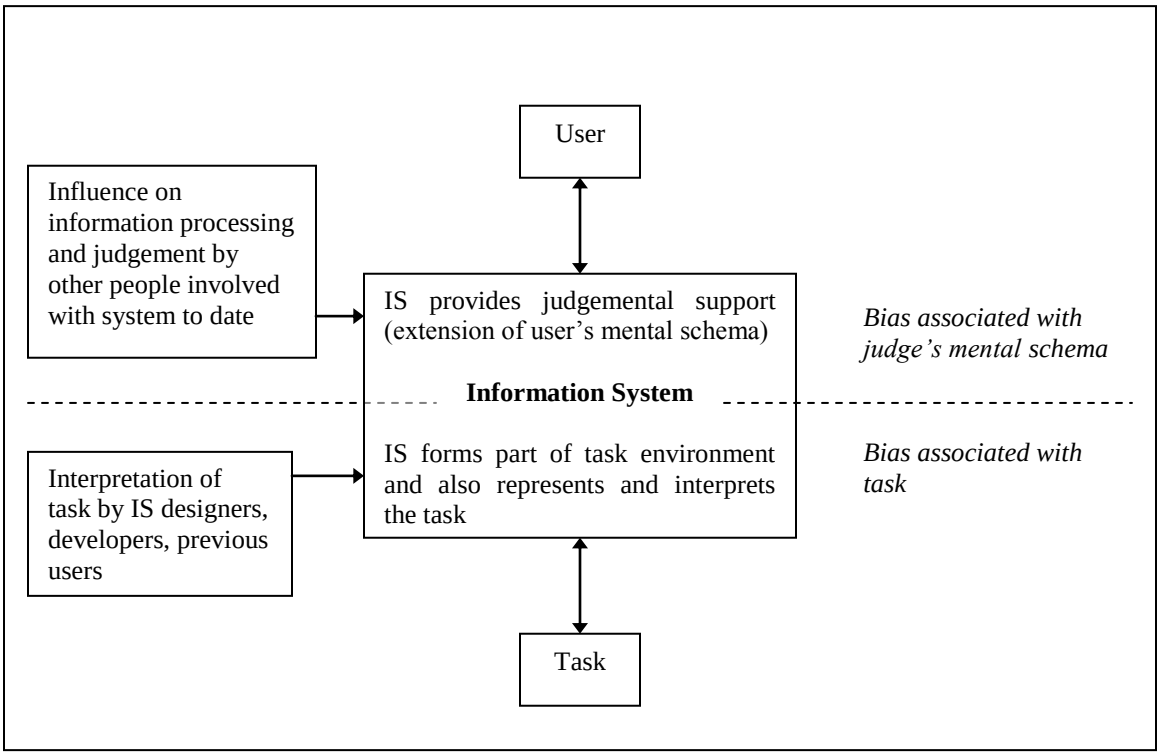

Figure 5.1: The role of the information system in judgement

Postscript: the romantic world view and social context of information systems

It has been argued in the previous paragraphs that the role of information systems in biases of information processing is heavily dependent on the design and use of the information system, as well as perceptions the user has of the system. The information system carries an interpretation of the task, which is again interpreted by the user. All along the way, biases can enter the situation and affect the user's decision-making.

It would appear as if the mechanistic world view that created the context for identifying and examining biases in the first place, falls short of recognising the environment in which information systems are used and/or where decision-making occurs.

The romantic world view that took shape in the nineteenth century originated in protest against the mechanistic view (Dahlbom and Mathiassen, 1993). Romanticists strived to understand the world within rather than the world without. Amongst others, they were influenced by Descartes, who distinguished between the way the world is and the way people perceive it. People's understanding of the world is an interpretation, which is not less valid than the world 'out there'. People's representation of the world is a construction. The romantic view would recognise the notion that an information system is a construction; that it contains an interpretation of the task environment. It would also recognise the unique interpretation of the contents of the system by the user.

Another characteristic of the romantic view is its acceptance of contrasts and contradictions (Dahlbom and Mathiassen, *op cit*). It would be sympathetic to the idea that the social context of decision-making determines what is 'logical', rather than a static measure such as probability theory. What is 'logical' in one context might not be logical in another.

Du Plooy (2003) builds on the romantic world view. He argues that information systems are social systems. The social context in which information systems operate goes beyond the fact that they are used and interpreted by humans. The technology itself, the hardware as well as the software, carries the stamp of the values and assumptions of the humans that developed them. Du Plooy describes the human environment in which information systems are used. He lists six aspects of this environment, namely the social contexts of individuals, groups, and the organisation, as well as the social contexts of tasks, the technology and the larger environment.

It appears that the romantic world view as well as Du Plooy's social context provides a more appropriate context than the mechanist view for investigating the role of information systems in human judgement, and thus the environment in which biases relating to the use of information systems can be formed, enlarged or reduced. The rational model of decision-making, under which biases were identified, originates in the mechanist world view. Ironically, within the romantic world view, the notion of biases itself can be drawn into question.

5.4. Studies of Biases and Information Systems

From the information systems literature, seven studies were found where the occurrence of biases was investigated in an information systems context. These studies are summarised below. In four instances, experiments were set up to test the robustness of a specific bias in a tailor-made information system setting. In the other three instances, the point of departure was a particular kind of information system (an expert system, an executive information system and a battle management and radar tracking system) and the biases expected or found to occur with the use of such a system was discussed.

5.4.1. Countering the anchoring and adjustment bias with a decision support system

Anchoring and adjustment refers to the situation where people are requested to estimate a value, and they bias their estimates towards a suggested initial value. George *et al* (2000) built a decision support system to test the anchoring and adjustment effect as well as the user's response to warnings (Fischhoff's debiasing of faulty judges, level 1). Subjects had to estimate the value of a house, based on information normally considered by a realtor. Some were provided with a listing price and others not. The listing price provided was either significantly higher or lower than suggested by the other real estate information. When the subject's appraisal was too close to the anchoring value (the suggested listing price), a warning was issued by the system. George *et al* found that the judgement of people who received anchoring values was significantly influenced higher or lower, depending on the value of the anchor. Subjects who received warnings to this effect adjusted their appraisals only slightly, and their responses still reflected the anchoring bias. George *et al* show that the mere use of a decision support system, even with a moderate attempt at debiasing, does not lead to more rational decision-making.

5.4.2. Overcoming the bias of base-rate neglect with a different mental representation

The study of Roy and Lerch (1996) focuses on the bias of base-rate neglect, encountered when people do not take prior information into account when calculating a probability. Roy and Lerch show how different presentations of a problem, including a graph, can assist in reducing base-rate neglect. They build on the thesis of Todd and Benbasat (1991, 1992), who argued that in order to provide better decision support systems, one must first identify the cognitive processes that are in need of support. The decision aid then needs to reduce the cognitive effort required to solve the problem, or overcome the limitations of human information processing while supporting the underlying cognitive processes.

Roy and Lerch argue that the mental representation invoked by a problem description is a cause of bias. Often, a causality of events is read into a problem, and subjects use a narrative mode to 'explain' the story to themselves (as also observed by Boland (2001)). People end up rephrasing the question to fit their story and end up answering the wrong question. In the case where base rates occur in the problem description, people end up calculating the wrong probability (eg. the probability that a patient shows particular symptoms, given that they have a certain disease, instead of the probability that a patient has a certain disease, given that the particular symptoms are displayed). Re-presenting the verbal question into formats of tables and graphs removes the causality factor and is shown to reduce biases. The graph they use is a sketch of 100 blocks, in a 10 by 10 square. It displays probabilities in terms of natural frequencies (out of 100) rather than using Bayes' rule with its unintuitive formulation of conditional probabilities. (Chase *et al* (1998) also report success with the use of natural frequencies in overcoming base-rate biases). Of all graphical and tabular alternatives considered, the graph *without* an accompanying verbal version of the question result in the lowest occurrence of bias. Roy and Lerch argue that when a verbal description is present, people may regard graphical displays as redundant, ignore them and not receive their benefit.

Roy and Lerch's experiment does not include information technology as such, but they propose that their conclusions have implications for the design of better decision aids in general, and in particular the way information is presented in an information system.

5.4.3. Investigating the role of multimedia in changing first impression bias

First impression bias refers to the notion that people are strongly influenced by the first piece of information they receive and tend to evaluate subsequent information based on the first impression. Lim, Benbasat and Ward (2000) conducted an experiment to investigate whether multimedia could reduce first impression bias. After receiving negatively biased information about a manager, two groups were requested to do a performance appraisal of the manager. Both groups were given the results of an interview with the manager: one group received a text copy of the interview and the other a video clip (both of these in an electronic, hypermedia format). The video-based performance appraisal resulted in a much more positive review of the manager compared to the text-based appraisal. The first impression bias was reduced in the group using the video clip, but not eliminated.

Lim *et al* argue that the video-based or multimedia information is communicated through a richer language, providing information in more vivid ways and thus making it more salient than the original piece of text with biased information. Also, the complementary cues of body language accompanying the manager's words can make his verbal message appear more authentic. Lim *et al* see a benefit for organisations to use more multimedia-based systems for decision-making. They do admit that the complementary cues or body language could be manipulated or could possibly do injustice to a 'good' message delivered by someone with poor presentation skills.

5.4.4. The influence of automation bias on decision-making

Automation bias refers to the finding that people relax their vigilance in the presence of automated decision support, such as an autopilot. Skitka *et al* (1999) firstly ascribe this behaviour to social loafing and diffusion of responsibility (the automated aid is treated as a 'team member' with whom responsibility is shared), and secondly, to submission to authority, assuming that people treat the decision aid as a source of authority that is possibly also regarded to be smarter than themselves.

Skitka *et al* set up an experiment with a flight simulator. All subjects (experimental and control groups) had to monitor a number of gauges and respond to them. In addition, the experimental group had a decision aid that automatically monitored the gauges and displayed messages with

instructions to respond to the values of the gauges. However, the decision aid was programmed with a few errors, about which the experimental group was told.

Skitka *et al* found that under the correct functioning of the decision aid, the experimental group outperformed the control group in responding to the flight simulator. However, under faulty conditions of the decision aid, the group with the decision aid performed substantially worse, despite the fact that they had access to the readings of the same gauges as the control group. The experimental group made errors of omission (not responding to the gauges when they should have, and the decision aid failed to prompt them) as well as errors of commission (executing faulty instructions given by the decision aid). The automation bias found in the experiment did not seem to improve with training.

According to Skitka *et al*, the dilemma faced by environments dependent on automated decision aids (such as aeroplanes, intensive care units and nuclear plants) is that they are dependent on these systems but that they cannot afford for system-related errors to go unnoticed by human operators, because of the high risks involved.

5.4.5. The Vincennes disaster, expectancy bias and the role of the Aegis system

During the Iran-Iraq Gulf War in 1988, a United States military ship, the USS Vincennes, shot down an Iranian passenger airliner with 290 civilians, mistaking it for a fighter plane on the attack. Several investigations were made into this incident, at least four of which official. The discussion that follows is informed by Fisher and Kingma (2001) and Klein (1998).

Who was to blame for this high profile embarrassment? The Vincennes was one of the US Navy's most technically advanced ships, equipped with Aegis, the world's finest battle management system, and acclaimed for its radar tracking proficiency (Fisher and Kingma, p 113). Investigations showed that the system by itself did not make any errors in identification. Rather, according to Fisher and Kingma, errors manifested themselves in the *human-computer interface*. The Aegis system's display was correct but not easy to read: altitude could not easily be inferred from the radar screen. In addition, the system recycled the radar track number of the Iranian passenger plane and re-assigned the number to a far-away fighter plane. The Vincennes captain asked for the status of the approaching plane, referencing it by the previous track number. He was

told it was a fighter plane, descending and increasing in speed. The captain gave the instruction to fire.

Klein refers to the fact that several of the incident reports ascribe the error to an expectancy bias, incurred by the crew: they expected the airliner to be a fighter plane and then mistook it for one, despite conflicting information from the radar display. According to Klein, this is an unfair allegation. He ascribes the incident to a misunderstanding between competent decision-makers and the Aegis system (similar to Fisher and Kingma), as well as among the crew themselves.

This case shows that even where no local or individual errors have been made, an error could still occur because a series of misunderstandings: the captain not realising the radar track number had been re-assigned, the crew interpreting his words literally rather than responding to the captain's intended question, and the difficulty to infer from the radar screen that the approaching plane's altitude was different to that of the far-away fighter plane. Two of the misunderstandings were a direct result of the information system design: the system assisted with the introduction of an expectancy (whether viewed as a bias or misunderstanding) and was not helpful in clarifying it.

5.4.6. Reinforcement of biases through the use of an executive information system

Rai *et al* (1994) argue that a number of features of a modern executive information system (EIS) have the ability to reinforce decision-making biases. These features include:

- the ability to select and monitor critical success factors;
- exception reporting;
- graphs using vivid colours and visualisation techniques;
- drill-down capacity;
- large volumes of data;
- real-time data; and
- electronic communication features (e-mail).

The *availability bias* is incurred when decisions are based on information that is easily retrievable, rather than a more complete set of relevant information. Salient graphics, exception

reporting and custom-made interfaces to monitor critical success factors can enhance the availability bias.

Regression effects are encountered when undue attention is given to outliers. Exception reporting and associated graphic displays (flashing outliers in red) can cause managers to focus their attention on outliers (which according to statistical theory will regress back to the mean) at the cost of following the trends and performance of the bulk of the entities to be monitored.

Overconfidence occurs when a person displays more confidence than is warranted by the supporting information. The large volumes of data in an EIS, of which much might be redundant, leads to increased confidence. Also, with a customised user interface, displaying only information of interest to the manager, the illusion can be created that the manager has an accurate overview of the situation. Drill-down functionality, providing access to minute detail, can create the illusion that a manager is in control of all the levels in the organisation.

To reduce decision-making biases, Rai *et al* recommend the incorporation of Silver's decisional guidance principles, referring to suggestive and informative guidance, in the design and use of an EIS. They also advocate the more frequent use of analytical or statistical tools, which they believe do not feature strongly enough in executive information systems.

5.4.7. Biases in the development and use of an expert system

Shore (1996) discusses a case study of the development of an expert system, and indicates how biases could enter the system and the decision-making processes influenced by it. He discusses the influence of subject-matter experts, knowledge engineers, end users as well as the people validating and maintaining the system. While Shore's use of the word 'bias' is fairly loose and general, it seems to be founded on Tversky and Kahneman's (1981) work describing the problem framing bias.

The expert system described by Shore was developed to assist underwriters at a large insurance company. Shore argues that when one moves higher up in the decision-making hierarchy, decision-making becomes less structured and more open to interpretation, thus more prone to biases. The system he refers to has been developed to support such decision-making, so the occurrence of biases could be expected. Shore discusses two kinds of errors that can occur when

using an expert system. The first type is when the problem framing or user's cognitive style differs from that represented in the expert system. The result is that the user does not trust the system and does not make use of it when they should. The second type of error is when the system becomes outdated or inaccurate and the user is not aware of it. In Hogarth's terminology, a bias of generality occurs. In the second case, the user trusts the system when it shouldn't be trusted.

In the case studied, the first type of problem was addressed by the increased joint effort of knowledge engineers and subject-matter experts in designing the system. They analysed many cases that the system might have to deal with, looking at various ways to solve the problem. This was found to reduce the 'problem framing' bias. The second type of error calls for additional effort to update and maintain the expert system. The possibility of a 'learning system' to deal with the second type of error was investigated. It was found to be theoretically attractive but practically not feasible.

5.4.8. A summary of the biases studied in an information systems context

Tversky and Kahneman's classification of heuristics and biases (Kahneman *et al* (1982), Tversky and Kahneman (1981)), as discussed in paragraph 4.2.1, will be used when summarising the biases discussed in an information systems context as found in the literature. The automation bias does not appear in Tversky and Kahneman's classification – biases as a result of automation have not been studied in the psychology literature considered.

Representativeness

- Insensitivity to prior probability or base rates: Roy and Lerch (1996).
- Misconceptions about regression towards the mean: Rai *et al* (1994).
- Overconfidence: Rai *et al* (1994).

Availability

- First impression bias: Lim *et al* (2000).
- Familiarity and salience: Rai *et al* (1994).

- Expectancy bias: Klein (1998) (contested).

Anchoring and Adjustment

- George *et al* (2000).

Problem framing

- Shore (1996).

Automation bias

- Skitka *et al* (1999).

Although the literature that was found on the occurrence of biases in information systems settings is rather sparse, one can see from the above list that the literature has a broad coverage in terms of the types of biases investigated.

5.4.9. Conclusions from the experiments and case studies

Are there any general lessons to be learnt from the mentioned literature, in terms of how biases could be reinforced or reduced through information systems?

Introducing and reinforcing biases through information systems

Skitka *et al*'s work on the automation bias shows how biases that did not otherwise exist can be introduced as a result of using information systems. George *et al*'s study concludes that the mere use of information technology, even with moderate attempts at debiasing, does not reduce biases as such. Rai *et al* shows how the design of an information system can allow users to 'hang themselves', or to choose to use information in a biased fashion. The story of the Vincennes shows the unintended consequences of a seemingly logical information systems design: the system led to the introduction of a false expectancy (viewed by some as a bias) and did not help to clarify it. Thus, an information system can assist in introducing as well as reinforcing biases.

Shore as well as Rai *et al* studied decision-making on a strategic level, where decision-making is less structured. It is typically such decision-making that is supported by an MIS, an EIS (as discussed by Rai *et al*), and also in Shore's case, an expert system. Shore argues that unstructured decision-making is more open to bias, and this is confirmed in Rai *et al*'s study. Thus, when information systems are used at strategic level, the chances of biases being introduced or reinforced are even greater.

George *et al* shows that information systems by themselves do not reduce biases. At least in this case, the argument that information technology, as an artefact of the mechanistic world view, can assist people to be more rational, is refuted.

Reducing biases through information systems

Roy and Lerch use a different mental representation of information, to better assist with human information processing. Lim *et al* reduce biases by the use of multimedia instead of text. Multimedia communicates with the user in a different, richer and more humane language. Shore shows how problem framing can be overcome by making explicit the cognitive styles of subject-matter experts on whose reasoning the system is built, and by examining alternative ways to present the problem solving process in an expert system. The mentioned three studies all indicate an awareness of the way information is communicated and represented. It appears as if there were deliberate attempts to assist the human cognitive processes in need of support, as argued by Todd and Benbasat (1991, 1992).

Were any of Fischhoff's (1982) debiasing strategies used? The debiasing schedule he suggests for faulty judges, namely warnings, feedback and training were showed to be ineffective in the studies by George *et al* and Skitka *et al*. The suggestions for debiasing of faulty tasks were not referred to. However, some of his suggestions for addressing the mismatch between judge and task, such as considering alternative situations (Shore) and offering alternative problem formulations (Roy and Lerch), showed positive results. It needs to be pointed out that Roy and Lerch (1996, p 235) state that they went further than merely to reformulate or reframe (thus to reword) problems; they used another language of logic altogether.

Conscious vs. unconscious biases

The cases that were studied focus on unconscious side effects of interpreting and processing information. The heuristics and biases literature also views biases as an unconscious side effect of judgement or decision-making. Yet, systems can be manipulated consciously to lead to decision-making biases (eg. the design of graphs, reports or any other displays) of which users might be unconscious. Hogarth's (1980) work alludes to the conscious manipulation of information, and Tversky and Kahneman's (1981) discussion of problem framing implicates people that sell insurance and people in the gambling industry. One can conclude that an awareness is required not only of unconscious biases, but also of consciously manipulated information, and ways to uncover and reduce such biases.

The case studies and the framework for the role of information systems in judgemental biases

The framework for investigating the role of information systems in judgemental biases (Fig 5.1) is confirmed by Shore's study that shows how all the different people interacting with an expert system can introduce biases. It is also helpful in explaining how, in the case of the Vincennes shootdown, the system was inadequate in the way that it represented the task environment.

The framework shows that an information system becomes an extension of a user's (decision-maker's) mental schema. The literature that reports on successful debiasing emphasises the significance of the way in which information systems are used to assist with the user's cognitive information processing. Again, one is reminded of the claim by Todd and Benbasat that information systems should support the cognitive processes of the decision-maker.

5.5. Conclusion

Information systems and in particular decision support systems have been developed to supplement human information processing. In this chapter, the role of information systems in decision-making has been investigated, with the focus on the role of information systems in introducing, reinforcing or reducing biases of decision-making. In doing so, the literature pertaining to the main research question of the study has been addressed, namely:

- How does information system – supported decision-making get affected by biases?

It has been found that information systems have the ability to introduce new biases (such as automation bias) and to reinforce biases. It has also been shown in one case that the mere use of information technology does not reduce biases. On the other hand, there were reports of the successful use of information systems to reduce biases. In the latter respect, the systems did more than just automating information; information was communicated in a different manner, whether by using a different problem formulation, mental representation or communication media.

The next chapter will report on a practical case study of biases and information systems, after which the study will be concluded.

6. CASE STUDY: EXCEPTION REPORTING

6.1. Introduction

This chapter will report on a case study where exception reports have been developed as part of a project to provide management information for a government department. Information systems and data analysis software play an important role in defining the exceptions and developing the reports.

In the previous chapter, a summary of the work of Rai *et al* (1994) has been given. Rai *et al* argued that exception reports in an EIS setting could lead to undue attention being given to outliers. This bias is also referred to as 'regression effects', and does not acknowledge the law of regression towards the mean. As mentioned in Chapter 4, outliers are due to regress closer to the mean over a period of time (Hogarth, 1980).

Is the exception reporting done for the government department open to this bias? Can the bias be observed in the way the reports are used by decision-makers? What other biases can be found in the decision-making processes supported by the Ops Room team? What is the role of information systems in enhancing or reducing these biases? These issues will be addressed here.

The discussion will commence with some background on outliers and a review of literature on the bias related to regression effects. Subsequently, background on the project will be given, after which the information collecting process will be discussed. The occurrence of outliers in the case study will be investigated, and the questions mentioned above will be addressed.

6.2. Outliers

What is the Ops Room project team's view on outliers? Why do datasets contain outliers, and how should they be dealt with? According to one of the statisticians interviewed (a team member of the project to be discussed, see paragraph 6.4.2), there are several ways of identifying an outlier:

- A visual examination of the data, when data points are plotted on a graph. Data points falling outside a visual cluster defined by the majority of the data can be regarded as outliers.
- Using statistical measures, such as the 5^{th} and 95^{th} percentiles (thus classifying the bottom 5% and the top 5% of data points as outliers) or the standard deviation (eg. regarding values that are three standard deviations away from the mean as outliers).
- Setting particular targets, such as sales targets, and regarding individuals or groups deviating significantly from the target as exceptions or outliers.

The statistician was uncomfortable with simply regarding highest and lowest points in the data as 'outliers'. According to him, data following a normal distribution (or another distribution indicated by the behaviour of the population) will inevitably have points at the edges of the curve, and that is acceptable. He indicated that one should be careful in labelling a data point as an outlier.

Two possible reasons were given for explaining the phenomenon of outliers. Firstly, structural differences in the entities observed can result in behaviour different to that of the majority of the group (that is, a group of outliers may mean that these points belong to a separate group). In this case, entities should be regrouped according to their different attributes and analysed separately. Another possible reason for observing outliers is data errors. Order-of-magnitude errors or data capturing errors are not uncommon. Hogarth (1980, p 33) states that unreliable data often display greater variability than a similar dataset free of errors.

Analysing vs. managing outliers

The statistician observed that the way an analyst deals with outliers will and should be different from a manager's response. An analyst may identify outliers and subsequently remove them from a dataset in order to analyse the characteristics of the dataset, such as fitting a curve to it. In such a case, outliers are regarded as 'noise.' Outliers may also be simply explained away, such as by arguing that a normal distribution would always include 'outliers'.

A manager, however, needs to deal with the outliers. A teacher cannot explain away the marks of the pupil who failed the test by stating that the class' results follow a normal distribution (which indicates that somebody will inevitably fail); he or she has a moral obligation to ensure that the

pupil's performance is addressed. Similarly, a manager cannot afford to say that underperformance is allowable, since it might affect the financial well-being and future of an entire organisation.

The bias of regression effects refers to situations where a manager over-reacts to outliers in an attempt to diligently manage them. This bias will be discussed in more detail in the next section.

6.3. Review of literature: regression effects

The bias associated with regression effects is discussed by Tversky and Kahneman (1974), Kahneman and Tversky (1973), and Hogarth (1980). Rai *et al* (1994) refer to this bias, which they call "undue attention to outliers" in the context of exception reporting in EIS. The bias of regression effects has been mentioned in Chapter 4, but will now be discussed in more detail.

Tversky and Kahneman (1974) discuss "misconceptions of regression" as one of the biases associated with representativeness. How representative of a dataset is a particular outlier?

Hogarth (1980) displays a graph with 22 data points, indicating the performance of an individual over time. The points are all either slightly above or slightly below the mean, which is also displayed. The last data point is a bit further below the mean than the other points. Hogarth argues that the most likely prediction for the next observed performance of the individual will be the average, rather than the last data point. "Consequently, exceptionally good performances are liable to be followed by lower (ie. average) ones, and exceptionally bad performances will be bettered" (p 20). According to Hogarth, we all have our good and bad days, which can be accentuated by specific circumstances. However, there is a danger that a manager who displays the totally reasonable behaviour of managing by exception will over-react to this 'explainable' behaviour of the individual. Tversky and Kahneman (1974) discuss a similar example and then points to the psychological implications of a manager's reaction. In a case study where flight training was investigated, experienced instructors noted that praise for an exceptionally good landing is typically followed by a poorer landing the next time round, while harsh criticism after a rough landing usually leads to an improvement on the next landing. They concluded that "verbal rewards are detrimental to learning, while verbal punishments are beneficial, contrary to accepted psychological doctrine. …. Thus, the failure to understand the effect of regression leads one to overestimate the effectiveness of punishment and to underestimate the effectiveness of reward.

…. Consequently, … by chance alone, one is most often rewarded for punishing others and most often punished for rewarding them" (p 10).

Rai *et al*'s (1994) work that indicates how the characteristics of an EIS lend itself to biases, has been discussed in Chapter 5. It has been argued that exception reporting and associated graphic displays (eg. flashing outliers in red) can cause managers to give undue attention on outliers. What is true for exception reporting in an EIS, should hold for exception reporting in general. In the following case study, exception reporting forms an important component of the management information that is provided as part of the project to be discussed.

6.4. Exception reporting at a government department

6.4.1. Background and introduction

A government department has recently initiated an Operations Room (Ops Room) function that aims to use existing departmental information as well as information processing and -analysis technology in order to provide better intelligence for decision-makers in the department. The concept is similar to that of a War Room. In particular, the Ops Room has to report on the performance of close to 500 government offices around the country, and where a total of almost 10 000 people are involved. Although information analysis and performance measurement has been done before, the Ops Room provides for the first time an integrated and dedicated management information function.

Part of the Ops Room's brief is to do exception reporting. The Ops Room has to give regular feedback as to the best and worst performing offices in the country. These offices perform an important function for the citizens of the country, and are essential for the legitimacy of the government. As such, the performance of the offices is not only of concern to the government department, but to various other stakeholders as well. In the discussion that follows, the way that the exception reporting is performed will be described. Later on, this process will be investigated for evidence of the regression bias as well as other potential biases.

6.4.2. Information collection

The researcher has been a member of the Ops Room project team since its initiation, and has been previously involved with a number of other projects for the same government department. Thus, it is believed that the functions of the department as well as the context of the Ops Room project are understood. In order to investigate the exception reporting function, the team members immediately involved with the design and drafting of exception reports have been consulted. Sources of information were discussions at project team meetings as well as a number of informal discussions with team members, including colleagues from the consulting organisation (mostly statisticians by profession), people from the government department as well as other roleplayers on the project team. Among the project team members are two people who used to be managers and decision-makers in the government offices under investigation. In order to investigate decision-making biases, discussions were held with some of the government decision-makers making use of the reports from the Ops Room.

6.4.3. The exception reports

To date, the performance of the best and worst offices have been measured in three different ways.

Historically, the only information available for performance measurement were inspection reports. Offices would be visited more or less once every two years by officials from the Inspectorate function in the department. An in-depth inspection or audit would be performed of all the functions of the office. This included an audit of the administrative (housekeeping) processes of the office, as well as the office's efficiency in performing their core functions. The inspectorate's physical presence in the offices allowed them to obtain 'rich' information, such as to form impressions of the atmosphere in the office, perform 'drill-down' investigations where they suspected non-compliance, and have interactive discussions with a variety of people working at the office. Drawbacks of the inspection reports were that they were extremely labour-intensive to complete and that the information for a large number of offices was dated by the time the performance of the various offices were compared. Also, the inspection reports did not contain quantified performance ratings; the best and worst offices were selected based on impressions confirmed by the reports.

More or less at the start of the Ops Room project, office productivity information was made available by a third party role-player involved in the offices. The role-player collected information nationally and on a monthly basis for its own purposes. The productivity information was collected for only a selected, yet significant function of the offices (where the role-player was involved). It provided a quantitative measure of performance, and indicated how well the selected function met its own targets. The Ops Room subsequently started using this dataset as indicator of the best and worst offices. The dataset included a fair amount of variables for the statisticians on the Ops Room project to play around with, and a fairly robust model has been developed for calculating the best and worst performing offices. Advantages of this method of exception reporting include:

- the boundaries or cut-off values for determining the best and worst offices are data-driven. Rather than using absolute targets, percentiles classifying offices according to three different variables jointly determine the 'exceptions'.
- Since information is available on a monthly basis, the average for a number of months can be used to calculate exceptions. In this way, regression effects or undue attention to outliers (as indicated by single data points) are minimised.

Disadvantages of using this performance measure are the following:

- Only a selected function of the office is measured, and only for productivity-related performance. Adherence to processes, as were included in the inspection reports, is not addressed at all. Performance thus relates to only a certain part of the entire function of the offices.
- The relationship between the role-player or body collecting this set of data and the Ops Room is strained. When queries arise around the quality of the data or the way it should be interpreted, the role-player is often not willing to discuss the dataset.
- Inconsistencies in the dataset indicate potential problems with data quality. The strained relationship mentioned above makes it difficult to resolve this issue.

A third set of information to measure office performance, including the best and worst offices, became available as a result of an intervention in a number of the offices. To date, 44 offices countrywide have implemented a detailed performance tracking information system. The same

kind of information is collected through this intervention as is contained in the database of the third party role-player discussed above, and for the same selected function of the offices. Thus, the representativeness of the data of the entire office function is similarly limited. However, the data quality of the information in this system is believed to be significantly better than the third party dataset.

6.4.4. Comparison of results of the three datasets

Although the list of best and worst offices produced by the analysis of the three sets of information are mostly consistent, it has happened that an office appears as 'worst' on, for example, the list derived from the inspection reports, and simultaneously appears as a 'best' office on the list derived from the third party dataset. This has been a major concern for the team members of the Ops Room project. To date, it has been possible for people with detailed knowledge of the concerned offices to explain these differences. It is possible for an office to perform well on the variables measured in one dataset, and to perform badly on other variables measured in a second dataset. These examples show the danger of accepting any of the produced lists of 'best' and 'worst' offices as a true reflection of reality, and of the danger of acting on analysed data without taking cognisance of the context of the offices and the limitations of the data.

6.4.5. Managing by exception in the government department

The department appears to be very careful in using the exception reports as a mechanism for managing the offices. On the one hand, there is a realisation that the reports represent only a part of the picture. On the other hand, the environment in the offices and between this department and other roleplayers involved in the offices is so politically loaded that the department is careful not to worsen relationships by implicating people for not performing. As a result, the exception reports have to date not been made part of a feedback loop between the department and the offices.

6.4.6. Conclusion

In the Ops Room environment, concerns for the possible error due to regression effects or undue attention to outliers appears to be overshadowed by other challenges. These include the lack of a representative data set, potential data errors and the politicised decision-making environment. The inclusion of a number of qualified and experienced statisticians on the Ops Room project team appears to have led to a situation where at least the quality of statistical analysis and interpretation as appears in the Ops Room reports, are scientifically valid.

Ops Room's doctrine is to do the best they can with data they have. The situation with data collection and quality checking is gradually improving. Concurrent to the 'make do with what we have' exercise has been a more 'future-driven' initiative to define a scorecard of performance indicators that should ideally be measured in the offices. Such a scorecard needs to be defined in an all-inclusive process to ensure ownership. More or less a year after launching the scorecard initiative, attempts are still being made to get all the roleplayers on board to start defining the scorecard. Decision-makers from the department are starting to realise that they will need to follow a more pro-active approach in developing the scorecard, possibly compromising at least partially on the political inclusiveness of the process.

6.5. Applying the theory on heuristics and biases

Possible or perceived biases on the Ops Room project will be investigated with reference to the different theoretical contributions on heuristics and biases in paragraph 4.2.

6.5.1. Tversky and Kahneman's biases

Avenues for regression bias

As mentioned in paragraph 6.3, Tversky and Kahneman associate the regression bias with heuristics of representativeness. It has been argued that the regression bias is anticipated with exception reporting. How does the Ops Room fare in terms of the development and use of the exception reports?

Regression bias would have been possible with the inspection reports, since they were infrequent and could easily be interpreted in isolation. However, the same bias is unlikely to occur with the current set of performance data (the third party dataset) because of the way the dataset is analysed. The exception reports contain averaged performance values, so that the 'best' and 'worst' offices are those that consistently perform 'better' or 'worse', rather than being shooting stars or offices that just experienced a bad month. Also, the boundaries for defining outliers are derived from the data itself rather than being externally imposed, and as such are more realistic.

Representativeness in general

Although Tversky and Kahneman's work focuses on specific biases associated with the incorrect use or interpretation of numerical data, the liberty will be taken to use their term 'representativeness' here in a more general sense.

A representativeness bias can occur if the exception reports are regarded as more representative of the underlying situation than they really are. The potential for this bias to occur, is high. Firstly, the exception reports are based on a subset of the total function of the offices. Secondly, the dataset that is primarily used to report best and worst offices is that of the third party role-player, which is problematic to verify and suspected to contain an unacceptable amount of errors. Perceptions of error are shielded by the fact that an organised entity exists to deal with information processing and analysis, and that this entity is perceived to be well organised and competent. Also, the large volume of data analysed can lead to the perception that the analyses will be more accurate. It has been observed that the decision-makers have a high regard for the work of the Ops Room team, and that they are keen to prove to stakeholders that they know what is going on in the department's offices, based on the reports generated by the Ops Room. Although this behaviour is justified (the Ops Room team is indeed doing good work, and in a politicised environment the decision-makers need data to justify their opinions) it is indicative of a bias of representativeness.

Availability in general

It is perceived that the exception reports of the Ops Room are biased towards data that is readily available. Since the available data is unrepresentative, the bias of availability is in this case just another way to view the bias of representativeness.

6.5.2. Hogarth's biases of information processing

Hogarth's biases of information processing were listed in Table 4.1. Of the listed biases, the following are observed in the Ops Room's information processing function.

Hogarth mentions a bias of *data presentation* with respect to both information acquisition and information processing. The manner in which the office performance dataset is presented by the third party role-player that has captured it, has an influence on the way it is interpreted by the statisticians of the Ops Room project. Secondly, the exception reports generated by the Ops Room convey a message of thoroughness and analytical accuracy. Thus, the way in which the exception reports are presented can bias decision-makers to believe the results are more correct than they really are.

Consistency of information sources is a bias of information processing that reflects the fact that people can have a false sense of security when they have large amounts of information, much of which is redundant. This is true for the Ops Room, as mentioned above.

Illusion of control is a bias associated with information output. According to Hogarth, planning or forecasting can induce a feeling of control over the future. Strictly speaking, the Ops Room have not done forecasting to date. Yet, it is believed that the analyses and exception reports provided to the decision-makers at the government department can lead to the illusion that they have better control over the offices that they need to manage.

6.5.3. Biases associated with the decision-making process

The biases discussed by Russo and Schoemaker in a decision-making context (see paragraph 4.2.3), and that were found to be relevant to the Ops Room project, largely overlaps with what has already been presented above. For example:

Problem framing: difference in problem frames can be observed between the data capturers, the analysts and the decision-makers respectively.

Intelligence gathering: as mentioned, intelligence is based on the most readily available data.

Choice: as mentioned, false confidence is found in the sophisticated methods of data analysis.

6.5.4. The social context of decision-making

A number of heuristics and biases occurring in a social context have been discussed in paragraph 4.3. Of these, the *one-of-us-effect* was observed in the decision-making environment of the Ops Room. The actions of roleplayers external to the project were found to be interpreted according to whether they were 'one-of-us', or whether they were sympathetic to and supportive of the work of the Ops Room.

6.6. The role of information technology with respect to biases

The information systems used on the Ops Room project (spreadsheets and a statistical package) allows for sophisticated statistical data analysis. A database has been developed earlier in the project, but has been found to be less flexible for manipulating data and understanding the relationships between variables than the spreadsheets have been. The information technology that is used enables the deriving of data-driven performance boundaries and the averaging of monthly data over a longer period, thus decreasing the possibility of regression bias or unrealistic performance targets, and subsequently decreasing the possibility of 'skew' attention to outliers.

A disadvantage of information technology in the Ops Room context is that it only works with quantified data or numbers to represent effectiveness of the offices. Qualitative impressions as found in inspection reports are lost. Also, some of the data analysts have never visited any of the government offices themselves. Consequently, a large distance potentially exists between the dataset and the real-life situations that it represents.

Information technology in the Ops Room has assisted in the successful countering of the regression bias. However, it indirectly promotes the more general bias of representativeness. Also, the sophistication of the technology and the fact that large volumes of data can now be analysed can result in a false sense of security, as discussed previously.

From the perspective of the framework for information systems and judgement

In Chapter 5, a framework has been developed for investigating the role of information systems in judgement (Figure 5.1). The framework shows the mediating role of the information system between user and decision task, as well as the influence of the various roleplayers contributing to the information system, on judgement and consequently on biases. In the Ops Room, the interpretation and representation of the decision task by external roleplayers, such as the ones providing the performance data for the selected function of the offices, influences the work of the analysts, whose interpretations of the task as well as the decision-making process in turn influence the judgement of the department's decision-makers. During the process of data analysis, some of the external roleplayers were not available which made it difficult to clarify assumptions and have left room for misinterpretations and biases.

6.7. Conclusion

During the discussion of the case study, the main research question has been implicitly addressed, namely:

- How does information system – supported decision-making get affected by biases?

It has been shown that the bias of regression effects has been successfully countered in the analysis and reporting of the data in the case study. However, the role of information systems in enhancing or reducing biases in this particular case is limited. Apart from the regression bias, a number of other possible biases have been identified that can impact the success of decision-making based on exception reporting in the department concerned. It was found useful to view biases in a more general sense, such as discussed in the literature on biases and decision-making, rather than to work with Tversky and Kahneman's narrowly defined biases related to probability theory.

The case study also shows that reducing biases and providing high quality analyses does not necessarily improve the rationality of decision-making in a politicised environment. Analysts need to acknowledge the existence of a political model of decision-making as opposed to an idealised rational model, as discussed in Chapter 3. At the same time, analysts have a

responsibility to communicate their results clearly and reduce the potential of reports being misinterpreted and false impressions being created when people do not appreciate the limitations of data analysis.

7. FINDINGS AND CONCLUSIONS

7.1. Introduction

This chapter will summarise the findings of the study, provide a succinct criticism of the main assumptions encountered during the course of the study, and suggest possibilities for further research.

7.2. Summary of the findings of this study

7.2.1. Models of decision-making

The first derived research question to be addressed, was:

- How do people make decisions?

A number of decision-making models from the literature have been discussed in Chapter 3. From these models, the following insights were gained.

The *rational model* of decision-making is a normative model towards which many analysts strive. The Subjective Expected Utility provides a way to calculate the optimum or best decision. The psychologists that have researched decision-making biases, such as Tversky and Kahneman, make use of probability theory and in particular Bayesian statistics (rather than the SEU) to provide a benchmark for the best decision. Other decision researchers, such as Russo and Schoemaker, and Hammond *et al*, have derived their own prescriptive models of decision-making. The close resemblance between their models and the rational model indicate that they also advocate rational decision-making. Hogarth's model of judgement, another derivation of the rational model, shows decision-making that is based on the cognitive school of psychology, taking into account both the decision-makers mental processes and cues from the environment.

Simon's model of *bounded rationality* describes rational decision-making but take into account the constraints imposed by human abilities and needs as well as the structure of the decision

environment. It provides a context for understanding the concept of *heuristics* or short cuts that people take when they make decisions. Tversky and Kahneman saw biases as the negative side effects of heuristics.

The *organisational procedures* view alerts one to the routinised thinking that becomes part of an organisation's operating procedures and is likely to become embedded in information systems. These procedures can also be regarded as heuristics, with the same benefits and dangers attached.

Gary Klein's Recognition-Primed Decision model, an example of *naturalistic decision-making*, describes how people make use of pattern recognition and classification in their decision-making. According to some researchers, these skills are what humans naturally excel at, rather than the arithmetic skills measured in the heuristics and biases experiments. Other natural skills that feature in Klein's model include those of mental simulation and storytelling.

The *political model* of decision-making describes aspects of how many organisations function in reality. In a politicised environment, rational methods are of limited use, unless they can be used as a basis for positioning. As argued in the case study in Chapter 6, analysts need to take cognisance of the political environment of the decision-makers whom they provide with reports and information systems. They need to communicate their findings clear enough so that it is understandable and sellable to decision-makers and their clientele, but also to prevent misinterpretation or misuse of findings.

7.2.2. Heuristics and their application

In response to the question

- What is a heuristic?

a number of definitions from the literature have been provided in chapter 3. As opposed to the mathematicians' use of the term (referring to algorithms that give satisfactory rather than optimal solutions to computationally complex problems), the decision researchers' definition refers to *informal reasoning strategies*, which are often employed unconsciously in order to simplify and deal with complex decisions.

The subsequent question, namely

- How are heuristics used in the process of decision-making?

was explored in Chapters 3 and 4 by investigating the use of heuristics as described in the literature. Mentioned below are some of the heuristics that were found and discussed.

Tversky and Kahneman's heuristics, namely availability, representativeness, adjustment and anchoring, and problem framing, have a *statistical foundation*. They performed numerous laboratory experiments where they gave their subjects multiple-choice questions on probability theory, and found them to use the heuristics mentioned.

The following heuristics were observed to be used in *strategic organisational decision-making*: simplification, using past case experience, imitation, risk avoidance, searching for satisficing rather than optimal decisions, and cooperation.

Heuristics were also found to be used in a *social context*, such as social imitation and the one-of-us-effect. As shown by Dietz and Stern, some socially valid heuristics, resulting from eg. driving passions, altruism, religion, propaganda and the mass media, are likely to contravene rationality.

7.2.3. Enter biases

It was found that when heuristics are studied by researchers adhering to normative models of behaviour, judgement or decision-making, the emphasis is placed on the deviation from optimal decision-making as a result of the use of heuristics. On the other hand, researchers emphasising descriptive models of decision-making, such as the organisational procedures view or naturalistic decision-making, describes the usefulness of heuristics within the mentioned model of behaviour. The difference in viewpoints was observed in researchers' response to the question:

- What is the effect of heuristics on the outcomes of decisions or judgements?

Tversky and Kahneman's work focus on the errors of judgement, or biases, resulting from the use of heuristics. They show that novices as well as experts make classical mistakes when having to estimate probabilities. Hogarth follows their thinking from an information processing point of

view. He takes the various stages of information processing, which he terms information acquisition, information processing, output and feedback, and discusses for each step the possible biases when heuristics are used to deal with that step. More recent contributions that follow the same line of thinking are those of Russo and Schoemaker, and Hammond *et al*. These authors focus on biases associated with decision-making. They apply their work in an organisational context, giving managers advice on typical traps that they can fall in, and how to improve their decision-making. They prescribe steps for decision-making similar to the rational model. The impression left by the above authors can be summarised in the words of Slovic *et al*, quoted in Chase *et al*, (1998): "It appears that people lack the correct programs for many important judgemental tasks… it may be argued that we have not had the opportunity to evolve an intellect capable of dealing conceptually with uncertainty."

The stream of research that defends human decision-making has different yardsticks to those of the rational model for measuring the appropriateness of decisions. For example, Dietz and Stern argue that people's analytical decision-making abilities might not be very advanced when it comes to arithmetic and algebra, but people are very sophisticated when it comes to pattern recognition and classification. Klein, whose recognition-primed decision model relies strongly on the latter capabilities, has shown empirically that people are competent decision-makers. Dietz and Stern further emphasise the social context of decision-making: they believe it is more important for the human species' survival to be socially adept than to be good with arithmetic or to be able to maximise utilities. Various other examples were provided that support Dietz and Stern's view concerning the social context.

Given the contrasting views above, the position that has been taken in this study is that, since rational decision-making is held to be the norm by many, biases will be found, whether these are referred to as psychological traps, flaws in judgement or general mistakes in decision-making. It is acknowledged that the occurrence of biases is possibly over-emphasised in some of the literature, or not viewed in context.

Biases and information systems

Information systems and in particular decision support systems have been developed to supplement human information processing and to assist with human decision-making. It could

therefore be assumed that information systems, when used as decision aids, would have some influence on decision-making biases. In chapter 5, the role of information systems in decision-making has been investigated, with the focus on the role of information systems in introducing, reinforcing or reducing biases of decision-making. In doing so, the main research question of the study has been addressed, namely:

- How does information system – supported decision-making get affected by biases?

It has been found that information systems *have the ability to introduce new biases* (such as automation bias) *and to reinforce biases*. It has also been shown in one case that the mere use of information technology does not reduce biases. On the other hand, there were reports of the successful *use of information systems to reduce biases*. In the latter respect, the systems did more than just automating information processing; information was communicated in a different manner, whether by using different problem formulations, mental representations or richer communication media.

A framework has been developed, based on the work of Hogarth, to show the role of information systems in judgement and in particular with respect to biases. The framework shows the information system in a mediating role between user and task. It indicates the twofold role of information systems in judgement, namely to act as an extension of a decision-maker's mental schema (and thus provide support with the cognitive process of decision-making) as well as its role in representing and interpreting the task environment, of which it forms part. The framework also indicates the influence of other people on how the task environment is represented and interpreted and the way in which the information system provides judgemental support.

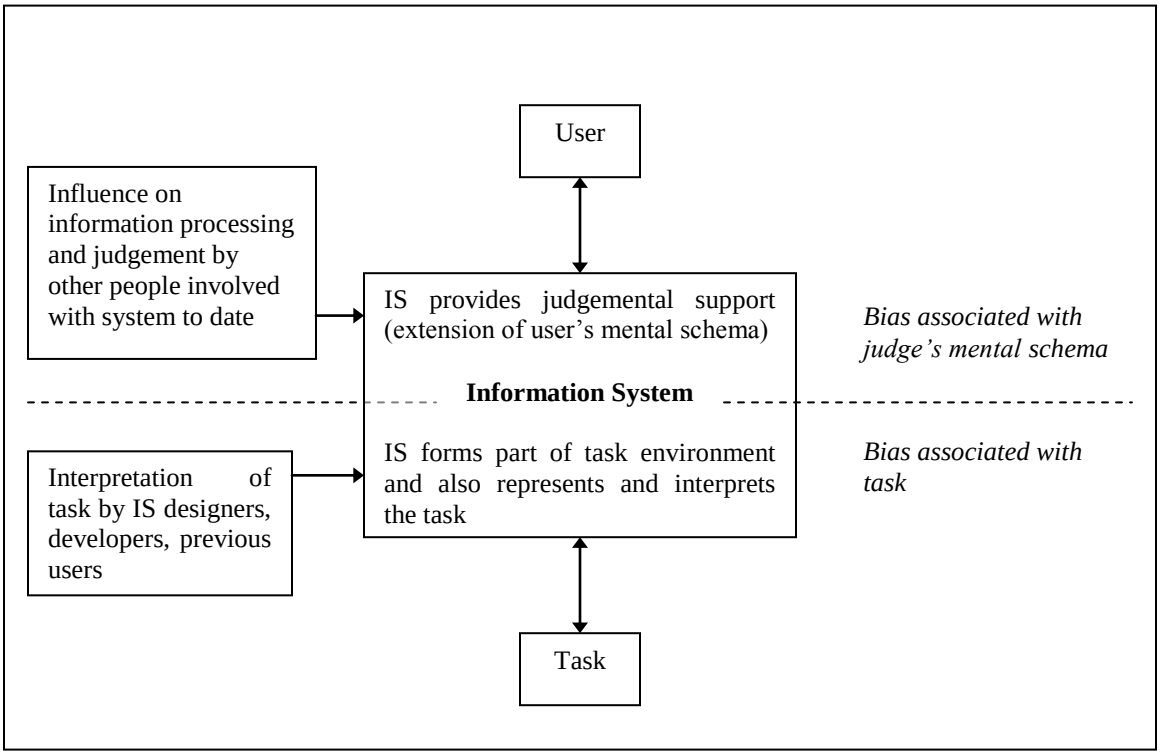

Figure 7.1: The role of the information system in judgement

From the perspective of the case study

In chapter 6, a case study was analysed to investigate the role of information technology and -analysis with respect to a particular bias. The reported project aimed at providing management information for a government department. Exception reporting formed an important part of the project. Since the bias of regression effects has been associated with exception reporting in the literature, evidence was sought of this bias on the project.

It was shown that the bias of regression effects had been successfully countered in the analysis and reporting of the data in the case study. However, the role of information systems in enhancing or reducing biases in this particular case is limited. Apart from the regression bias, a number of other possible biases have been identified that can impact the success of decision-making based on exception reporting in the department concerned. It was found useful to view biases in a more general sense, such as discussed in the literature on biases and decision-making, rather than to work with Tversky and Kahneman's narrowly defined biases related to probability theory.

The case study also showed that reducing biases and providing high quality analyses does not necessarily improve the rationality of decision-making in a politicised environment.

The framework shown in figure 7.1 was found to be applicable, highlighting the influence of various roleplayers on possible biases of decision-making. In particular, the data received contained the assumptions of a number of other roleplayers. The fact that some of these roleplayers were either not known or not available made it difficult to clarify assumptions, leaving room for misinterpretations and biases.

7.3. Criticism of assumptions made in heuristics and biases literature

A number of final claims will be made regarding assumptions that were encountered in the literature considered for this study.

Rational decision-making

Biases lead to errors of judgement as compared to a model of perfect rationality in decision-making. However, rationality is not always the benchmark. Dietz and Stern argue that it is more important to adapt socially than to be a super-calculator. In an organisational environment, alternative models of decision-making may prevail, such as the garbage can or political model. Under a political model, being politically naïve and perfectly rational will bring one nowhere; at best it will cause frustration.

Information processing

The cognitive model of psychology recognises internal mental processes (driving the individual rather than the other way around) as well as behavioural factors, which are determined by the environment. It is claimed that the cognitive model does not adequately provide for theories of motivation, ‘irrational’ drives such as religion and peer pressure, and the narrative mode of cognition.

Laboratory experiments and reality

An impression is created that laboratory experiments present a neater and more controllable version of reality. What if it does not represent reality at all? A number of debiasing techniques, such as those described by Fischhoff, are human activities that occur spontaneously in real life, but are difficult to perform in laboratory settings. The significance of self-organising and self-correcting behaviour in real life (that is not possible in laboratories) is underestimated in the heuristics and biases literature.

7.4. Suggestions for further research

Biases and group decision-making technology

Group decision-making is a field in its own right that has received very little attention in this study. Group decision-making is a significant component of organisational decision-making: managers do not make decisions behind closed doors. In the South African government environment, participative decision-making has become the norm.

Group support systems are technologies specifically catering to assist group decision processes. Group decision-making does not only have to be supported by specialised software that calls decision-makers together in one room and forces them to communicate via computers. Web-enabled software, such as chat rooms, is used to create communal spaces and communication facilities for virtual communities. Groupware, such as e-mail, is frequently used in an organisational context as a decision support tool, allowing people to share feedback on ideas and to have heated debates.

Group decision-making has its own set of associated biases, such as groupthink and social loafing. Based on the literature searches performed for this study, it is perceived that the topic of how group decision-making biases are affected by the use of group support systems is even less researched than the topic of how decision-making biases in general are affected by the use of information technology. It is therefore suggested as an area for further research.

Investigating the role of biases in real-life settings

As mentioned above, the majority of biases reported are identified and measured in laboratory settings. The case study in chapter 6 has shown that other factors can have a more significant impact on decision-making than biases. These include all the non-rational components of decision-making, such as the political environment, and other 'noise' factors such as the unreliability of data. In real situations, biases might not be the main concern of decision-makers. The investigation of more case studies, such as the Vincennes disaster, the expert system case study, and the study on exception reporting in chapter 6, could assist in understanding the 'real' role of biases in decision-making.

Further investigation of computer-induced biases

Examples of computer-induced biases were found in the study of automation bias, the Vincennes disaster (where the information system design and display were unhelpful), and the expert system, in which people placed a false trust. It is suggested that the biases mentioned, namely social loafing in an automated setting, false trust in systems and errors resulting from unhelpful user interfaces, as well as others, might be more widely present in organisational settings, and that their presence should be investigated.

7.5. Conclusion

This chapter has summarised the conclusions of the study. In response to the main research question, namely

- How does information system – supported decision-making get affected by biases?

it has been found that information systems have the ability to introduce new biases and to reinforce biases. Information systems can also reduce biases, but this requires innovate thinking on the way information is represented and the way human decision-making processes are supported. It has also been found that in the real world, other than the laboratories where biases are usually measured, other constraints on rational decision-making, such as politics or data errors, can overshadow the effects of biases.

REFERENCES

Banville, C and Landry, M (1989), Can the Field of MIS be Disciplined? in: Galliers, R (1992): *Information Systems Research: Issues, Methods and Practical Guidelines*, Blackwell, Oxford, pp 61 – 88.

Boland, RJ, Tenkasi, RV and Te'eni, D (1994): Designing Information Technology to Support Distributed Cognition, *Organisation Science*, Vol 5 No 3, pp 456 – 475.

Boland, RJ (2001): The Tyranny of Space in Organisational Analysis, *Information and Organisation*, Vol 11 (2000) pp 3 – 23.

Buchanan, JT and Corner, JL (1997): The Effects of Anchoring in Interactive MCDM Solution Methods, *Computers and Operations Research*, Vol 24 No 10, pp 907 – 918.

Burrel, G & Morgan, G (1979): *Sociological Paradigms and Organisational Analysis*, Arena, Hampshire.

Chase, VM, Hertwig, R and Gigerenzer, G (1998): Visions of Rationality, *Trends in Cognitive Sciences*, Vol 2 No 6, pp 206 – 214.

Cohen, MD, March, JG and Olsen, JP (1972): A Garbage Can Model of Organisational Choice, reprinted in: March, JG (1988): *Decisions and Organizations*, Basil Blackwell, Oxford.

Dahlbom, B and Mathiassen, L (1993): *Computers in Context*, NCC Blackwell, Malden, MA.

Das, TK and Teng, BS (1999): Cognitive Biases and Strategic Decision Processes: an Integrative Perspective, *Journal of Management Studies*, Vol 36 No 6.

Dietz, T and Stern, PC (1995): Towards a Theory of Choice: Socially Embedded Preference Construction, *Journal of Socio-Economics*, Vol 24 No 2, pp 261 – 279.

Du Plooy, NF (1998): An Analysis of the Human Environment for the Adoption and Use of Information Technology, unpublished PhD thesis, University of Pretoria.

Du Plooy, NF (2003): The Social Responsibility of Information Systems Developers, in: Clarke *et al* (eds): *Socio-technical and Human Cognition Elements of Information Systems*, Information Science Publishing, Hershey.

Du Plooy, NF, Introna, LD and Roode, JD (1994): Notes on Research in Information Systems, Working Paper, Department of Informatics, University of Pretoria.

Fischhoff, B (1982): Debiasing, in: Kahneman, D, Slovic, P and Tversky, A (eds) (1982): *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, pp 422 - 444.

Fisher, CW and Kingma, BR (2001): Criticality of Data Quality as Exemplified in Two Disasters, *Information and Management*, Vol 39, pp 109 – 116.

Fox, CR and Levav, J (2000): Familiarity Bias and Belief Reversal in Relative Likelihood Judgement, *Organisational Behaviour and Human Decision Processes*, Vol 82 No 2, pp 268 – 292.

George, JF, Duffy, K and Ahuja, M (2000): Countering the Anchoring and Adjustment Bias with Decision Support Systems, *Decision Support Systems*, Vol 29, pp 195 – 206.

Gilovich, T, Griffen, D, Kahneman, D (eds) (2002): *Heuristics and Biases: The Psychology of Intuitive Judgement*, Cambridge University Press, Cambridge.

Hammond, JS, Keeny, RL and Raiffa, H (1998): The Hidden Traps in Decision Making, *Harvard Business Review*, Sept – Oct 1998.

Hammond, JS, Keeny, RL and Raiffa, H (1999): *Smart Choices: A Practical Guide to Making Better Decisions*, Harvard Business School Press, Boston, Massachusetts.

Harvey, JT (1998): Heuristic Judgement Theory, *Journal of Economic Issues*, Vol 32 No 1, pp 47 – 63.

Hogarth (1980): *Judgement and Choice: The Psychology of Decision*, John Wiley & Sons, New York.

Huber, GP (1981): The Nature of Organisational Decision Making and the Design of Decision Support Systems, *MIS Quarterly*, June 1981.

Introna, LD (1997): *Management, Information and Power*, Macmillan, London.

Jennings, D and Wattam, S (1994): *Decision Making: An Integrated Approach*, Pitman Publishing, London.

Kahneman, D and Tversky, A (1973): On the Psychology of Prediction, Psychological Review, Vol 80, pp 237-251, reprinted in: Kahneman, D, Slovic, P and Tversky, A (eds) (1982): *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, pp 48 – 68.

Kahneman, D, Slovic, P and Tversky, A (eds) (1982): *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.

Kahneman, D, and Tversky, A (eds) (2000): *Choices, Values and Frames*, Cambridge University Press, Cambridge.

Keen, PGW and Scott Morton, MS (1978): *Decision Support Systems: An Organisational Perspective*, Addison-Wesley, Reading, Massachusetts.

Klein, G (1998): *Sources of Power: How People Make Decisions*, MIT Press, Cambridge, Massachusetts.

Krabuanrat, K and Phelps, R (1998): Heuristics and Rationality in Strategic Decision Making: An Exploratory Study, *Journal of Business Research*, Vol 41 (1998) pp 83 – 93.

Kreitner, R and Kinicki, A (2001): *Organizational Behaviour*, 5th Ed, Irwin McGraw Hill, Burr Ridge, IL.

Lim, KH, Benbasat, I and Ward, LM (2000): The Role of Multimedia in Changing First Impression Bias, *Information Systems Research*, Vol 11 No 2, pp 115 – 136.

Lipshitz, R, Gilad, Z and Suleiman, R (2001): The One-of-Us Effect in Decision Evaluation, *Acta Psychologica*, Vol 108 (2001), pp 53 – 71.

March, JG (1988): *Decisions and Organizations*, Basil Blackwell, Oxford.

Marsh, B (2002): Heuristics as Social Tools, *New Ideas in Psychology*, Vol 20 (2002), pp 49 – 57.

Nicolau, AI (1999): Social Control in Information Systems Development, *Information Technology and People*, Vol 12 No 2, pp 130 – 147.

Online Dictionary (2003): Decision, Judgement, http://dict.die.net

Pfeffer, J (1981): *Power in Organisations*, Pitman Publishing, Marshfield MA.

Rai, A, Stubbart, C and Paper, D (1994): Can Executive Information Systems Reinforce Biases? *Accounting, Management and Information Technology*, Vol 4 No 2, pp 87 – 106.

Roode, JD (1994): Implications for Teaching of a Process-Based Research Framework for Information Systems, Working Paper INF WP/009, Department of Informatics, University of Pretoria

Roy, MC and Lerch, FJ (1996): Overcoming Ineffective Mental Representations in Base-Rate Problems, *Information Systems Research*, Vol 7 No 2, pp 233 – 247.

Russo, JE and Schoemaker, JH (2002): *Winning Decisions: How To Make the Right Decision the First Time*, Piatkus, London.

Shore, B (1996): Bias in the Development and Use of an Expert System: Implications for Life Cycle Costs, *Industrial Management and Data Systems*, Vol 96 No 4, pp 18 – 26.

Simon, HA (1956): Rational Choice and the Structure of the Environment, *Psychological Review*, Vol 63 No 2, pp 129 – 138.

Simon, HA (1977): *The New Science of Management Decision* (2nd ed), Prentice Hall, Englewood Cliffs, New Jersey.

Simon, HA (1979): Rational Decision Making in Business Organisations, *American Economic Review*, Vol 69 No 4, pp 493 – 513.

Skitka, LJ, Mosier, KL and Burdick, M (1999): Does Automation Bias Decision-making? *International Journal of Human-Computer Studies*, Vol 51, pp 991 – 1006.

Slovic, P, Finucane, M, Peters, E and MacGregor, DG (2002): Rational Actors or Rational Fools: Implications of the Affect Heuristic for Behavioural Economics, *Journal of Socio-Economics*, Vol 31, pp 329 – 342.

Stair, RM and Reynolds, GW (2001): *Principles of Information Systems: A Managerial Approach*, Course Technology, Cambridge, MA.

Todd, P and Benbasat, I (1991): An Experimental Investigation of the Impact of Computer Based Decision Aids on Decision Making Strategies, *Information Systems Research*, Vol 2 No 2, pp 87 - 115.

Todd, P and Benbasat, I (1992): The Use of Information in Decision Making: An Experimental Investigation of the Impact of Computer Based Decision Aids, *MIS Quarterly*, Vol 16 No 3, pp 373 - 395.

Tversky, A and Kahneman, D (1974): Judgement under Uncertainty: Heuristics and Biases, *Science*, Vol 185.

Tversky, A and Kahneman, D (1981): The Framing of Decisions and the Psychology of Choice, *Science*, Vol 211, pp 453-458.

Walsham, G (2001): *Making a World of Difference: IT in a Global Context*, Wiley, Chichester.

Wright, G (1984): *Behavioural Decision Theory: An Introduction*, Penguin Books, Harmondsworth, England.