A critical discussion of Jon Elster's arguments about rational choice, infinite regress and the collection of information

How much information should you collect before making a decision?

Dissertation for the Cand. Polit. degree

Hans O. Melberg

Department of Economics
University of Oslo

May 1999

A critical discussion of Jon Elster's arguments about Rational Choice, Infinite Regress and the Collection of Information

How much information should you collect before making a decision?

by Hans O. Melberg

SUMMARY

Jon Elster has argued that (a) the optimal collection of information creates an infinite regress (we have to collect information about how much information to collect about how much information to collect …) and (b) it is often very difficult to estimate the net value of more information because we have to make decisions in situations that are fast changing, unique or novel. Taken together these two arguments imply, according to Elster, that it is often impossible to make a rational choice and this, in turn, means that economists should focus more on theories of satisficing (people acting when they have an alternative that is "good enough") and less on theories based on optimization. The purpose of the dissertation is to examine these arguments.

The answer to whether rational choice implies an infinite regress and whether it is impossible to rationally collect the optimal amount of information clearly depends on the definition of rationality (Chapter 2). To examine this question, I first present the standard economic theory (expected utility theory). This theory, however, places few restrictions on what probabilities the agents are allowed to use (and still be labeled rational). In response to this some authors (like Jon Elster) argue that rationality implies not only that we chose the optimal action for given beliefs, but also that these beliefs are optimal for given information *and* that the collection of information must be optimal. Others, like Russell Hardin, argue that it is not intuitive to demand that a decision must be based on an optimal amount of information before it is labeled rational. The issue is important because if Hardin is right (that rationality does not demand optimal collection of information), the problem of infinite regress in the collection of information does not even get off the ground. I argue, however, that Hardin is wrong on the intuitiveness of excluding information-gathering (Is it rational to buy a house without collecting information?). There are also good methodological reasons (like parsimony) for making the optimal collection of information a requirement of rationality (Why make a different assumption about the choice of information than the choice of action?).

Chapter 3 is mainly descriptive. I first present the standard economic theory for determining the value of information (based on Hirshleifer and Riley). The presentation is not only descriptive since it turns out that there are several factual mistakes in their account. Next follows an overview of Elster's arguments why the standard theory is circular (infinite regress) and in any case not very useful since our estimate of the relevant probabilities often is non-existent or very weak. I end the chapter on a more argumentative note. Elster's cannot both argue that it is *impossible* to collect an optimal amount of information and at the same time argue that it is *sometimes possible* (for instance, he admits that it is possible to make an optimal decision in highly stereotyped decisions like the amount of information to collect before a medical diagnosis). Moreover, I question the appropriateness of Elsters reference to S.G. Winter to justify the infinite regress argument. Winter himself does not use the term and he has never provided a formal proof. In fact, in an article from 1975 Winter only says that the problem of optimization "may involve" the logical problem of self-reference (which is far

from proving that this is so, or that this self-reference creates an infinite regress, or - the third step - that the infinite regress is a problem i.e. that it does not converge).

Regardless of the consistency of the two arguments together and the appropriateness of the reference, both the infinite regress argument and the problem of estimation can be evaluated separately. In Chapter 4 I try several possible interpretations of the infinite regress argument. Some are rejected as not being compatible with Elster's argument (radical skepticism because of infinite regress in induction, infinite regress in deliberation - as opposed to search, and infinite regress in deciding how to decide because optimization itself is costly). Elster's problem is infinite regress *in the collection of information.* I offer two possible interpretations of this (see Figure 4.1 and Figure 4.3 in the dissertation). First, the demand for an optimal amount of information always require the solution of a new optimization problem (optimal collection of information about information and so on). Second, infinite regress makes the set of available options at any point in time infinite. One could collect information (say, about the quality of a house), or one could collect information about information (buy a book about how to collect information about houses), or one could collect information about information about information (a magazine with a review of several books about how to collect information about houses) and so on ad infinitum (theoretically). On the first problem, I believe time-constraints ends the infinite regress. On the second, time-constraints may also eliminate the regress, but there is a complication because of the possibility of "saving" time. Collecting information at a high level may tell you to collect less information at a lower level than you initially believed. I end the chapter by presenting a more formal example of infinite regress that is less focused on Elster (and more on Winter's arguments).

In Chapter 5, I first try to structure Elster's many arguments about the problem of estimating the net value of information. I distinguish between three types of probability (objective, theoretical and subjective), three types of problems (non-existent probabilities, weak probabilities and biased probabilities) and three types of implications (use maximin when uncertain; randomization is better - since it is more honest - than trying to maximize expected utility when you have weak beliefs; do not waste resources seeking information about the accuracy of the second decimal when the first is unknown). I then discuss the validity and internal consistency of Elster’s arguments under four headings. First, the argument about the non-existence of probabilities is weak because it relies heavily on the classical notion of probability as relative frequencies (in contrast to theoretical and subjective probabilities). Second, I disagree that randomization is better than maximization of expected utilities as a general decision-rule because *in aggregate* even small differences in probability is significant. Third, biased probabilities are not relevant to the debate about the theoretical possibility of rational choice. Fourth, I use a digression into the economic theories of search to argue against the view that we cannot conduct a rational search because our knowledge is limited in a way that is comparable to "being lost in a forest." The last argument is not entirely successful since it turns out that optimal search strategies require some kind of initial assumption. It is not possible to do a rational search for information when you are "radically uncertain" (but one may question whether it is possible to be "radically uncertain").

As mentioned in the first paragraph Elster claims that the argument about problems in the collection of information shows that rational choice theories are inherently indeterminate and that economist should focus more on behavioral theories, such as satisficing. The argument is often accompanied by an argument against Friedeman's "as-if" justification for optimization (the argument being that the environment changes so fast that the selection mechanism does not have time to weed out the non-optimal agents). The purpose of Chapter 6 is to discuss these two arguments. First, it follows from my argument from Chapter 5 that I do not believe infinite regress in the collection of information is a good justification for behavioral economics. There may, however, be other infinite regress problems (e.g. in deliberation and expectation formation in strategic environments) that resurrects the impossibility of rational choice. The computational demands implied by the problem of estimation may also justify a more behavioral approach. Second, the argument that the "as-if" justification is invalid because the environment changes fast, overlooks the possibility that some agents are adaptive precisely because the environment is fast changing. Finally, it is possible to use the "as-if" argument without the literal analogy with natural selection in biology. One simply adopts the instrumentalist position that unrealistic assumptions are not important as long as the predictions are accurate. The discussion in Chapter 6 is anchored in a brief digression on the nature of economics since this is what the discussion is about i.e. what kind of questions economist should try to answer and how they should do so.

In conclusion (Chapter 7), I do *not* claim to have answered the general question of behavioral vs. traditional economics. I do, however, claim to have pointed out some of the weaknesses in Elster's arguments which are relevant to that debate. I also point out some areas where more work is needed, focusing especially on Hey's suggestions that we should follow "reasonable" rules when deciding how much information to collect " (as opposed to trying to follow optimal rules that may be very bad if we make small mistakes).

A critical discussion of Jon Elster's arguments about rational choice, infinite regress and the collection of information

How much information should you collect before making a decision?

Dissertation for the Cand. Polit. degree

Hans O. Melberg

Department of Economics
University of Oslo

May 1999

Preface

In the preface to *Ulysses and the Sirens*, Jon Elster (1979/1984, viii) writes that "to fail is always to fail *at something*, and it leaves you with a knowledge of the kind of thing you unsuccessfully tried to do." The same can be said about this paper for two reasons. First, I failed to do what I originally set out to do. Second, I failed to give a comprehensive, formal and reliable answer to the question I ended up discussing.

Originally I wanted to "explain" economic fluctuations and my approach was based on the belief that a good explanation must have microfoundations (for more on microfoundations, see Janssen 1993). As I researched this topic, I encountered two major difficulties. First, there were serious problems in the measurement of the business cycle, which in turn meant that an empirical test of the theory was difficult. Second, the justification of my psychological microfoundations approach was based on Jon Elster's argument that it is often impossible to make a rational decision—for instance about investment—since rationality cannot tell us how much information we should collect. For this reason theories of economic fluctuation could not be based on rational choice alone and we were forced to use more psychological theories. However, as I examined the argument more closely I became more and more convinced that it was weak. In short, I was left with an approach that was based on flawed justification and which in any case would be very difficult to test even if it could be build on better foundations. In this failure, however, there was one achievement: the discovery that Jon Elster's argument about the indeterminacy of rational collection of information may be wrong. I decided to build on this and made it the focus of my thesis.

Even after making this my focus, I failed to give a completely satisfactory answer to the underlying question of how much information to collect before making a decision. However, I believe the attempt to answer the question may still be valuable. First, while I cannot give a comprehensive positive answer, the negative aim of arguing against some theories is still possible. Second, given the cumulative nature of academic work, it is perfectly acceptable to drop the aim of comprehensiveness and focus on some aspects of the problem in question. Third, given inherent limitations of time, space and personal abilities I would end up with a very poor result if I aimed for the first best comprehensive and formal answer. For these and other reasons I continued to work under the title "How much information should you collect before making a decision?" even though I never labored under the illusion that I would be give a perfect answer.

I once read a joke on the theme that copying from one book was called plagiarism, while copying from several books was called a dissertation. In preparing this work I have not only used other people's ideas, but I have also asked people for comments and advice. My supervisor, Olav Bjerkholt, greatly improved the precision of the arguments and corrected several embarrassing errors in previous versions. Ole J. Røgeberg made detailed comments on many of the arguments—especially about subjective probability—which led me to revise some of my initial beliefs. Per Ariansen, Jon Elster, Jack Hirshleifer, Barton L. Lipman, Roy Radner, Tore Schweder and Atle Seierstad took the time to answer what sometimes must have seen like childlike questions, and for this I am grateful. Finally, I should thank Timur Kuran and Barry Weingast who sent me copies of forthcoming papers. Support from a project led by Pål Kolstø and financed by the Research Council of Norway, allowed me to go to the Association for the Study of Nationalities' conference in New York in 1998 to learn more about rational choice theories of ethnic violence. The discussion of Russell Hardin's arguments in chapter two owes much to that project. None of the above, of course, can be blamed for the remaining shortcomings of this paper.

Hans O. Melberg
Oslo, May, 1999.

(e-mail: hmelberg@hotmail.com)
(academic home-page: http://www.geocities.com/hmelberg/papers/papers.htm)

Contents

1 Introduction

1.1 What is the question?

The starting point of this paper is the following question: How much information should we collect before making a decision? The short answer is that we should collect information as long as the expected value of spending more resources on collecting information is greater than the expected cost. But, how do we know the expected value and cost of more information? To answer this we need to collect information, i.e. we have to collect information to determine how much information to collect. As the reader may already have understood, this apparently leads to an infinite regress. We must collect information on how much information to collect before we decide how much information we should collect and so on forever. This is *the problem of infinite regress* in the collection of information, and some authors—for instance Jon Elster and Sidney G. Winter—argue that it is a serious problem in the theory of rational choice.

The infinite regress problem is only one possible source of indeterminate answers when trying to decide how much information to collect. For instance, when we are in a unique situation we cannot determine the value of information from historical experience of similar situations, and hence there is—using the classical view of probability—no rational basis for estimating the value of information. I have labeled all the problems in this (residual) category *the estimation problem*.

The question is then whether it is correct that the infinite regress argument and the problem of estimation make it impossible to act rationally.

1.2 What kind of a question is this?

This section has two aims: To limit the scope of the inquiry and to locate the questions I want to discuss within the larger map of questions and problems in economics. The point of departure is the outline presented in Figure 1.1. In the paragraphs below I will only provide the minimum explanation necessary to understand the figure. It is the task of the other chapters (especially chapter two) to put more meat on the bones.

First of all, there is a very general disagreement about the assumptions on which economists should base their theories, that is the discussion between those who define

Figure 1.1: Locating and limiting the question

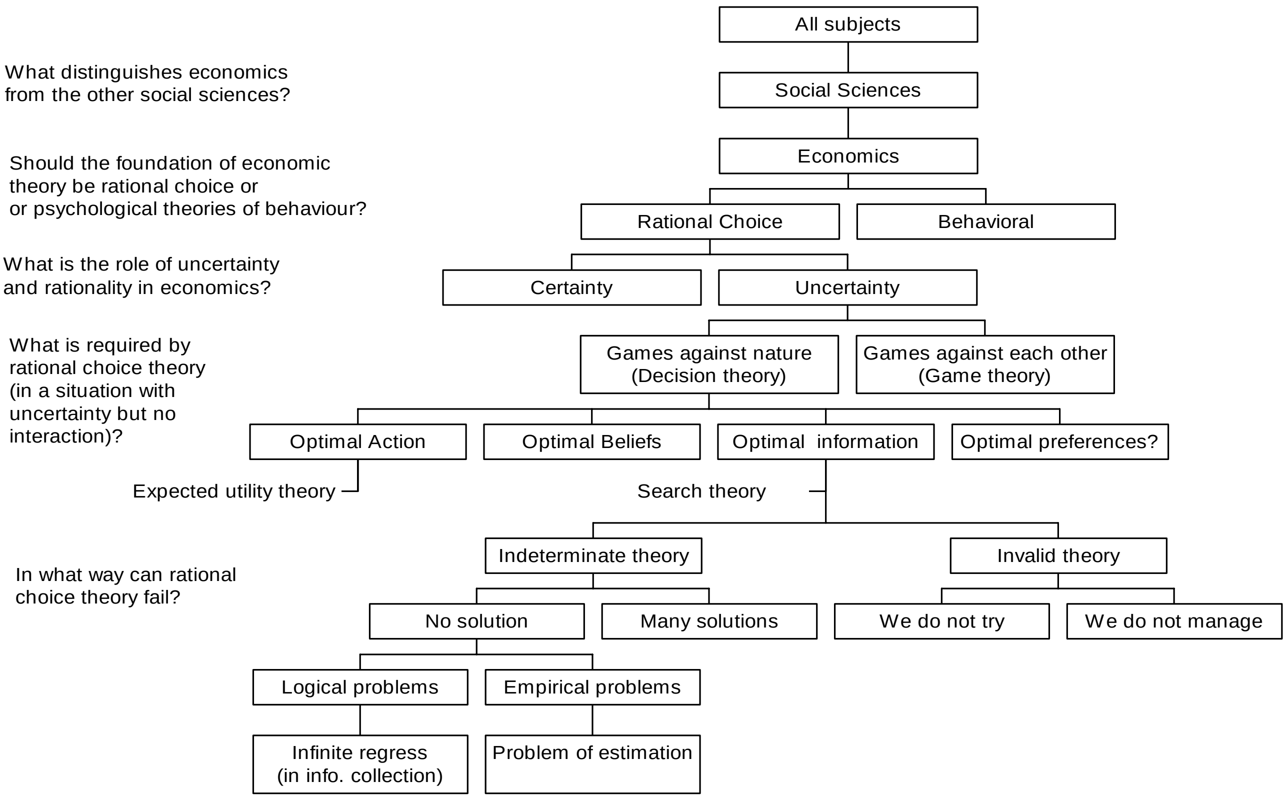

economics as the application of rational choice analysis vs. those who focus on more psychological or behavioral foundational assumptions.[1] The two problems discussed in this paper—infinite regress in the collection of information and the problem of estimating the net value of information—are two of very many issues in this debate.

The complexities of the definition of rational choice depends, among other things, on what kind of environment we are assuming: certainty vs. uncertainty; one-person vs. multi-person situations. For some time I worked on the problems of indeterminacy both in a multi-person and a one-person world. There is, for instance, a problem of finding solution concepts in game theory that are strong enough to generate unique predictions. There is also the problem of backward induction and infinite regress of the type "I believe that you believe that I believe..." In the end I decided that the focus was too broad, and limited myself to problems in one-person decision theory with uncertainty.

[1] The distinction between rational choice economics and behavioral economics sounds strange when one realizes that rational choice is simply one of several behavioral theories. Thus, it would be more correct to locate "rational choice" as a sub-category under "behavioralism."

The next level in the figure indicates the various demands that could be made before one agrees that the decision is rational. Is it enough that the choice of action is rational for given beliefs, or should we also require beliefs, preferences and the collection of information to be optimal? This is the main topic of Chapter 2.

If we agree on the definition of rational choice, there are only two possible arguments against this theory. First, it may be impossible to act rationally, for instance because the theory demands that the agent must have information which it is in principle impossible to have. I have followed Elster's terminology and labeled the problem indeterminacy. Second, the theory may be invalid in the sense that people do not act according to the theory even when it is possible to do so. I have limited my attention to problems of indeterminacy i.e. the argument to be evaluated is not of the type "this is not how people act", but rather "it is impossible to act according to the theory because the theory does not present a solution." Whether the theory is indeterminate or invalid depends, of course, on the definition of rationality.

The two general reasons for the failure of rational choice theory (invalid or indeterminate) can be divided into two sub-groups. A theory can be invalid because people *do not try* to behave rationally (believing, for instance, that adherence to norms and moral rules is the correct way of deciding how to act). The second reason for judging a theory invalid could be that people *do not manage* to behave rationally even when they try (due to, for instance, cognitive limitations; see Melberg 1998 for some examples). Indeterminacy may exist when the theory has *no solution* and when there are *many solutions* (for instance many Nash equilibria). I have already limited my topic to "indeterminacy," but I will further limit myself to the problem of "no solution" as opposed to "many solutions."

There may well be diminishing return to even further classification, but I believe that the two problems discussed in this paper do not really belong to the same class. They are both about why it may be impossible to collect an optimal amount of information, but the problem of estimation—unlike the problem of infinite regress—appeals to epistemological difficulties. More information may reduce the estimation problem, while the problem of infinite regress is inherent to the process of optimization. That is why I have labeled one "logical problems" and the other "empirical problems." I do not attach much prestige to these last two labels as long as the two problems are kept separate the reader is free to label them as he want (if he wants).[2]

[2] Part I of Arrow et al. (1996) is devoted to "rational choice and logical problems", but there is no clear definition of the term "logical problems." One intuitive meaning could be arguments of the type "the theory/argument is not correct because the implications of the various premises contradict each other." There is no possible world in which all the statements could be true at the same time. See also Elster (1989c, pp. 40-43) for a discussion of true logical paradoxes (like Russell's paradox).

Finally, I discuss only one (of many) "logical" and only some (of many) "empirical" problems that (supposedly) makes it impossible to collect an optimal amount of information. It is here—at last—that the problem of infinite regress in the collection of information and the problem of estimation can be located. Even more specifically, my main (but not exclusive) argument is focused on Jon Elster's interpretation of the problem of infinite regress and estimation.

1.3 Why try to answer this question?

What makes a question worth asking and answering? First of all, the answer should not be obvious or the obvious answer incorrect. Second, the question should be important. Third, it must be possible in principle to give an answer. In this section I will try to relate these requirements to the problem of deciding how much information to collect. The aim is both to demonstrate that I am not flogging dead horses (i.e. there is disagreement today) or arguing against straw men (i.e. I show that there are people who think the collection of information does not represent a problem).

1.3.1 The answer should not be obvious, or the obvious answer should be incorrect

One good way of demonstrating that the answer is not obvious is to show that the "experts" disagree. For instance, on the question under consideration Roy Radner (1996) seems to believe that the gathering of information can be solved rationally. As he writes:

> It is convenient to classify the costly (resource-using) activities of decision-making into three groups:
>
> 1. *observation*, or the gathering of information;
> 2. *memory*, or the storage of information;
> 3. *computation*, or the manipulation of information.
>
> [...]
>
> 4. *communication*, or the transmission of information.
>
> Of these activities, and their related costs, the first, second, and fourth can be accommodated by the Savage paradigm with relatively little strain, although they do have interesting implications. *(Radner 1996, p. 1363)*

In apparent contradiction to this view, we may quote Jon Elster (1985) who believes the problem of information collection is significant. He writes:

> In most cases it will be equally irrational to spend no time on collecting evidence and to spend most of one's time doing so. In between there is some optimal amount of time that should be spent on information-gathering. This, however, is true only in the objective sense that an observer who knew everything about the situation could assess the value of gathering information and find the point at which the marginal value of information equals marginal costs. But

> of course the agent who is groping towards a decision does not have the information needed to make an optimal decision with respect to information-collecting.[23] He knows, from first principles, that information is costly and that there is a trade-off between collecting information and using it, but he does not know what that trade-off is. *(Elster 1985, p. 69)*

The term contradiction may be too strong to describe the difference between the two quotations. Radner claims that information collection can be "accommodated by the Savage Paradigm" but he does not discuss whether mere consistency of subjective beliefs (which is what is required within the Savage paradigm) is sufficient to label the decision rational. Thus, the difference between Radner and Elster may be that Radner is willing to label a decision rational as long as it is based on consistent beliefs, while Elster places stronger demands on beliefs, e.g. that they should be rationally constructed for a given set of information. This is a topic I will discuss closer in the second chapter, which is a general introduction to rational decision-making. In any case, the quotations prove that there is a difference in the degree to which decision theorists view the collection of information as a problem (see sub-chapter 4.1 for more examples of conflicting views on this issue). Whether it is a substantial disagreement or a mere problem of labeling remains to be discussed.

1.3.2 The answer should have important implications

Not all non-obvious questions are worth asking. For instance, assume you have spent much time and effort finding the answer to a non-obvious question ("What is the twenty-first decimal of π?"), but that few or no important consequences follow from answering the question. It seems like you could have made better use of your time trying to answer a different and more important question.[3] Sometimes people react in this way to the opening paragraph of this paper. "So what"; "Who cares?" and "This is simply too abstract to be of practical use" were some of the comments. I disagree, but before I can explain why it is necessary to discuss the meaning of "importance."

Clearly, "importance" is a subjective term so what is important to you need not be important to me. Although the previously mentioned reaction ("who cares?") could be dismissed on this ground ("I don’t care what you say, it is important to me!"), I think this would be

[3] Olav Bjerkholt has pointed out to me that the concept of importance is more complicated than I initially believed. First, it is difficult to know in advance whether a question is important. What appears to be a purely intellectual problem or discovery may turn out to have very important applications. It may also be true that important discoveries are made as by-products while trying to answer "impossible" questions. Finally, working on "impossible" problems need not be useless when it leads to the conclusion that the problem is undecidable (which is useful knowledge).

wrong. It is wrong because I believe the reaction stems *not* from thinking that the implications are unimportant, but from being unaware of the true implications. Hence, I will try to persuade the reader by making some of the implications explicit.

We should distinguish between implications that derive their importance from being directly relevant for decision-making versus those implications that have intellectual importance. If I could give a good answer to the question of how to reduce the problem of unemployment, this would immediately be of importance to the welfare of many people (which may be one commonly agreed meaning of importance). The question of how much information to collect is important both for decision-makers and for academics.

Medics can exemplify a group of decision makers for whom the problem of information collection is crucial. Before deciding upon a diagnosis they have to make up their minds about how much information to collect. The choice of how much information to gather may determine the difference between life and death for some (see, for instance, comments by doctor Jan-Henrik Pederstad about the difficulties in distinguishing meningitis from a flu in *Dagbladet*, Tuesday 5 January, 1999, p. 21). Here is another example I recently encountered: The person in charge of the towing of a large platform construction was told by the engineers the day before the planned towing operation was scheduled to start that they wanted to collect more information before they went ahead with the operation (to the cost of $100 000 in addition to the delay in time). How should he decide whether to collect more information or not? Is it possible to make an optimal choice in this situation? (In this case, it was decided that the information was not worth the cost). In short, the question in this paper has direct relevance to many important decisions: medical diagnostics, buying a house, making an investment, choosing your career and so on.

Although I shall discuss some of the direct real-life implications in Chapter 5, this paper is mainly focused on the theoretical problem of information collection ("is it possible at all") and in this sense it is only a necessary conceptual preliminary to the practical question of "how much should be collected." The implication of this conceptual problem may not be directly relevant to decision-makers, but it is important in an intellectual sense in the debate between the "behavioralist" and the rational-choice economists. If it can be proved that it is logically impossible to make a rational choice, then the behavioralists will have a better case against the traditional economists.

In short, besides the inherent intellectual pleasure in satisfying our curiosity about the answer to a question and the direct importance for practical-decision-making, I believe that the question in this paper is important because of its role in the debate on the basic building blocks of economic theories. Should the foundation be the assumption of rational choice or should it be psychological theories about behavior. It is, however, important to note that I do not pretend to give general answers to that question. As illustrated by Figure 1.1, all I do is to discuss two of many issues that are relevant to the larger debate, without claiming that these two alone determine the argument.

1.3.3 Is it possible to answer the question and if so how?

Some questions are interesting (or non-obvious) and important, but there is very little hope of determining the answers with any degree of reliability and for somebody interested in deriving useful policy-advice this reduces the value of spending time on them. Hence, my third demand for it to be worth trying to answer a question is that it is at least in principle possible to give an answer.

Not all questions have objectively true answers. For instance, it is in principle impossible to give an objectively true answer to the question of whether vanilla or chocolate ice cream is the best ice-cream flavor. More seriously the question "What is just?" does not have a unique and objectively true answer.[4] As for the topic of this paper, it is true that there is no single definition of rationality that everybody accepts. This does not mean that it is impossible to discuss the question of rational collection of information in a scientific manner.

Imagine that I claim to have the correct definition of rational choice. Another person may then criticize this definition by pointing to some of its logical implications that I had not considered or by showing that the definition is not coherent (because all its elements cannot be true at the same time[5]). I may then agree that the definition was wrong, and revise it accordingly. Alternatively, I may claim that his implication does not represent a counterintuitive example that should lead us to revise the theory.[6] Finally, I may try to prove that the claimed implication does not follow from my definition. There is nothing "unscientific" about this

[4] It is easy to understand that questions involving taste and values do not have unique and objectively true answers. A question may, however, be impossible to answer conclusively even if it does not involve taste. I discovered this while reading Newman's (1987) entry on Ramsey in *The New Palgrave. A dictionary of economics* and I make no claim to understand the proofs (Gödel's incompleteness theorem) or the example itself (but I do understand another example of the same general problem, Russell's paradox). For what it is worth, Newman writes about Ramsey's (unsuccessful) work on the Continuum hypothesis which—it was later discovered by Paul Cohen—is an undecidable proposition.

[5] Arrow's Impossibility theorem is a good example how a limited number of requirements can be proved to be logically incompatible.

[6] As an example of a result that leads some people to revise the theory and others to revise their behavior, consider Allais Paradox. Allais (1987) himself believes the paradox shows that the theory is wrong. Other's believe it shows that people are wrong (do not act rationally), not the theory of rational behavior (e.g. L. Savage, who himself gave the "wrong" answers).

debate. It is true that the process of confronting each other with the implications of the definition need not lead to a unique definition of rationality, but that is different from arguing that the process itself is unscientific (for more on this, see Melberg 1996, pp. 475-477 and Elster 1993, pp. 180-181).

Having established that it is in principle possible to discuss the problem in a scientific manner, it remains to be argued exactly how we should answer the question. I have previously admitted that I shall proceed by dividing the question into smaller parts and then select only some of these for closer investigation. Moreover, I have chosen to do an in depth examination of a few arguments from a few authors instead of a short discussion (or survey) of many arguments using unnamed ("some economists say")—or worse: imagined—opponents. As a small compensation for the restricted scope, I have made extensive use of footnotes to direct the interested reader to relevant literature. Finally, on the question of formal and abstract mathematics vs. verbal reasoning, I have opted for a mainly verbal style. These choices were made mostly out of necessity.

One final aspect should be commented on since it may seem peculiar to some. I think it is important to state the weaknesses of my own arguments and sometimes I will indicate the degree to which I am unsure. This is not only a question of academic honesty. By telling the reader about my own uncertainty, I make it easier for those who want to scrutinize my argument. This is important because, as mentioned in the preface, science is a cumulative effort and there is no reason to make this cumulative work more difficult than it already is by hiding uncertainties behind confident language.

2 What is rational choice?

2.1 Introduction

If we disagree on the definition of rationality we may also give different answer to the question of whether it is possible to make a rational decision about how much information to collect. To avoid misunderstandings of this sort, it is necessary to discuss the concept of rationality in general. It is not necessary, however, to discuss all aspects of every possible definition of rational choice. I will limit myself to a short presentation of the standard theory in economics (expected utility theory) and a closer discussion of those elements that are relevant to the problem of infinite regress and the problem of estimation. It is, for instance, very important to discuss whether we should require that the collection of information and the construction of beliefs be optimal before we label the decision rational. Without this requirement it is not possible to speak of rationality creating an infinite regress in the collection of information. It is not important to discuss whether rationality also demands that we exclude certain types of preferences—such as acting on preferences the agent "knows are impossible to fulfill" (Nozick 1993, p. 144).[7] These discussions are often said to be about *substantial* rationality,[8] as opposed to the economists more *instrumental* concept of rationality (i.e. they are about what we should want, not only what we should do to get something we want). Although interesting, these deeper philosophical issues are not relevant to the infinite regress argument or the problem of estimation.

2.2 Rational Choice in Economics: Expected Utility Theory

Imagine that you have to select one action (x_i) from a set of feasible actions (X). Assume, moreover, that you are in a situation of uncertainty.[9] Which action should you choose?

---

[7] In his book *The Nature of Rationality,* Robert Nozick (1993, pp. 139-151) devotes a sub-chapter to discuss the many demands he wants to make on preferences before they should be called rational.

[8] Not to be confused with Simon's concept of *substantive* rationality. Substantial is the opposite of instrumental and the distinction is associated with Max Weber. Substantial rationality is a broadening of the concept of rationality to include preferences (see Sen 1987). Simon's (1987, p. 17) concept of substantive rationality is the opposite of procedural rationality and represents a narrow concept of rationality in which the process is viewed as insignificant—only the end result is important in judging something as rational or irrational.

[9] Uncertainty is here not distinguished from the concept of risk (se footnote 38). When this distinction is important, and it is not obvious from the context what the term means, I shall use phrases like "radical uncertainty" or "total ignorance" to indicate a situation in it is not even possible to assign numerical probabilites to the different

In its prescriptive variant expected utility theory says that we should choose that action which maximizes expected utility. As Schoemaker (1982) points out, the theory can also be used descriptively ("this is how people choose"), predictively ("I expect him to choose x since it is the act that maximizes expected utility") and postdictively i.e. it is used as a non-falsifiable assumption that guides research. Anomalies do not falsify the theory, but stimulates search for the unknown variable that makes behavior conform to the theory. My interest here is mainly in the normative aspects of the theory i.e. whether it can tell us what to do. In the following I shall thus present the basics of the theory with special emphasis on exactly what is required before the theory can be used prescriptively.

How do we calculate expected utility? To answer this we first specify our uncertainty as a list of possible "states of the world" (each state is denoted s_i and is a member of the set of possible states S). We then have a list of possible actions and possible states that together form the set of possible consequences (c_{xs}). A simple example is the following: You have to choose whether to bring an umbrella or not when you go for a walk (x_1= bring umbrella, x_2= not bring umbrella). There are two possible "states of the world": s_1= it will rain, s_2= it will not rain. Cross-tabulating this we have the following four possible consequences:

Figure 2.1: Calculation of Expected Utility

		Possible states (S)	
		s_1 Rain (probability p_1)	s_2 No rain (probability p_2)
Possible actions (X)	x_1 (Bring umbrella)	c_{11} (it rains and you have an umbrella)	c_{12} (you brought the umbrella, but is does not rain)
	x_2 (Do not bring umbrella)	c_{21} (it rains and you did not bring an umbrella)	c_{22} (you did not bring the umbrella and it did not rain)

The expected utility of an action is calculated by multiplying the utility of each possible consequence of an action with the (subjective or given) probability that the consequence will occur. Formally in our example:[10]

states. See Lawson (1988) and Davidson (1991) for some non-standard views on uncertainty and probability. See also Kelsey and Quiggin (1992) for a survey of theories of choice under uncertainty (both senses).

[10] Hirshleifer and Riley (1992, p. 13) argue that much confusion has been created by not distinguishing between the utility of consequences and the utility we expect to result from an action. To avoid this they suggest using $V(\cdot)$ to denote the utility of consequences and $U(x)$ to indicate the utility derived from the consequences of an act. I have followed this suggestion.

$$EU\ (x_1) = v(c_{11})\ p_1 + v(c_{12})\ p_2 \qquad (2.1)$$

$$EU\ (x_2) = v(c_{21})\ p_1 + v(c_{22})\ p_2 \qquad (2.2)$$

Or, more generally:

$$EU\ (x_i) = \Sigma\ v(c_{xs})\ p_s \qquad (2.3)$$

Maximization of expected utility then simply means that you choose that alternative which has the highest expected utility when it is calculated in the way described above.

So far all I have done is to describe exactly how one calculates the expected utility from an action. How can this procedure be justified as *the* rational way of making a choice? The answer is that the decision rule "maximize expected utility" (MEU) follows from what some people think are appealing axioms.[11] More specifically, define a prospect (y) as the pairings of consequences and probabilities associated with an action. For instance, the action "Bring an umbrella" is associated with the following prospect: $y_1 = [c_{11}, c_{12}; p_1, (1-p_1)]$. Having defined a prospect, we then assume that the preference over prospects satisfy the following:[12]

1. *Completeness* (For any prospects the agent must either prefer one to the other or be indifferent i.e. "I do not know" is not allowed.)

2. *Continuity* (When faced with a good, a medium and a bad prospect [$y_1 \geq y_2 \geq y_3$] there must be some probability that makes the agent indifferent between a lottery involving y_1 and y_3 and the sure prospect y_2: $(y_1, y_3; p, 1-p) \sim y_2$. This implies that lexicographic preferences are not allowed. [~ symbolizes indifference.])

3.*Transitivity* (if you strictly prefer y_1 to y_2 and y_2 to y_3, the you must also prefer y_1 to y_3)

4. *Independence* (If you are indifferent between y_1 and y_2, and $y_1 \geq y_3$, then it should also be the case that $y_2 \geq y_3$) (This is comparable to Savage's sure-thing principle).

Von Neumann and Morgenstern showed that if we accept these axioms, then it follows that we should use the MEU rule to choose between the possible actions. When these axioms are satisfied, the *ranking of actions* using the calculation of expected utilities correspond to the

[11] See Hampton (1994) for a dissenting view.

[12] Different authors present this in slightly different ways. See Haregreaves Heap (1992, p. 9), Machina (1987a, p. 86), Schoemaker (1982, pp. 531-532) and Schmeidler and Wakker (1987, p. 74). See also Dawes (1988, chapter 8) for a very good introductory presentation of von Neumann and Morgenstern's theory.

ranking of consequences.[13] The intuitive idea that it was reasonable to choose on the basis of the expected utilities of actions may have been around for a long time (at least since Bernoulli's solution to the St. Petersburg Paradox), but it was von Neumann and Morgenstern who rigorously proved that the MEU rule followed from what many believe are appealing axioms.

The Expected Utility Hypothesis has been extensively discussed, and especially the fourth assumption (independence) has been questioned. The purpose of this section, however, was not to present a detailed review of the debates around the hypothesis (see, for instance, Machina 1987b or Sugden 1991). Instead I simply wanted to describe the basics of the theory and make its assumptions explicit.

2.3 Should we demand more or less than Expected Utility theory?

What is the source of the probabilities used in calculation of expected utility and what—if any—restrictions should be place on the construction of probability estimates before we are willing to label the decision rational? Von Neumann and Morgenstern's theory says little about this and they simply take probabilities as objectively given. In contrast to the few restrictions placed on beliefs, they include some axioms that in no way are obvious demands of rationality. It is not rationality that requires us to have complete and continuous preferences (there is, for instance, nothing inherently irrational about non-continuous [lexicographic] preferences). However, this criticism should not be drawn too far since von Neumann and Morgenstern did not present their axioms as demands of rationality. They are rather conditions that must be satisfied if the MEU rule is justified as the rational way of making a decision.

If we focus on the construction of beliefs, we find that different authors disagree on the degree to which the construction of beliefs should be made a part of the definition of rationality. The debate has at least two aspects: One intuitive, the other methodological. First, what does our intuition tell us about the rationality of including the formation of beliefs and the collection of information in the definition of rationality. Second, one could use methodological criteria like parsimony and fruitfulness to justify (or deny) the claim that rationality

[13] The utility function is cardinal in the sense that any linear transformation of it will preserve the ranking and the ratio of the differences between the alternatives. It is not cardinal in the sense that 10 utilities represent twice the pleasure of 5 utilities. Moreover, the cardinality does not imply the possibility of interpersonal comparison of utility.

should be the working-hypothesis at all levels—be it the choice of action for given beliefs or the formation of beliefs.

To examine the first issue I have chosen to discuss two opposing views, that of Jon Elster and Russell Hardin. My argument is, in short, that Elster is right about the intuitiveness of demanding that our knowledge be rationally constructed before we label the decision rational.

2.3.1 The intuitive argument: Elster vs. Hardin

Elster's views on the definition of rational choice can be summarized by the following quotation:

> Ideally, a fully satisfactory rational-choice explanation of an action would have the following structure. It would show that the action is the (unique) best way of satisfying the full set of the agent's desires, given the (uniquely) best beliefs the agent could form, relatively to the (uniquely determined) optimal amount of evidence. We may refer to this as the *optimality part* of the explanation. In addition the explanation would show that the action was caused (in the right way) by the desires and beliefs, and the beliefs caused (in the right way) by consideration of the evidence. We may refer to this as the *causal part* of the explanation. These two parts together yield a first-best rational-choice explanation of the action. The optimality part by itself yields a second-best explanation, which, however, for practical purposes may have to suffice, given the difficulty of access to the psychic causality of the agent. *(Elster 1985, p. 71)*

According to this view there are two general demands that have to be met before we can use rational choice to explain an action: First, the demands of optimality. Second the demands of causality. The demands of optimality can be divided into three requirements: optimality in the choice of action from the feasible set, optimality of beliefs for a given set of information, and optimality in the collection of information. The two causal demands require that action and beliefs be caused "in the right way" given preferences, beliefs and evidence. For instance, assume it is rational for me to press a green button (not the red), and I do so. We would not call this a rational action if the reason I pressed the green button was that somebody pushed me and I accidentally hit it. The same goes for beliefs. I might, for example, make two errors when calculating probabilities, but these two errors could cancel each other out so the final belief is optimal. This is an example of evidence causing the beliefs in the wrong (non-rational) way.

I now turn to a critical examination of the opposing view i.e. the view that we should not demand that our estimates be constructed in a rational fashion before we label the action rational. To do so I shall use Russell Hardin's (1995) arguments from his book *One for All: The Logic of Group Conflict*. Some may believe it is a bit on the side to discuss the rationality of individual action in ethnic conflicts in a paper on economics, but I believe there are good reasons for focusing on Hardin: He is a well respected academic (so I am not attacking a soft target); he knows the general topic well (he is an authority on game theory, collective action and rationality); and he discusses the specific question head on (should we require beliefs and the collection of information to be rational before we label the decision rational?). The fact that the context is ethnic violence and not, say, investment, makes little difference to the principles involved.

The aim of Hardin's book, expressed in his own words, is "to go as far as possible with a rational choice account of reputedly primordial, moral, and irrational phenomena of ethnic and nationalist identification and action" (Hardin 1995, p. 16). A short summary of his theory of ethnic violence goes as follows. It is rational to identify with a group since it provides both security, material benefits and satisfies a psychic need to belong somewhere. Being a member of a group affects your beliefs since it tends to reduce awareness of alternative ways of doing things, as well as inducing the belief that what "we" do is the right thing to do (the is-ought fallacy). Given these beliefs, it becomes rational for people who want power to play on people's ignorance and the belief that we are "better" that the other groups. Finally, group violence happens when the leaders find it the best way of maintaining power (for instance to distract people from economic failure). Using nationalist propaganda, they create a belief that it is in people's self-interest to engage in a pre-emptive attack against the other group. Once violence starts there is a spiral that only increases violence, since it creates hate as well as an even stronger belief that one must destroy the other side before they kill us (and there is no way the parties can credibly promise not to discriminate or destroy each other).

Although this to some extent is a plausible story, we have to ask whether it is intuitive to label it rational. More specifically, is the formation of beliefs behind nationalism and ethnic violence rational? Hardin admits that beliefs used to explain group conflict are "not convincing, even patently not so in the sense that it would not stand serious scrutiny..." (Hardin 1995, p. 62, emphasis removed). But how can it be rational to act on beliefs that are obviously wrong? Hardin's answer is worth quoting in at length:

> One might say that the supposed knowledge of ethnic or national superiority is corrupt at its foundation. Unfortunately this is true also of other knowledge, perhaps of almost all knowledge of factual matters. [...] Hence, at their foundations there is little to distinguish supposed knowledge of normative from that of factual matters [...] Should we say that anyone who acts on such knowledge is irrational? We could, but then we would be saying that virtually everyone's actions are always irrational. It seems more natural to say that one's beliefs may have corrupt foundations but that, given those beliefs, it is reasonable to act in certain ways rather than others if one wishes to achieve particular goals.
> [...]

> Someone who carries through on an ethnic commitment on the claim that her ethnic group is in fact superior, even normatively superior, to others, may not be any more irrational than I am in following my geographic knowledge. She merely follows the aggregated wisdom of her ethnic group. *(Hardin 1995, pp. 62-63)*

In short, because all knowledge is corrupted at its base it "would be odd [...] to conclude that the action was *irrational when taken* if *it was fully rational given the available knowledge*" (Hardin 1995, p. 16, emphasis in the original).

Who is correct, Hardin or Elster? First of all, there are several internal inconsistencies in Hardin's argument. For instance, even if we agree that rationality demands only optimality for given information, it is difficult to see how people can believe that the individuals in their ethnic group descend from one "original" Eve. This is a common belief among nationalist (see Connor 1994). Hence, "patently false beliefs" do not require collection of information to be falsified; they may be irrational even for a given set of knowledge. I also fail to understand how he can write that it is odd to label an action irrational when it was rational for given beliefs and at the same time write that "a full account of rational behavior must include the rationality of the construction of one's knowledge set." (Hardin 1995, p. 16). Yet another inconsistency is revealed by his attack on communitarianism. Hardin writes:

> The chief epistemological problem with particularistic communitarianism is that it violates the dictum of the epigraph of this chapter: The important thing is not to stop questioning [...] To question such beliefs is to increase the chance of bettering them. *(Hardin 1995, p. 192)*

> Commonsense epistemology allows for variations in our confidence of our knowledge. My belief that concern for human welfare dominates concern for various community values or even for community survival is radically different from my belief that certain rough physical laws hold sway over us. *(Hardin 1995, p. 210)*

If it is true that all factual knowledge is corrupt at its foundation (his justification for not making the construction of one's beliefs a part of the definition of rationality), then we should put little faith in the proposition that we can increase the reliability of our beliefs and values by questioning them (which is his justification for rejecting communitarianism).[14]

Second, we might question the argument that all knowledge is equally corrupt at its foundation. As I shall discuss later (sub-chapter 5.5), Jon Elster and Leif Johansen have made similar claims, but they do not go this far. Is it really true that *all* our knowledge is so weak that none of the differences are worth seeking out or acting on?

Third, and perhaps most important, is the suggestion that to demand that we should construct the set of knowledge in a rational way must lead us to conclude that "virtually eve-

[14] Like Hardin I reject communitarianism as a political philosophy (but for different reasons), so my criticism is not politically motivated.

ryone's actions are always irrational." My immediate response would be that his argument leads to an equally odd conclusion: To reject the demand for rational construction of beliefs leads us to conclude that many intuitively irrational actions really are rational. For instance, imagine a person who buys a house or a used car without first collecting some information about its quality. Even if the decision was rational for the given (weak) information, it sounds odd to label the decision rational.

A better solution than to reject the demand that we should collect an optimal amount of information, I believe, is to at least make some demands on the collection of information. It may be true that we do not know the optimal level, but it is still possible to know that we should collect some information. In short, the demand is that: "One should collect an amount of evidence that lies between the upper and lower bounds that are defined by the problem situation..." (Elster 1985, p. 71). This may leave a large zone of indeterminacy in between, but at least it excludes some options as irrational. Moreover, this demand on information collection does not commit me to the position that almost all actions are irrational, as Hardin claims.[15] Finally, his argument makes it far too easy (and uninteresting) to prove that a phenomenon is caused by individually "rational" action. In sum, I believe it is intuitive to place demands both on the construction of beliefs and the collection of information before we label a decision rational.

2.3.2 The methodological arguments

In addition to the arguments about the intuitive appeal of including belief formation and information collection in the definition of rationality, we may add several methodological arguments. To understand these methodological arguments, it is useful to take a short look at the history of economic thought.

Roger E. Backhouse (1995) has described the modern trend in economics as follows:

> In the post-war period economic theory has been dominated by the attempt to explain economic phenomena in terms of rational behaviour. In macroeconomic this has taken the form of providing a microeconomic foundation for macroeconomic theories: deriving macroeconomic relationships as the outcome of individuals' optimizing subject to the constraints imposed by their endowments, markets and technology. There has been an aversion to what Lucas has termed 'free parameters': parameters describing individual or market behaviour that are not derived from the assumption of utility or profit maximization. *(Backhouse 1995, p. 118)*

[15] We might improve the discussion by distinguishing between rational, non-rational and irrational actions. When it is impossible to act according to the rules of rational choice, the choice is not irrational, but non-rational.

Another trend is the invasion of economic reasoning into subjects previously thought to be outside the scope of economics. Political science, Sociology and even psychology has been increasingly influenced by rational choice theories. Marriage, divorce, crime, ethnic violence and even suicide have all been subject to rational choice analysis (for some comments on this trend, see Stigler 1984, Hirshleifer 1985, Becker 1986 and Demsetz 1997).

In sum, there are at least three developments. First, the increasing emphasis on microfoundations. Second, the argument that the best microfoundation is rational choice. Third, the tendency towards economic imperialism. Taken together these three developments say something about what kind of methodological criteria academics, and particularly economists, regard as valuable.

The underlying methodological view is one that conceives of progress in a discipline as explaining as much as possible using as little as possible at the deepest level possible and in a way that can be quantified (and hence tested). Or to use the terminology of methodologists: We want universalism, parsimony, reductionism and quantifiability. Searching for microfoundations means going deeper, using rational choice is—arguably—due to a commitment to quantifiability and economic imperialism represent the attempt to explain more.

What is the relevance of this discussion for the definition of rational choice? Recall Elster's demand that a decision is not rational unless there is optimality in (1) the choice of action, (2) the formation of beliefs for given information, and (3) the collection of information. Hardin argued against (3), and possibly (2). My argument here is that the very same principles that inspires those who favor the use and extension of rational choice theory, also implies that the definition of rational choice should include both (2) and (3). The application of the principle of maximization to both the formation of beliefs and the collection of information increases parsimony, increases the scope of a single principle, provides deeper microfoundations and increases quantifiability.

The rational expectation revolution is itself an implicit indication that many economists have accepted a stronger definition of rationality than just optimality of action for given beliefs. Before this revolution, one "free parameter" was the assumption of either rigid or only backward looking adaptive expectations. In the 1970s Lucas, again in the words of Backhouse (1995, p. 123), argued that optimizing behavior "should be applied systematically to all aspects of macroeconomic models, including the formation of expectations..." The same methodological argument applies to the collection of information.

2.3.3 Sub-conclusion

I have argued in favor of Elster's definition of rationality and against Hardin. The argument had two main aspects. First, there is the theoretical presupposition based on parsimony that if we assume maximizing behavior in the choice of action for given beliefs, then we should also assume it when people form beliefs and when they collect information. Second, when faced with some concrete examples it sounded intuitively wrong to exclude the formation of beliefs and the collection of information from the definition of rationality.

3 Collecting an optimal amount of information

3.1 Introduction

Is it true that the problem of information collection cannot be solved rationally? To answer this question, I shall first describe the standard theory as presented by Hirshleifer and Riley (1992). I then discuss Elster's arguments against the standard theory. I shall argue that:

1. There has been a tendency in Elster's writings about the impossibility of collecting an optimal amount of information away from an emphasis of the logical problems of infinite regress and towards the empirical problems of estimating the net value of information (see sub-chapters 3.3.1 and 3.3.2).
2. Elster was wrong in focusing on the problem of infinite regress in the collection of information as an important problem in rational choice theory (Chapter 4). The argument is, moreover, in conflict with his later admission that it is sometimes possible to collect an optimal amount of information (see sub-chapter 3.4).
3. There are several potential problems with Elster's treatment of the problem of estimation (Chapter 5). First, the argument about non-existent probabilities relies heavily on the classical relative frequency view of probability (this view may be wrong and/or in contradiction with some of his other arguments, see sub-chapter 5.2). Second, I disagree that the existence of weak probabilities implies that randomization is a better strategy than maximization of expected utility (sub-chapter 5.3). Third, biased beliefs are not relevant to the discussion of whether it is possible in principle to make a rational decision (sub-chapter 5.4). Fourth, the economic theories of search reduce the force of the argument that we are so lost that no rational search is possible.

3.2 The standard economic theory: Hirshleifer and Riley

The choice about whether to collect information or not can be viewed as any other choice: We should try to collect more information when the expected utility of this alternative is higher than the expected utility of the other possible alternatives. But exactly how do we work out the expected value of more information?

To illustrate their general answers these question, Hirshleifer and Riley (1992, p. 173) use the example of an agent who believes there might be oil in a field. In this situation the agent has to decide whether to drill a test-well or go ahead with a major investment without collecting more information. The structure of the prior beliefs is given by the agent's beliefs about the geological structure of the land and there are three such structures (favorable geological structure: 0.9 probability of hitting oil, moderate: probability 0.3 of oil, hopeless: impossible to find oil). In the terminology of expected utility theory there are three "states of the world." Before drilling the agent believes that the probability of a favorable geological structure is 0.1, the probability of a moderate structure is 0.5 and the probability of a hopeless structure is 0.4. Finally, Hirshleifer and Riley assume that the result of the test drill is not conclusive i.e. the result is only "wet" or "dry." Whether the result is "wet" or "dry" depends on the geological structure, and the probability of "wet" if the true state is "favorable geological structure" is 0.9, compared to 0.3 probability of wet for "moderate" and 0 probability of wet if the true state is "hopeless." Given all this rather condensed information, we now ask three questions. First, how should you estimate the probability of oil given the result from a test-drill? Second, what is the value of doing the test-drill (i.e. gather information)? Third, how much should you be willing to pay for an information service (e.g. about the geological structure of the land)?

To impose some order on the information, Hirshleifer and Riley use three different matrixes: the likelihood matrix (L), the joint probability matrix (J), and the posterior matrix (Π). The likelihood matrix specifies the probability of each message given the state of the world, P(m|s); The joint probability matrix gives the probability of each combination of states and messages P(sm); Lastly, the posterior gives the probability of a state of the world give a message, P(s|m);. Using the information above we have (Table 3.1):

Table 3.1: The likelihood, joint probability and posterior matrix

	The likelihood matrix (*$L \equiv q_{m/s}$*)			*The joint probability matrix* (*$J \equiv j_{sm}$*)				*The posterior matrix* (*$\Pi \equiv \pi_{s/m}$*)		
		Message			Message		π_s		Message	
		Wet (m_1)	Dry (m_2)		Wet (m_1)	Dry (m_2)	Prior beliefs		Wet (m_1)	Dry (m_2)
States of the world (Geological structure)	Favorable (s_1)	0.9	0.1	Favorable (s_1)	0.09	0.01	0.1	Favorable (s_1)	0.375	0.013
	Moderate (s_2)	0.3	0.7	Moderate (s_2)	0.15	0.35	0.5	Moderate (s_2)	0.625	0.461
	Hopeless (s_3)	0	1	Hopeless (s_3)	0	0.40	0.4	Hopeless (s_3)	0	0.526
				q_m	0.24	0.76				

At this point a short summary of the logic behind these calculations may be in order. The probability of both a message of wet and a favorable geological structure is, by definition, the probability of a favorable geological structure multiplied by the prior probability of that you will receive the message wet if the structure is favorable:

$$j_{sm} \equiv \pi_s \, q_{m|s} \qquad (3.1)$$

For instance, the probability of both "wet" and a favorable geological structure is: 0.1 * 0.9 = 0.09

To find the probability of a favorable structure given a message of wet, we simply divide the joint probability of wet and favorable by the overall probability of receiving "wet" (independent of structure):

$$\pi_{s|m} \equiv j_{sm} / q_m \qquad (3.2)$$

If we use the probability of a favorable structure given the message "wet", this is: 0.09/0.24 = 0.375

So far all we have done is to apply the definitions of probability. The next step is to note that we can use (3.1) and (3.2) to find a new expression of the posterior probability of a state given a message:

$$\pi_{s|m} \equiv \pi_s \, (q_{m|s} / \Sigma_s \, \pi_s \, q_{m|s}) \qquad (3.3)$$

This is often called Bayes' Theorem, but since it is a simple combination of two definitions it is sometimes more appropriately called Bayes' Rule. The intuition is quite simple. After the test drill you have two pieces of information relevant to the estimation of the probability of the various geological structures. First, the result of the test drill. Second, the prior beliefs about the geological structure. The final rational estimate is a combination of the two and Bayes' rule tells you how to rationally combine the two pieces of information.

Having considered the answer to the first question (how information should affect your beliefs) in some detail, it remains to answer the question of how much the information is worth ($\omega_{n,}$ i.e. worth of information measured in utilities). Hirshleifer and Riley (1992, p. 180) first define this as the difference between the expected utility you will receive when choosing an action based on current information vs. the expected utility of choosing an action after receiving information.

$$\omega_n = U(x_m; \pi_{s|m}) - U(x_0; \pi_{s|m}) \qquad (3.4)$$

However, the answer is slightly more complicated since the authors assume that people can only buy an "information service" (μ) and not one piece of information. Hence, in the oil-case we could buy a test, but we could not buy the result "wet" since the result of the test may be both wet and dry. Thus, the value of information is the expected difference in utility: the sum of the utility difference between your best action with and without information for each message (the sum of each difference multiplied by its probability; $\pi_{\bullet m}$ symbolizes the revised probabilities after receiving information):

$$\Omega(\mu) = E\ \omega_m = \Sigma_m\, q_m\, [U(x_m; \pi_{\bullet m}) - U(x_0;\ \pi_{\bullet m})] \qquad (3.5)$$

To work out the precise answer in terms of oil, we have to make some assumptions about the costs and gains. Assume the following payoffs: Drilling and wet: \$1 000 000; Drilling and dry: -\$ 400 000; No drilling: -\$50 000 (relocation costs). The example is also easier if we assume that the agent is risk-neutral (since this implies that the utility function is linear in income). Hirshleifer and Riley then ask how much one would pay for a geological analysis before the test drill. If we follow their example, they assume that the likelihood matrix of the geological analysis is as follows:

Table 3.2: The likelihood and posterior matrix (of the geological analysis)

		The likelihood matrix		*The posterior matrix*	
		Message		Message	
		Wet (m_1)	Dry (m_2)	Wet (m_1)	Dry (m_2)
States of the world (Geological structure)	wet	0.6	0.4	0.486	0.136
	dry	0.2	0.8	0.514	0.864

First we calculate the payoff from the best action before drilling. The expected value from drilling (x_1) is -\$64,000 [(0.24 * 1,000,000) - (0.76 * 400,000)]. This is worse than the expected loss of no drilling (-\$50,000). Hence, the optimal action before receiving information is "no drilling" with an expected payoff of -\$50,000.

The next step is to calculate the expected value of the optimal action after receiving information. If the message is "dry", the optimal action is "no drilling" with payoff -\$50,000. Hirshleifer and Riley then write that:

> … if the message is "wet," expected gain from "drilling" (action x = 1) becomes 0.486 (\$1,000,000) - 0.514(\$400,000) = \$140,400. So the expected value of the information is \$140,400 + \$50,000 = \$190,400. This is the value of the message service. *(Hirshleifer and Riley 1992, p. 183)*

Unfortunately, this conclusion did not correspond to my calculations and it turns out that there was a mistake in the book, as Jack Hirshleifer confirmed in an e-mail when I asked him. To find the correct answer we first have to find the marginal probability of the messages "wet" and "dry" (q_m), which I calculated to be (Table 3.3):

Table 3.3: The joint probability matrix (of the geological analysis)

		The joint probability matrix		
		Message		Prior probability
		Wet	Dry	π_s
State of the world	Wet	0.144	0.096	0.24
	Dry	0.152	0.608	0.76
	q_m	0.296	0.704	

To get the precise money value of the information service in this case, we note that as long as the utility function is linear in income (as assumed), we can replace the expressions for utility in (3.5) with money values. Define c^*_{sm} as the income from the best action in state s after receiving the message m. On the other hand, c^*_{s0} indicates the income from the best action without receiving more information. If we then use these expressions instead of utility in equation (3.5), we have that the money value of the information service is:

$$\Omega(\mu) = \Sigma_m q_m \Sigma_s \pi_{s|m} v(c^*_{sm}) - \Sigma_m \Sigma_s \pi_{s|m} q_m v(c^*_{s0}) \qquad (3.6)$$

$$\Omega(\mu) = \Sigma_m \Sigma_s \pi_{s|m} q_m v(c^*_{sm}) - \Sigma_s \pi_s v(c^*_{s0}) \qquad (3.7)$$

To find the value of information in our case, consider, first, the gain if the message is "dry" (the probability of this message is 0.704), In that case both the pre-message and post-message optimal action is the same (no drilling) and there is no gain in expected utility. If the message is "wet" (and the marginal probability of this message is 0.296) the optimal post-information action is "drilling" i.e. a change from the pre-message optimal action (no drilling). One could then simply calculate the money difference between the two alternatives and multiply this by the probability of the message "wet" to find the value of information. Once again, however, there is a mistake in Hirshleifer and Riley (this time in the last edition from 1995). Hirshleifer told me that the new edition reads as follows: the value of information "is 0.296($140,400 + $50,000) + 0.704(0) = $56,358" (personal communication). Although the probabilities are correct this time, the number $140,400 is wrong. The sum is meant to indicate the expected value from drilling after receiving the message "wet" and they incorrectly write that this is "0.486 ($1,000,000) - 0.514($400,000) = $140,400". The correct figure is: 0.486 ($1,000,000) - 0.514($400,000) = $280,400. The value of information then becomes: 0.296($280,400 + $50,000) + 0.704(0) = $97,798.

More generally, the maximum a person should be willing to pay for an information service is the ξ that solves the following:

$$\Sigma_m \Sigma_s \pi_{s|m} q_m v(c^*_{sm} - \xi) = \Sigma_s \pi_s v(c^*_{s0}) \qquad (3.8)$$

Figure 3.1: The value of information (based on Hirshleifer and Riley 1992, p. 183)

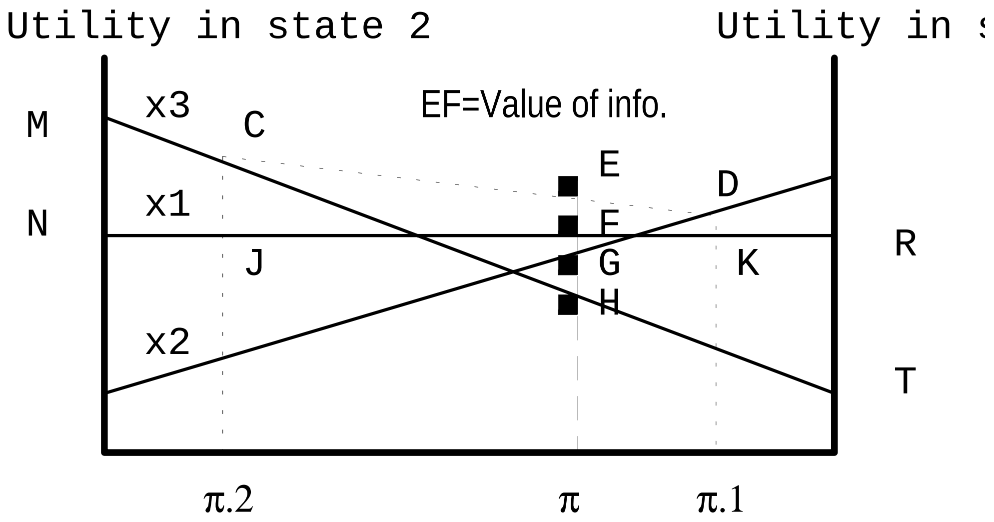

The value of information can also be visualized in a figure (Figure 3.1). In the figure there are three possible actions (x_1, x_2 and x_3) and two possible states (s_1 and s_2). If state 2 occurs the agent receives M if he has chosen x_3 and N if he has chosen x_1. On the other hand, if state 1 occurs, the agent would receive T if he had opted for x_3 and R for x_1. Clearly your choice of action would depend on your probability estimate of whether state 1 or 2 was most likely. Moreover, you would be willing to pay to receive information about the true probability. Imagine, for instance, a message service with two possible outcomes ($\pi_{.2}$ and $\pi_{.1}$). Initially the probability π indicates that the best action is x_1 with expected utility of F (the lines indicate the expected utility of the various actions at the different probabilities). If the message received is m=1, then the best action is x_2 (expected utility: D).[16] If the message is m=2, then the best action is x_3 (expected utility of C). If m=1, the expected gain in utility from having chosen x_2 over x_1 is the distance DK; similarly if m=2 the expected gain from having chosen x_3 over x_1 is given by CJ. The overall value (measured in utility) of the information service is the distance EF.

This concludes my treatment of the standard theory of rational choice in economics and the standard frame for determining the value of information. The presentation in Hirshleifer and Riley is relatively detailed, but it cannot be accused of spending much time on the underlying philosophical problems involved (e.g "how do we know the probabilities used to estimate the net value of information?") It is to these I now turn. More specifically, I want to investigate two problems as they have been described by Jon Elster. In the section below I simply present the arguments, leaving the task of evaluation to the next two chapters.

[16] Once again there is a mistake in Hirshleifer and Riley's book. They write (1992, p. 183) that the optimal action if m=1 is x_1, but their diagram clearly shows that x_2 is better.

3.3 Elster's arguments

3.3.1 A list of quotations

To enable the reader to follow the discussion, I have summarized Elster's writings on the impossibility of collecting an optimal amount of information in Table 3.4.

Table 3.4: Elster's arguments about the possibility of collecting the optimal amount of information

Year	Source	Key pages	Is "infinite regress" mentioned?	Reference to Winter?	Quotation
1978	*Logic and Society*	162 (173)	Yes	Yes (quoted)	One might argue that "… satisfaction emerges as a variety of m ing information are taken into account. [176] Winter, then, in creates an infinite regress, for how do you solve the problem 'choice of a profit maximizing information structure requires i profit maximizer acquires this information, or what guarantees *ster 1978, p. 162, quoting Winter 1964)*
1979/ 1984	*Ulysses and the Sirens*	58-60, 135	Yes	Yes (quoted)	"Take the case of a multinational firm that decides not to enter t costs of the operation would exceed the benefits.[51] Then we information to acquire before taking the decision not to acqui market. Unless one could prove (and I do not see how one cou converges to zero or at any rate rapidly becomes smaller for e tures, this argument not only has the implication that in every d tion stops and you simply have to make an unsupported choice, *the action as possible*. Why, indeed, seek for precision in the s *(Elster 1979/1984, p. 59)*
1982	Rationality Encyclopedia chapter	112-113	Yes	Yes	Many "argue that firms are profit-maximizers because otherw powerful because it is backed (Winter [6b] *[sic.]* by an infin planned profit-maximizing. The argument, briefly, is this. In or As information is costly, it would be inefficient to gather all the settle for the optimal amount of information. But this means th solved, only replaced by a new one, that immediately raises the s
1983	Review of Nelson and Winter in *London Review of Books*)	5, 6	Yes	Yes (quoted)	"The Nelson-Winter attack on optimality is therefore a two-pro cannot optimise *ex ante*, since they do not have and cannot get t they would need an optimal amount information, but this lead infinite regress. On the other hand, we cannot expect firms to does not operate with the same, speed and accuracy as it does ments strike at the root of neo-classical orthodoxy." *(Elster 198*
1983	*Explaining Technical Change*	139-140	Yes	Yes (quoted)	"One of his [S. Winter] contributions is of particular interest an notion of maximizing involves an infinite regress and should b pears to me unassailable, yet it is not universally accepted amo uniquely defined behavioural postulates." (Elster 1983b, p. 139)
1983	*Sour Grapes*	17-18	Yes	Yes (quoted)	"The demand for an optimal amount of evidence immediately le

1985	"The nature and scope of rational choice explanations" (book chapter)	69	No	Yes	"In most cases it will be equally irrational to spend no time on doing so. In between there is some optimal amount of time th however, is true only in the objective sense that an observer wh the value of gathering information and find the point at which costs. But of course the agent who is groping towards a decisio optimal decision with respect to information-collecting.[23] H costly and that there is a trade-off between collecting informatio off is." *(Elster 1985, p. 69)*
1986	Introduction to the edited book: *Rational Choice*	14, 19	No	No	"It is not possible, however, to give general optimality criteria fo "The non-existence of an optimal amount of evidence arises ... value of the search for information." *(Elster 1986, p. 19)*
1987	*"The possibility of rational politics" (article)*	72-73	No	No	"Indeterminacy of how much evidence of collect, given desire Typically, collection of information has associated with it know certain benefits. To assess the uncertain elements is itself a co have to act, without any illusion that the decision is in any sense
1989	*Solomonic Judgements*	15-16	No	No	"Sometimes it is impossible to estimate the marginal cost and be of battle who does not know the exact disposition of the enemy tially great, cannot be ascertained. Determining the expected form numerical probability estimates concerning the possible ene
1989	*Nuts and Bolts*	35-38	No	Yes (bibliographical essay).	"Deciding how much evidence to collect can be tricky. If the sit we know pretty well the costs and benefits of additional infor urgent, like fighting a battle or helping the victim of a car acci (*Elster 1989b, p. 35)*
1993	"Some unresolved problems …" (article*)*	182-183	No	No	"Suppose than I am about to choose between going to law scho career but of life style. I am attracted to both professions, but I c a lifetime, I might have been able to make an informed choice b make a rational decision." *(Elster 1993, p. 182)*
1999	*Strong Feelings*	144-145 178-179	No	Yes	"Clearly, it will often be irrational not to invest any time in co occasions when there is a danger of gathering too much inform optimal level of search, a 'golden mean.' Whether one can kno which I shall not discuss here [9]" *(Elster 1999, p. 144-145, the ter 1964)*

3.3.2 Distinguishing the two main arguments

First of all, the quotations indicate that there has been a shift in Elster's emphasis. From 1978 to 1983 the argument against the possibility of collecting an optimal amount of information was based on the infinite regress argument. After the important article from 1985, the focus turns to the empirical problems of estimating the value of information when we are in novel situations and when the environment is fast changing. As mentioned in the introduction I have labeled the first argument "the infinite regress problem" and the second "the estimation problem." Both arguments are used by Elster to argue that it is impossible to collect an optimal amount of information.

3.3.3 S.G. Winter and Elster's argument

The infinite regress argument is clearly inspired by S.G. Winter. The key quotation is:

> The "... choice of a profit maximizing information structure itself requires information, and it is not apparent how the aspiring profit maximizer acquires this information, or what guarantees that he does not pay an excessive price for it." *(Winter 1964, p. 262)*

Three things should be noted about this quotation. First, Winter does not use the term "infinite regress," nor does he do so in any of the articles I have read (Winter 1964, Nelson and Winter 1964, Winter 1971, Winter 1975, Winter 1987). In fact, in Winter's article from 1975 the problem is said to be a *potential* "self-reference," not infinite regress. The distinction is important because the existence of self-reference need not involve the problem of infinite regress.[17] Moreover, he explicitly admits that he has not proved that the problem is one of self-reference. As he writes:

> ... the optimization whose scope covers all consideration including its own costs -- sounds like it may involve the logical difficulties of self-reference. To demonstrate this -- to prove logically that there is no superoptimization -- would require the development of a formal framework within which the statement could be interpreted. That would be an interesting project. But, whatever the outcome of that project, it is clear that "optimization" as ostensively defined by pointing at appropriate portions of decision theory literature does <u>not</u> involve self-reference. *(Winter 1975, p. 83)*

Second, the focus in the quotation from 1964 is on how somebody can acquire information about the value of more information. The term "not apparent" indicates some reservation whether the argument really is purely logical (it is impossible) or empirical (it is difficult). Third, there is something odd about the last sentence ("what guarantees that he does not pay an excessive price for it"). Rationality does not demand that we never pay more than the true value of something. The question is whether we were justified in believing that the information was worth the costs *when the decision was made*. It may turn out that the information was less valuable than we believed, but—as Elster argues in *Sour Grapes* (Elster 1983c, pp. 15-19)—it is possible to make rational mistakes.

Thus, an investigation into the sources leads me to question the appropriateness of Elster's reference to Winter when discussing the problem of infinite regress. True, Winter has frequently written about problems around the collection of information, but on the specific question of infinite regress he only mentions that there *may* be a problem of self-reference. He

---

[17] Thanks to Per Ariansen who answered my questions on this and provided me with the following example of a self-referential sentence with no infinite regress problem: "All sentences have a truth value."

does not argue that this is the case. Even if Winter had proved that there was a self-reference, it is not enough since he must also prove that the self-reference involves a vicious infinite regress.[18] Proving infinite regress involves, first, the demonstration that there is a self-reference. Second, showing that the self-reference creates an infinite regress. Third, proving that the infinite regress is "vicious" (that it does not converge). Elster himself does not provide such a three-step argument, and his main reference explicitly admits that he has not proved even the first step (self-reference) necessary to demonstrate "vicious" infinite regress.

3.4 Are the arguments consistent?

There is, at the very least, a tension between Elster's argument on the infinite regress problem and the estimation problem. When discussing the estimation problem, Elster admits that it is sometimes possible to choose what approximates the optimal amount of information:

> Information is useful, but costly to acquire. Ideally, the rational agent would strike a balance between these two considerations: he would acquire information up to the point at which the marginal cost of acquiring information equaled *[sic.]* its expected marginal value. In some areas of decision making these calculations can be carried out with great accuracy. Thus 'To detect intestinal cancer, it has become common to perform a series of six inexpensive tests ('guaiacs') on a person's stool. The benefits of the first two tests are significant. However, when calculations are done for each of the last four tests to determine the costs of detecting a case of cancer (not even curing it), the costs are discovered to be $49 150, $469 534, $4 724 695 and $47 107 214, respectively. To some these calculations suggest that the routine should be reduced, say to a three-guaiac test'."[19] *(Elster 1989a, p. 15-16)*

This is a problem because it is inconsistent to argue that it is logically impossible to collect an optimal amount of information and at the same time argue that the problem is sometimes solved empirically. To what extent is this a problem in Elster's writings?

First of all, I am hesitant about using the label "contradiction" because Elster himself does not explicitly write that the problem of infinite regress represents a "logical problem" in the theory of optimization. On the other hand, consider the following quotations:

> The demand for an optimal amount of evidence *immediately* leads to an infinite regress. *(Elster 1983c, p. 18, my emphasis)*

> … firms **cannot** optimise *ex ante*, since they do not have and **cannot get** the information that would be required. Specifically, they would need an optimal amount information, but this

[18] Academic honesty demands that I admit the following. First, my own argument in chapter four does not provide the "formal framework" that Winter claim is necessary. Second, the mentioned Per Ariansen who provided details on the distinction between self-reference and infinite regress, also argued that the problem of including the cost of optimization in the optimization itself is one of those self-reference problems that leads to a vicious infinite regress (I am not sure I agree with him on this).

[19] Elster quotes from a book by P.T. Menzel (*Medical costs, moral choices*. New Haven, Conn.: Yale University Press, 1983) and he notes that the conclusion is controversial because it depends on how much we believe a life is worth.

> leads to a new optimisation problem and hence into an infinite regress. *(Elster 1983a, p. 6., boldface added)*

> One of his [S.G. Winter] contributions is of particular interest and importance: the demonstration that the neoclassical notion of maximizing involves an infinite regress and should be replaced by that of satisficing. *(Elster 1983b, p. 139)*

As for the strength of these arguments, Elster writes that the argument "appears to me unassailable" (1983b, p. 139). He also thinks that S.G. Winter and himself have provided the sketch of an "impossibility theorem" (1979/1984, p. 59) and that the infinite regress problem represents an argument "against the very *possibility* of planned profit-maximizing" (1982, p. 112, my emphasis)." Hence, the tendency of the argument seems to be that it is logically impossible to choose an optimal amount of information. To the extent this is true, Elster's early argument is in tension with his later emphasis on the empirical nature of the problem i.e. that it is *often difficult* to form a reliable estimate about the net value of information.

4 Evaluating the infinite regress argument

4.1 Infinite regress: An overview

The problem of infinite regress has a strange place within the intellectual history of decision making theory. Several decision theorists mention the problem briefly, but few discuss it in any detail.[20] Some dismiss it as a fruitless problem, others believe it has the power to overturn traditional economics. Some are unsure about whether there is an infinite regress problem at all, other think it is obvious that the problem exist, and some of these believe the problem has an obvious solution.

Among those who mention the problem but dismiss it as fruitless, are Howard Raiffa and Leonard J. Savage:

> People often ask, "How do you know whether or not it is worth the effort to make a formal analysis of the decision problem? Is this a decision problem itself? Can you do a decision analysis of whether it is worth doing a decision analysis?" I don't know anyone who can give definitive answers to these questions, and I suspect one runs into a messy and explosive infinite regression if he tries to incorporate considerations of these questions into the formal structure of a decision-theoretic model. *(Raiffa 1968, p. 266)*

> It might ... be stimulating, and it is certainly more realistic, to think of consideration or calculation as itself an act on which the person must decide. Though I have not explored the latter possibility carefully, I suspect that the attempt to do so leads to a fruitless and endless regression. *(Savage, quoted from Conlisk 1996, p. 687)*[21]

Modern theorists have followed this example, either ignoring the problem or dismissing it briefly. For instance, M.C.W. Janssen (1993, p. 14) writes: "In order to avoid a discussion of the conceptual difficulties related to this infinite regress, I will not be concerned with the information-gathering process in what follows. The analysis starts at the point where agents are assumed to possess some specified amount of information." Another example of a quick dismissal of the problem (but for very different reasons than Raiffa) is Russell Hardin (1988, p.

---

[20] For some examples of works where the problem of infinite regress and costly information & optimization are mentioned (except for Winter and Elster), see: Conlisk (1996, p. 686); Hardin (1988, p. 4); Hodgeson (1994, p. 425); Hodgeson (1997, p. 667); Hoogduin and Snippe (1987, p. 436); Johansen (1977, p. 144); Pingle (1992, p. 8); Raiffa (1968, 266); Resnik (1987, p. 11); Savage (1967, p. 308); Shulman (1997, p. 143).

[21] Savage was less dismissive, but equally brief, in a later article (from 1967): "A person required to risk money on a remote digit of π would, in order to comply fully with the theory [of expected utility] have to compute that digit, though this would really be wasteful if the cost of computation were more than the prize involved. For the postulates of the theory imply that you should behave in accordance with the logical implications of all that you know. Is it possible to improve the theory in this respect, making allowance within it for the cost of thinking, or would that entail a paradox." (Quoted from Hacking 1967, p. 311 and Lipman 1991, p. 1105. The original quota-

4) who—unlike Raiffa—believes "It should embarrass philosophers that they have even taken this objection [the quandary of unending calculation] seriously" since the solution is obvious: "we *satisfice*, we do not maximize." If asked "How do you know that another ten minutes of calculation would not have produced a better choice?" Hardin answer is simply "You do not" and "At some point the quarrel begins to sound adolescent" (Hardin 1988, p. 4). He may be right that the quarrel sounds adolescent (in his ears), but I do not believe name-calling of this type is a very good argument against taking the problem seriously.

Even among those who take the problem seriously, few enter into a detailed discussion. For instance, Leif Johansen has written than:

> The question of how far to go in the direction of perfecting the analyses is in itself an optimization problem, but a peculiar one in that it can itself not be subjected to analysis ... at least in the last instance. Should one try to analyse the question of how to strike an optimal balance between perfection and simplification, then the same question could be raised in relation to *this* question, and so on. At some point a decision must be taken on intuitive grounds. *(Johansen 1977, p. 144)*

In fact, after extensive search I found only four major papers with a detailed discussion of the problem of infinite regress: Mongin and Walliser (1988), Smith (1991), Lipman (1991) and Vassilakis (1992). There is some related work on "beliefs about beliefs" and common knowledge problems (Sargent 1991), but I did not enter into this literature as I restricted myself to one-person decision theory. There is also some relevant work within search theory, for instance Baumol and Quandt (1964). Finally, there is an extensive philosophical literature on the problem of induction and infinite regress. These issues, however, are somewhat peripheral and the four papers above are the only detailed discussions of the core problem. Altogether I therefore have to agree with John Conlisk (1996, p. 687) when he writes that "Given the vast number of expositions of choice theory, it is remarkable how infrequently the regress issue is mentioned ..."

An examination of the four articles reveal that they are discussing slightly different topics as well as taking very different approaches. Elster's focus is on infinite regress as a cause of the impossibility of collecting an optimal amount of information. Smith's focus is not on the collection of information, but on how to decide how to decide. The same is true of Mongin and Walliser, but they approach the question in a much more formal manner than Smith. Lipman also tries to tackle the same question, but he notes that his approach is very different from Mongin and Walliser. In sum, given the lack of papers, the importance of the

---

tion appears in Savage 1967, p. 308, but it appears that a printing mistake has omitted several lines in the quotation.)

problem and the difference in the existing papers, the issue of infinite regress is virtually crying out for a detailed and possibly unified treatment.

Unfortunately, I do not possess the abilities to provide a unifying frame. Of the three possible levels of the infinite regress problem[22] and the various approaches, I have chosen to focus primarily on Elster and secondarily on Winter and Lipman. I am in no position to use Vassilakis' arguments (in an e-mail even Bart Lipman admits that "Vassilakis is an extremely difficult paper to follow"), but I shall occasionally draw upon Mongin & Walliser and Smith.

4.2 Elster’s argument: A visualization

The infinite regress problem is presented as follows by Elster in an article from 1982:

> In order to maximize, say, profits, you need information. As information is costly, it would be inefficient to gather all the potentially available information. One should rather settle for the optimal amount of information. But this means that the original maximization problem has not been solved, only replaced by a new one, that immediately raises the same problem. *(Elster 1982, p. 112)*

Visualized the argument may look like this:

Figure 4.1: Elster's infinite regress argument

```
1. Collect an optimal amount of information (the first maximization prob-
lem)
2. To do (1) we must first collect information about how much information
it would be optimal to collect (the second maximization problem).
3. To do (2) we have to collect an optimal amount of information about how
much information we should collect before we decide how much information to
collect (the third maximization problem).
4. ...
...
```

Since the chain goes on forever, the argument is that the original problem has no rational solution. My question is this: Is it really true that we have to collect information before we decide how much information to collect? Is this not to demand that the agent always should know something that he does not know?

[22] First, deciding how to decide. Second, in belief-formation. Third, in the collection of information. These are clearly related and interact with each other, but conceptually distinct.

4.3 The first counterargument: Do the best you can!

Imagine the following reply to Elster's infinite regress argument as visualized in Figure 4.1: At any point in time you simply have to base your decision on *what you know at that time*. This includes the decision about how much information to gather. Previous experience in making decisions and gathering information may give you some basis for estimating how much information to collect (or it may not, but this is an empirical question). In any case, the rational decision is simply to choose the alternative—act or collect more information—that has the highest expected utility given your beliefs at time zero. The situation could be visualized as in Figure 4.2. At time t=0 you want to make a rational decision about what to do, either to "act now" or to "collect more information."

Figure 4.2: The choice situation

```
   Act
0<                      Act
   Collect Information <          Act
                        Collect <           Act
                                  Collect <
                                            Collect …
```

When the problem is visualized in this way, one infinite regress problem is simply that the branching could go on forever. This, in turn, means that it may be impossible to work out the expected utility of collecting more information, [23] and/or that the value of collecting more information may always be greater than "act."[24] In practice, however, there is little reason to expect an infinite regress problem in the collection of information. Many decisions simply cannot be postponed forever i.e. in the words of Holly Smith (1991) the decision is *non-deferral*. In fact, as he also notes, *all* decisions are non-deferral since all humans eventually die.[25] As long as this is the case, it seems rational simply to start at t=0 and base your choice of whether to collect more information on your beliefs about the net value of collecting more

[23] I use the word "may" because even infinite series may converge (e.g. the St. Petersburg Paradox).
[24] Buridan’s Ass problem: The ass that died standing between two stacks of hay because he could not decide which one it would be best to choose.
[25] This assumes selfishness. If people are altruistic they may seek to do the best they can on behalf of an organization, a cause, or "human mankind." In this case there may be some decisions that can be postponed forever since your death does not mark the end of the utility to be gained from the decision..

information. We avoid the infinite regress since it would not be rational to include the options after your death (or after the time limit) in the calculation.[26]

Second, even if the decision could be postponed forever, the benefits of collecting more information might decrease and as such the problem has a solution in the limit. Of course, the real question is not only whether the problem has a solution in the limit, but whether it is possible for the agent to know this and the precise trade-off (costs and benefits of information) that enables him to make the rational choice about whether to collect more information or not. The answer is to base your decision on the best possible beliefs about the value of more information at time t=0. Should I collect more information? Yes, if my beliefs (based on all my past experiences up to t=0) tell me that more information has a higher expected utility than acting now. Is it possible to estimate the expected net value of more information? The intuitive answer is simply to use your previous experience up to t=0 to form beliefs about how much more information is worth. Of course, this is easier said than done, but sometimes you may compare with similar situations in the past (classical view of probability); sometimes you may use a theory which is developed using past data to predict the value of more information; and, finally, some would argue that it is rational to base your decision simply on your subjective beliefs regarding the value of more information (see Chapter 5, especially sub-chapter 5.3.1, for more on this).

In sum, to get Elster's infinite regress in the collection of information off the ground (if the visualization in Figure 4.2 is correct) we must assume an immortal agent with a zero rate of discounting[27] (or acting on behalf of something which is immortal and not impatient), a decision that can be postponed forever and the impossibility of rationally realizing that the value of information converges. A problem based on these assumptions does not appear very significant in the real world.

4.4 Arguing against the counterargument: Radical skepticism and costly deliberation

A number of counterarguments can be made to the above argument that it is rational to base your choice of whether to collect more information on your current beliefs and that this "solves" the infinite regress problem. Some of these arguments are "internal" to Elster in the

[26] One could argue that economists often assume agents who live infinite lives, and that human mortality is not considered a "proper" counterargument. See footnote 31 and the discussion in the text for more on "proper" arguments against the infinite regress.

[27] Thanks to Olav Bjerkholt who gave suggested that "zero rate of discounting" also was a precondition for infinite regress as visualized in Figure 4.2.

sense that he discusses the issues, and these will be dealt with later (for instance that the beliefs are too weak, that the numerical beliefs do not exist or that the beliefs are biased). Some of the other arguments, however, are "external" i.e. arguments that may make the problem reappear but only in ways that—I shall argue—are not consistent with Elster's arguments. Two of these will be dealt with here. First, the argument that it is never rational to base your beliefs on past experience since it is impossible to "prove" the rationality of induction. Second, there is the problem that we are not logically omnipotent, so even forming the best beliefs entails a cost that we should consider in the optimization process.

The problem of induction and the possibility of radical skepticism was introduced by David Hume and has been extensively discussed among philosophers ever since. Once again we find that an infinite regress is the cause of the problem: We need to justify the method by which we go from the past to the future, and this justification, in turn, needs to be justified and this also needs to be justified and so on (for more on the philosophy of knowledge and the problem of infinite regress, see Nozick 1981, ch. 3, especially pp. 268-280). For instance, in the case of medical diagnostic we may try to justify the inference that the third test costs $49 150 for each case of cancer discovered. The radical skeptic then argues that even if this cost is a true description of the past, you do not know whether it will hold in the future. Trying to defend yourself you claim that the general method behind your statistical method of processing information from the past to predict the future has succeeded quite well in the past for a great many different cases, not only medical diagnosis. But the skeptic counters this by the repeating his argument: Even if the general method has performed well in the past you cannot prove that it will continue to perform well in the future. There is, the skeptic claims, no rational basis for induction.

We may deny the paradox at the very start (denying the regress) or we may say that it is possible to stop the regress at some level (Nozick 1981, pp. 275-280, claims that this is possible). In the current context, however, it is not necessary to go further into the problem of induction in philosophy. It is true that radical skepticism is one possible counterargument to the statement that "it is rational to base you decision on the expectations formed on the basis of your beliefs at t=0", *but the argument is not available to Elster*. He accepts that it is sometimes possible to collect close to an optimal amount of information (as in the example of medical diagnosis). By so doing he rejects radical skepticism and since the current context is an evaluation of Elster's arguments we do not need to spend more time on radical skepticism as a cause of indeterminacy in the rational collection of information.

To discuss the second possible counterargument, we need to follow Smith (1991) and distinguish between deliberation and search. Deliberation is costly because we have limited cognitive abilities and it takes time and effort to work out the best beliefs for a given set of knowledge. Search is the activity of gathering more information. The advice "base your action (on whether to collect more information) on the best possible beliefs you can form given your knowledge at t=0" ignores the costs of deliberation. It may not even be rational to form "the best possible beliefs" once we take into account the cost of deliberation. Moreover, deciding how much deliberation to conduct also raises an infinite regress problem—you have to deliberate on how much deliberation to do on how much deliberation .... Thus, the infinite regress is resurrected.

Once again I will claim that this resurrection does not correspond to Elster's presentation of infinite regress. Elster explicitly writes about *infinite regress in search* and not deliberation.[28] It may be true—as Smith (1991) claims—that most authors who discuss the problem of infinite regress do so in the context of deliberation. I am reasonably sure, however, that he misleads us by quoting Elster to exemplify an author who discusses infinite regress in deliberation. Thus, while there may be an infinite regress in deliberation, this is not the problem Elster focuses on and I do not have to deal with infinite regress in deliberation to argue against him.

One might argue that the problems of infinite regress in deliberation and search are so similar that solving one implies that the other can be solved. I disagree since I found infinite regress in deliberation much more difficult to handle than infinite regress in the collection of information. In short, it makes sense to say that "I'll stop collecting information now since given my current beliefs about the value of more information 'act now' has a higher expected utility than 'collect more information'." When faced with the problem of deliberation, I am less certain. It does not make to sense to say that "I'll stop deliberating now because I know without thinking that the expected costs of more deliberation is larger than the expected gains" because just calculating these costs and gains means that you have done some thinking (and you cannot know that this was worthwhile). Hence, my argument against Elster's views on infinite regress in search does not imply that I dismiss the potential problem of infinite regress in deliberation.

[28] It is possible to argue that Elster includes deliberation as well (Elster 1978, p. 162, quotation includes "cost of acquiring *and evaluating information*", my emphasis), but it is not plausible to interpret this as his main argument given the emphasis on search in the other quotations (Table 3.4).

4.5 An alternative interpretation: Information about information about ...

Figure 4.1 presented one interpretation of Elster's argument. There is, however, however, a second possible interpretation of his argument which may make the problem reappear. Consider the visualization presented in Figure 4.3.[29] Here the problem at t=0 is that the set of possible actions is infinite. The choice is not only between "act now" and "collect more information" since the option "collect more information" is really a general category which includes an infinite set of alternatives. Hence, at t=0, one option is to act right away, another is to collect information directly relevant to the problem; A third option is to collect information about how much information you should collect. Fourth, you may collect information about how much information to collect before you decide how much information you are going to collect. And so we could go on forever.

```
Figure 4.3: The infinite set of alternatives

Possible alternatives at time t=0:
1. Act
2. Collect information
3. Collect information about how much information you should collect
4. Collect information about how much information you should collect to de-
cide how much information to collect.
…
…
```

If the problem is visualized in this way it is less obvious that the non-deferral of decisions can solve the problem. Among all the feasible alternatives at t=0 we want to choose the one that has the highest expected utility. If the set of feasible actions is infinite, then we do not know for certain whether some alternative "far down" would have a higher expected value.

4.6 Trying to argue against the alternative interpretation

One possible counterargument against the infinite regress in Figure 4.3, could be that it is not feasible (given limitations in the human mind) to go very deep. For instance Lipman's "solution" to the infinite regress involves a restriction on the feasible set of computations which in his words "can be viewed as a restriction the complexity of the calculation the agent can carry out" (Lipman 1991, p. 1112). One such restriction is a limitation on how "deep" people can think (Lipman informed me in an e-mail that this is not exactly the type of computational re-

[29] This interpretation is inspired by, but not equal to, Lipman (1991).

striction he uses). Most people are not able to go beyond three or at most four levels of reasoning of the type "I know that you know that he knows ..."[30] Even experts in strategic thinking cannot go further than about seven levels. This might indicate either that this information is not very valuable, or that it is difficult to utilize it given our cognitive limitations, which in turn makes the information less valuable. If we define rationality as "doing the best we can", then considering your own cognitive limitations may "solve" the infinite regress problem. Given your own cognitive limitations it becomes rational not to go very deep! How convincing is this argument?

I am unsure, but I do not think we should allow the argument about limited human cognitive abilities much weight in the current context. The key question in this paper is whether it is possible to collect an optimal amount of information. "Possible" may be interpreted to mean "it is feasible given *unavoidable* constraints." The reason I am reluctant to use limited human cognitive abilities as an argument against the infinite regress argument, is that we may overcome (at least some) of our cognitive weaknesses (i.e. they are not unavoidable). And, as Savage (1967) argues, we want to use the theory of rationality to police our own decisions—as a tool to find the best possible decision and to improve the way we make decisions. Including human limitations makes it too easy to label actions rational, reduces the utility of the theory as a guide to what we should do, and—finally—it seems to me to be a case of mislabeling to argue that human cognitive weaknesses can "solve" a *logical* problem.[31] In sum, I do not want to use this argument against the infinite regress problem.

I should admit that there is a potential weakness in the argument above. While it is true that we may work to avoid some of our cognitive weaknesses, it is vain to believe that we can overcome all. This means that a normative theory of rationality faces a dilemma. It will not give very good advice if it does not take into account human weaknesses ("jump 90 feet" is not very useful advice when it is impossible). On the other hand, if it takes all our weak-

[30] Information about information about information may not be comparable to "I know that you know I know." Imagine the case of buying a house. Most people want to collect information about the house. Some also collect information about what kind of information (and how much?) they should collect (e.g. books about how to collect information before buying a house). We could easily imagine information about this information e.g. a magazine that reviews several books about how to collect information. (But can we find information about how many books to read before determining how much information to collect?) This is three levels deep and it is still not too difficult to imagine. Maybe information about information about information is easier to imagine than "I know that he know that I know?" Be that as it may, we seldom find that the regress goes beyond three or four levels (empirically speaking).

[31] The reader may recall that I previously allowed the argument that humans are mortal as a part of the solution to the infinite regress argument (see section 4.3). Is this inconsistent with the rejection of the solution based on the argument that humans have weak cognitive abilities? I do not think so. The key difference is that mortality is an unavoidable fact of life, while cognitive weaknesses can be (at least to some extent) reduced i.e. they are not unavoidable to the same extent that death is.

nesses into account it is less useful as a tool for improving our decision making abilities; It is the gap between the ideal and reality which tells you that you should become better at jumping or that you should improve you cognitive abilities so you are not fooled by framing and so on. This distinction is also important as an answer to a potential criticism of my definition of rationality (sub-chapter 2.4). One might argue that if it is impossible to collect a rational amount of information (on one definition) then that definition must be incorrect since a correct definition of rationality cannot prescribe something that is impossible. If this is the case it seems strange, as I have done, to first define rational choice and then examine whether it is determinate. No correct definition can be found *before* we know whether the definition prescribes the impossible.

On both accounts—the rejection of cognitive weaknesses as an argument and the inclusion of information collection as a requirement of rational choice—one possible answer is to put emphasis on the role of rationality in telling us how to improve the way we make decisions, not only which alternative to pick from the feasible set. This is, however, one of those arguments on which I will—as promised in the introduction—indicate that I am very uncertain. (See Melberg 1996, p. 474, for an attempt to use the theory of the second-best to examine when ideals should take account of the feasibility of its advice.)

Given the doubts about using cognitive weaknesses to "solve" the infinite regress problem, we might try a second argument against the infinite regress in Figure 4.3. It is the same argument that was used to "solve" the problem as visualized in Figure 4.2. That is, *if the decision is non-deferral* the set of relevant alternatives is also constrained. True, one could always choose to collect information at some very deep level at time t=0, but as long as we know that time is limited the value of doing so is zero since after collecting this information we have to go through all the other levels before we finally make a decision. After collecting information about how much information to collect we have to go out and collect the information. Since this process is time-consuming, time constraints limits the depth of the feasible set than needs to be considered.

Unfortunately, also this argument has a weakness. Collecting information about how much information to collect may in fact reduce the overall time spent collecting information. For example, after collecting information about how much information to collect, you may find that the optimal amount of information one level down is very low, even zero. I am unsure about the implications of this problem for the existence of the infinite regress argument.

4.7 A formalized example

As mentioned in the introduction I have not been able to create a general and formal proof for the non-existence of infinite regress in the collection of information. On the other hand, a general proof is not needed to disprove the argument that infinite regress makes it *logically* impossible to make a rational decision. As philosophers are found of pointing out; one black swan is enough to disprove the argument that all swans are white. Hence, all I need is one example of an optimization problem that can be solved even if we consider the potential causes of infinite regress described in the previous sections. The purpose of this section is thus to try one formal example of optimization in order to see what kind of infinite regress problems we might encounter and whether they can be solved. This sub-chapter is not as closely connected to Elster work as the previous chapters (but note that optimization costs are mentioned in Elster 1987, pp. 72-73), and it is closer to the spirit of the problem presented by S.G. Winter (1975, *not* 1964).

To develop the argument I want to use a very simple model described by Mark Pingle (1992) in his article about how people react when the optimization process is costly. The agents are assumed to derive utility from leisure (s) and food (F):

$$U = U(s, F) \qquad (4.1)$$

The agent wants to maximize this function, but he is limited by a time constraint and a budget constraint. The total time allocated to leisure and work (L) cannot exceed the time endowment (T). The time allocated to work gives income (W pr. unit time) that can be used to buy food at price P and the agent cannot buy more food than his income allows. Formally we have the following constraints:

$$T = s + L \qquad (4.2)$$

$$PF = WL \qquad (4.3)$$

The "rational" choice of leisure vs. work can easily be found by doing the appropriate mathematical operations.[32] However, it seems reasonable to assume that this process itself is time consuming. In other words, the process of optimization is itself costly and we should consider this cost in the optimization problem. To incorporate this cost we use the symbol D to denote the time spend deliberating or searching for more information. The time constraint (4.2) then becomes:

$$T - D = s + L \qquad (4.4)$$

[32] The solution is then given by the two constraints (4.2 and 4.3) and the following condition: $U_1/U_2 = W/P$

Now consider three types of agents:

a) "Hyperrationality:" Those who do the calculations necessary to find the optimum without considering the deliberation/search costs

b) "Myopic rationality:" Those who consider the optimization costs, but treat them as constant

c) "Higher-order rationality:" Those who realize that the optimization costs themselves are variable and try to choose a decision-rule that balances the costs and benefits associated with the various deliberation/search levels.

The situation is visualized in Figure 4.4. The best option in the feasible set is A. It is impossible to get there if you try to calculate since this (by assumption) is costly. It is nevertheless a part of the feasible set since if one had chosen a different decision rule one might end up in A (such as instant randomization). A hyperrational agent could end up either in B or C, depending on whether his choice variable is food consumption or leisure. A myopic agent would do better, ending up in D. A person of "higher-order rationality" would realize that the decision cost is a variable that should not be treated as a fixed constant. Instead he might try to find the optimal deliberation or search time. More precisely, he has to formulate a new decision rule that tells him the extent to which he should deliberate/search before making a decision.[33] The optimal rule should balance the average savings in decision-time against the average cost of worse decisions (as a result of reduced deliberation).

Figure 4.4: Costly optimization (based on Pingle 1992, p. 8

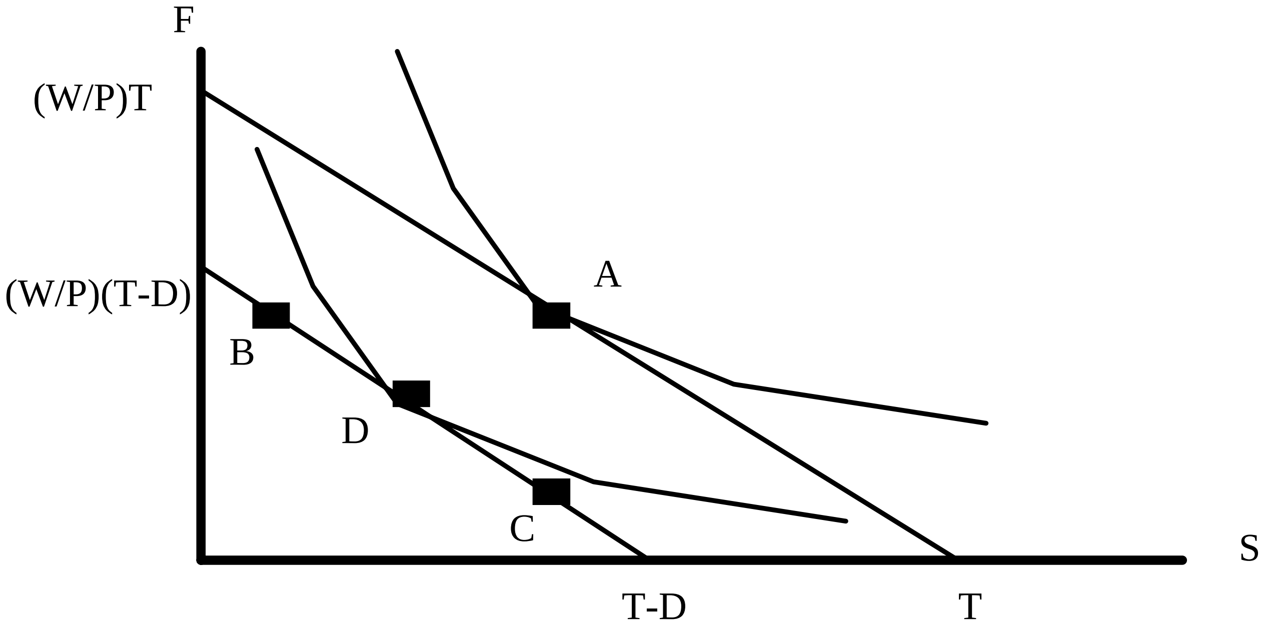

Figure 4.5: Costly optimization, variable deliberation cost

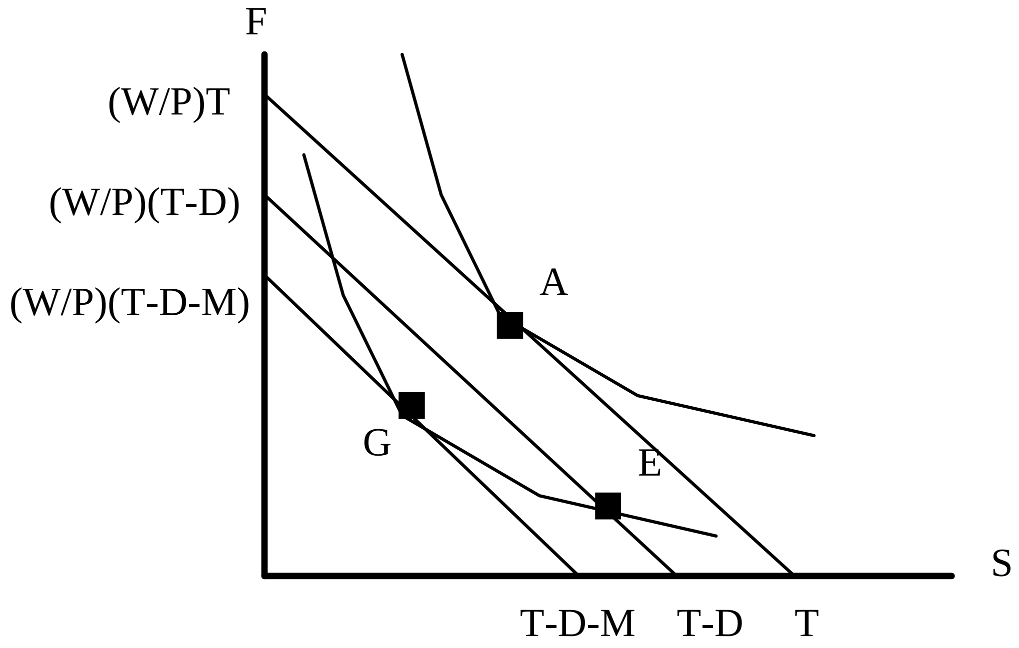

Consider, for instance, a rule that produces the average situation illustrated by point E in Figure 4.5. The same level of utility could have been achieved with even less time available (point G); i.e. the consumer could throw away some time and still reach the same amount of utility. Pingle labels the amount of time that could be thrown away "misuse-costs" (M). Thus, the optimal choice rule would minimize the "cost-of decision" expressed by the following equation:

$$C = D + M \qquad (4.5)$$

In optimum the expected marginal cost in utility of increasing the decision time must be equal to the expected marginal benefits in terms of reduced "misuse costs."

The problem is that the new optimization problem (choosing a decision-rule before you start to apply it) is itself costly. To distinguish the different costs, we might use D_1 to symbolize the costs associated with the process of choosing what decision-rule to apply, D_2 to symbolize the cost of the process when choosing how to choose how to choose what decision rule to use, D_3 and D_n should then be self-explanatory (although very cumbersome to express verbally).[34]

Using this notation one alleged problem of infinite regress is that it is impossible to consider all the costs in an optimization problem. If we want to reach a decision we simply

[33] I use the term "more precisely" because it is in general not optimal to spend a fixed amount of time on search and deliberation. Rather, the optimal decision-rule would involve a sequentially optimal search rule (see chapter 5.5.1)

[34] The same argument could be applied to M (misuse costs) which might be different at different levels. It was not necessary to include this to demonstrate Pingle's argument, so I have not done so.

have to stop the regress at some point and take a decision without considering the cost of the process itself at that level. A person of higher order level rationality (e.g. second order), would face (4.6) instead of (4.5) (now using D_0 instead of D):

$$C = D_0 + D_1 + M \qquad (4.6)$$

Since he ignores D_2 the result cannot be optimal. But trying to include D_2 (a person of "third order rationality") by setting:

$$C = D_0 + D_1 + D_2 + M \qquad (4.7)$$

does not produce optimality because one then ignores D_3. In this way we could continue without ever being able to take account of all costs.

One potential solution would be to argue that the problem converge on a solution. For instance, Pingle argues that the higher level problem is more difficult than the previous level and hence has a higher decision-cost. He writes: "It follows that decision-costs act to limit the extent to which rationality can be displaced to higher levels" (Pingle, 1992, p. 11). I have already argued against using the cognitive limitations of the agents as a part of the solution. However, Pingle's conclusion is also wrong for another reason. If the gain in terms of reduced misuse-costs becomes larger at each level in the hierarchy this gain may outweigh the loss in terms of increased deliberation costs (when the problem becomes more complex). Consider, for instance, learning about learning. Even if the cost of doing so is higher than "just learning", the payoff may be (at least over some interval) very high.

Barton L. Lipman (1991) presents a slightly different interpretation of the infinite regress problems. Return to the original choice of how much time to allocate between leisure and food. The participants in Pingle's experiment were told that their utility function was Cobb-Douglas with the following parameters:

$$U(s, F) = 0.3 \log(s) + 0.7 \log(F) \qquad (4.8)$$

Infinite regress is created by the fact that people are not perfectly calculating robots that require no (or almost no) time to work out the values that would maximize this function. The appropriate way to model this, Lipman argues, is to view the agents as being uncertain about the logical implications of their own knowledge. This means that one option in the set of possible actions is to engage in some activity that reduces the uncertainty about the implications (search or deliberation). Formally, assume O is a set of actions and that your uncertainty about the consequences of the actions in this set can be reduced by various options defined by the set M. The set M and O together constitute the feasible options at t=0. However, as soon as you include M, you must also recognize that you are uncertain about the effectiveness of the activities that are intended to reduce the uncertainty of the alternatives. Moreover, you may take actions to reduce this uncertainty and this set of activities should also be included in the set of possible options. But, and the reader may have gotten the hang of it by now, including a new set of options also requires us to include the activities to reduce the uncertainty of the method to reduce uncertainty of the methods to reduce uncertainty ... and so on. Now the infinite regress is interpreted as our inability to include all the feasible actions in the set of possible options. Slightly more formal, define U as an operator that gives the uncertainty associated with the set of actions (M) and define V as the operator which indicates what we can do about this uncertainty. Then we have the following regress:

$$M, V[U(M)], V[U[V[U(M)]]], V[U[V[U[V[U(M)]]]]], \ldots \qquad (4.9)$$

The question of whether this infinite regress is a problem, depends on whether it is possible to find a set of options (O) that includes all the possible activities the agent (perceives) he could engage in to deal with uncertainty at different levels:

$$O = V[U(O)] \qquad (4.10)$$

In other words, we are looking for a fixed point, a set of options that includes the activities one could engage in to deal with the uncertainty of those options. Lipman presents several possible assumptions that makes a fixed point possible, but he ends up favoring an alternative that works by restricting the complexity of the computation the agent can perform. I have already argued that I am reluctant to use such limitations to make rational choice theory determinate. This should not be interpreted as a criticism of Lipman since my aim in this paper is slightly different from his.[35] He wants to create a model that *predicts* human behavior while I want the concept of rationality to give use advice on how to behave i.e. I use the model prescriptively not predictively. Since real agents have limited cognitive abilities, this seems to be a perfectly acceptable restriction when we want to predict. The very need for an external restriction however, indicates that the concept of rationality alone is not powerful enough to yield a unique prescriptive solution.

The somewhat formal example in this section does not amount to a black swan. It is at most gray since the both Pingle and Lipman introduce "non-rational" restrictions to solve the infinite regress while a genuine black swan would not rely on cognitive weaknesses to solve the infinite regress argument. If it had not been for the fact that this is a thesis in economics, I could have been brave enough to continue the analogy and write that the swan is not even swimming in the right lake. But, since this is economics I shall restrict myself to note that Pingle's example deals with the cost of optimization, not—as Elster—the collection of infor-

[35] Thanks to Barton L. Lipman who made this point clearer to me in an e-mail.

mation. Lipman's example is closer to Elster's concerns about uncertainty and the collection of information, but he deals more with deliberation (working out the implications of your beliefs) than the collection of new information—which is Elster's topic.

4.8 Sub-conclusion

I have so far tried to understand Elster's argument on the impossibility of collecting an optimal amount of information because of the infinite regress problem. My conclusion is that the argument does not demonstrate a significant problem in rational choice theory. Empirically the conditions under which it may arise are very restrictive and I do not think it constitutes a logical proof against the very possibility of choosing an optimal amount of information. I want to note, however, that I have only discussed this problem in the context of how much information to gather. There are at least two other categories of infinite regress problems that I have not discussed. That is, first, to decide how to decide. And, second, to form beliefs using a fixed set of information (for instance, the problems involved in reasoning like "I know that you know that I know ..."). Moreover, I have excluded some possible interpretations that make the problem reappear because they were inconsistent with Elster's interpretation. There may be significant problems here for rational choice theory, but that was not the topic of the section above. Ignoring the other problems I focused on possible interpretations of Elster views and found them—with one exception—that the problem was not significant. The one possible exception is what we should do when gathering information at a deep level also may reduce the optimal amount of information at lower levels.

5 The problem of estimation

5.1 What are the arguments?

It may be impossible to collect an optimal amount of information even if there is no logical problem of infinite regress. For instance, when we are in a unique situation we cannot determine the value of information from historical experience of similar situations, and hence there is (on the classical view of probability) no rational basis for estimating the value of information. I have labeled these problems *the estimation problem* and I have characterized it as Elster's second main argument against the possibility of collecting an optimal amount of information. As argued in chapter three there is a shift towards this line of argument after Elster's article from 1985. In that article, and later, Elster does not use the term "infinite regress" and he does not quote S.G. Winter. Instead, the argument focuses on the problems involved in the formation of probabilities.

Elster's general position is that "beliefs are indeterminate when the evidence is insufficient to justify a judgment about the likelihood of the various outcomes of action. This can happen in two main ways: through uncertainty, especially about the future, and through strategic interaction" (Elster 1989b, p. 33).[36] More specifically the following two quotations illustrate some of the causes of the problem according to Elster:

> Deciding how much evidence to collect can be tricky. If the situation is highly stereotyped, as medical diagnosis is, we know pretty well the costs and benefits of additional information.[9] In situations that are *unique, novel and urgent*, like fighting a battle or helping the victim of a car accident, both costs and benefits are highly uncertain. *(Elster 1989b, p. 35, my emphasis)*

> In many everyday decisions, however, not to speak of military or business decisions, a combination of factors conspire to pull the lower and upper bounds [on how much information it would be rational to collect] apart from one another. The situation is *novel*, so that past experience is of limited help. It is *changing rapidly*, so that information runs the risk of becoming obsolete. If the decision is *urgent* and *important*, one may expect both the benefits and the opportunity costs of information-collecting to be high, but this is not to say that one can estimate the relevant marginal equalities. *(Elster 1985, p. 70, my emphasis)*

To impose some order on the following discussion, I want to make a distinction between three types of probability, three types of problems and three types of implications.[37]

[36] I will not deal with strategic uncertainty in this chapter since I have limited my topic to one person decision-problems.

[37] I should admit that I found it difficult to penetrate Elster's views on probability. First, the discussion is a mix of two different points (Elster himself probably does not confuse them, but he does switch quite fast between the two). The first point, which I agree with, is that there is a tendency for people to spend too many resources trying

On probability, we may follow Elster (1983b, pp. 195-199) and distinguish between the following concepts of probability according to their source: objective probabilities (using relative frequency as source), theoretical probability (the source of the estimate is a theory such as a weather prediction), and subjective probability (degrees of belief as measured by willingness to make bets on the belief).

As for the three problems, I want to make a conceptual distinction between non-existent probabilities, weak (but unbiased) probabilities and biased probabilities. Elster seems to argue that both non-existence and weak probabilities represent indeterminacy, but I believe it is important to distinguish between the two since the question in this chapter is whether it is *impossible* to form beliefs about the value of information.

Finally, I want to separate the following three implications related to the arguments about probabilities. First, the advice that uncertainty makes it rational to use the maximin strategy. Second, that it is intellectually honest to use a strategy of randomization in situations of radical uncertainty. Third, that uncertainty implies that we should not seek more information since it is wasteful to spend resources deciphering the second decimal when we cannot know the first. The paragraphs below elaborate on each of these three distinctions.

In *Explaining Technical Change*, Elster (1983b, p. 185) argues that there are "two forms of uncertainty [risk and ignorance] that differ profoundly in their implications for action. [...]. To this analytical distinction there corresponds a distinction between two criteria for rational choice, which may roughly be expressed as 'maximize expected utility' and 'maximize minimal utility'."[38] More specifically, the argument is that the choice between fossil, nuclear and hydroelectric energy source should be determined not by trying to assign numerical probabilities to the outcomes, but by selecting that alternative which has the least worst consequence (maximin). To justify this principle, Elster appeals to a paper by Arrow and Hurwicz (1972). Hence, one implication of the impossibility of estimating probabilities—

to make the best possible decision (collecting information and so on). The second point is whether it is rational to act on weak beliefs and that it would be no less rational to toss a coin to decide what to do. The two points are both evident in the following quotation: "… I shall argue that the notion of subjective probability is less useful for a theory of rational decision-making than is argued in the Bayesian literature, *and* that it is often more rational to admit ignorance than to strive for a numerical quasi-precision in the measurement of belief" (Elster 1979/1984, p. 128).
Secondly, I was confused by the lack of distinction between weak beliefs and non-existent beliefs. Both are labeled indeterminate, but I found this unhelpful.

[38] "Decisions under risk are present when we can assign numerical probabilities to the various answers to the question 'What will happen?' Decisions under uncertainty imply that we can at most list the possible answers[2], not estimate their probabilities" (Elster 1983b, p. 185). Elster's emphasis on the importance of the distinction between risk and uncertainty, is in contrast to Hirshleifer and Riley (1992, p. 10) who argues that "Knight's distinction … has proved to be a sterile one."

Elster claims—is that we should use maximin instead of trying to form probabilities and then use the principle of maximizing expected utility.

In a different context, the argument is that intellectual honesty implies that we should use a strategy of randomization when we are in situations of ignorance:

> The basic reason for using lotteries to make decisions is honesty. Honesty requires us to recognize the pervasiveness of uncertainty and incommensurability, rather than deny or avoid it. Some decisions are going to be arbitrary and epistemically random no matter what we do, no matter how hard we try to base them on reasons. *(Elster 1989a, p. 121)*

The idea is followed up in a chapter discussing rules about child custody after a divorce in which Elster argues that it may be better to toss a coin than to make an impossible attempt to determine which of the parents will be best for the child.

A third implication of uncertainty, according to Elster, is that it is often wasteful to collect a lot of information: "it is often more rational to admit ignorance than to strive for numerical quasi-precision in the measurement of belief" (Elster 1979/1984, p. 128).

In sum, Elster presents a number of arguments about our inability to form reliable estimates and the implications of this inability (Table 5.1). Probabilities can be non-existent, weak or biased and this implies that it may be rational to use maximin when this is possible, and it is more honest to use randomization (when maximin is impossible) than try to maximize expected utility, and that it is irrational to collect information about the second decimal in a problem when the first decimal is unknown.

Table 5.1: An overview of Elster's arguments about the problem of estimation and their implications[39]

Probability concept	*Problem*	*Cause*	*Implication* [a]	*Justification*	*Example*
Objective/ Theoretical	Non-existent probabilities	Brute and strategic uncertainty	Maximin [b]	Arrow and Hurwicz proof (Best end result?).	Choice between fossile, nuclear and hydroelectric energy
Objective/ Subjective/ Theoretical	Weak probabilities	Brute and strategic uncertainty	Randomization/ Maximin	Intellectual honesty	Choice of career (forester or lawyer)
Subjective	Biased probabilities	Hot and cold mechanisms	Randomization?	Better end result	Investment choices?

(a) Implication for all: We should not waste time seeking information when such information is impossible to find or only weakly significant.[40]
(b) Assuming we know the best/worst possible outcome.

[39] Elster's discussion is in no way as "neat" as the table suggests. As mentioned in the introduction, the "estimation problem" is a residual category and the table represents *my* attempt to create some order in a rather unordered set of arguments.
[40] This implication is redundant in the sense that there is no need for advice beyond the already mentioned rule: Collect information as long as the expected benefits of doing so exceeds the costs.

I have chosen to discuss the validity of Elster’s arguments under four headings. First, how strong is the argument about the non-existence of probabilities (which involves a discussion of subjective and objective probability)? Second, how sound is the argument that randomization is preferable (since it is more honest) in situations of weak probabilities? Third, what is the relevance of biased probabilities to the indeterminacy of rational choice? Fourth, to what extent is it true that our knowledge is limited in a way that is comparable to "being lost in a forest" and does this make it impossible to conduct a rational search? (The answer to this question involves a digression into the theory of search in economics). Within these four headings I want to discuss both the validity of the arguments in isolation, and their consistency with Elster’s other arguments.

5.2 Does radical uncertainty exist?

The principle of maximization of expected utility (MEU) presuppose that the agent has or can form probabilities about the possible consequences of an action.[41] Hence, if it can be shown that these probabilities do not exist, it implies that MEU cannot be used in that situation. This means, As Elster argues, that uniqueness, novelty and rapidly changing environments are problematic for expected utility theory because we cannot use previous experience of similar situations to estimate the relevant probabilities. One possible counterargument is that Elster’s arguments about uniqueness and non-existence of probabilities are heavily dependent on the classical view of probability as relative frequency. If, for instance, we use the concept of theoretical probability it is perfectly possible to get reasonable estimates even from unique combinations of weather observations. Another, and in this context more significant counterargument, is the view that probabilities should be interpreted as measures of subjective uncertainty, in which case it is perfectly possible to speak about probability even in unique situations.

[41] Historically the MEU principle required "decision weights"—it was only later that it was established that these decision-weights should be interpreted as probabilities (Bernoulli).

5.2.1 *Subjective probabilities*

Elster, of course, is aware of this alternative view of probability, but he argues against the use of subjective probabilities.[42] The arguments are (rather crudely), summarized in the following list:

1. It denies the possibility of genuine uncertainty (Elster 1983c, pp. 19-20)
2. It leads to logical inconsistencies because it "may not be possible to make a non-arbitrary individuation of states" (Elster, 1987, p. 73 [note 13]; Elster 1983b, p. 197).
3. "It presupposes that we are somehow able to integrate our various fragments of knowledge and arrive at a stable and intersubjectively valid conclusion" (Elster 1983b, p. 199)

On (1) and (2)

Does subjective probability deny genuine uncertainty? Bayesians argue that it is always possible to translate my uncertainty into probability statements about the world. You simply elicit the subjective probabilities by forcing a person to choose between a given set of alternatives. For instance, suppose you had to choose between the following alternatives (the example is inspired by Elster 1979/1984, p. 129):

> A: If you correctly guess the twenty-first decimal of π you get $100, if you are wrong you get nothing.
> B: If you draw a red ball from an urn of (*p**100) percent red balls and ([1-*p*]*100) percent blue balls you get $100.

If the person prefers A to B one might infer that the person's subjective probability of being able to guess the decimal, is higher than p. One might then increase the percentage of red balls in alternative B and make the choose between A and B once again. If we continue this process we will eventually come to a point where the agent prefers B to A (or end up with the conclusion that the agent is certain that he can estimate the twenty-first decimal of π).

I am not convinced by this argument for the non-existence on genuine uncertainty. First, it seems to deny (by assumption) the very question we want to examine; we do not allow the agent to respond "I don't know!" Second, it assumes that the answer reveals what we want it to reveal. The inference that the choice reflects our subjective uncertainty is only valid if the agent really tries to maximize his expected utility when faced with the two alternatives. If the agent instead simply selects his answer at random (or using some other criteria), then the inference from his answer to his subjective probability is not valid.

[42] "... subjective probabilities are often unreliable and should not be used as a basis for action" (Elster 1983b, p. 197).

A Bayesian might argue that the problem could be solved by saying that total ignorance ("I don't know" in the example above) can be translated into the probability statement that "all outcomes are equally likely to happen." I find this an attractive proposal, but this is both conceptually and logically problematic. Conceptually, as Iversen (1984, p. 61) admits, "saying that each value is equally likely is to say something about the parameter and represents one step up from complete ignorance." As for the logical problem, imagine that you have to guess the value of X, and all you know is that X is somewhere between 0 and 5 (the example is from Iversen 1984, p. 61). If you use the assumption that complete ignorance means that all outcomes between 0 and 5 are equally likely, then the probability that X is less than 2.5 is 0.5:

$$P\,(X < 2.5) = 0.5 \qquad (5.1)$$

But, if you are ignorant about the value of X, you are also ignorant about the value of X^2. The possible range of X^2 is from 0 to 25 (since X goes from 0 to 5). This means that the probability that X^2 is less than 12.5 should be 0.5 (being ignorant about the value of X^2 we simply say that all outcomes between 0 and 25 are equally likely). In other words:

$$P\,(X^2 < 12.5) = 0.5 \qquad (5.2)$$

By taking the square root of both sides of the inequality above, we get:

$$P\,(X < 3.54) = 0.5 \qquad (5.3)$$

But this clearly contradicts the first statement that $P(X < 2.5) = 0.5$.

I am not sure how to respond to this problem. The general difficulty it tries to exemplify is that complete ignorance gives different results depending on the specification of the set of possible events. But this does not show or prove that complete ignorance is something that really exists. A Bayesian might "solve" the problem of inconsistency by arguing that also the specification of possible outcomes belong to the subjective realm. That is, we must simply use the states we believe are possible in the calculation and the proof that this is inconsistent compared to the results using a different set of possible states (more states) is not relevant (or does not prove irrationality). A person cannot be blamed for not using a set of outcomes he believed did not exist (given that this belief itself was rational). In line with this many statisticians and decision-theorists appeal to the principle of insufficient reason to argue that "complete ignorance" does not exist. If you know nothing, they argue, then all the options you believe are possible have the same probability (see, for instance Bhattacharyya and Johnson 1977, p. 67).

On (3)

The final argument is that subjective probabilities are not intersubjectively valid. I am unsure about what this means, but one interpretation could be that people given the same information might come up with different probabilities and it sounds wrong to argue that both are equally valid as a basis for calculating what they should do. (The underlying argument seems to be that two different estimates cannot both be equally rational since there is only one truth). A Bayesian could make several responses. First, Bayesian and classical estimates may converge over time even if people have different initial priors (People starting with different beliefs about the amount of red and blue balls in an urn will revise their beliefs as the are allowed to see the color of selected balls using Bayes' rule). This means that there is a degree of intersubjective validity even for Bayesian estimates.

Second, given the differences in background knowledge it is perfectly possible that two rational people come up with different probability estimates. People will differ in their background knowledge because they have encountered different information in their lives and this is reflected in their prior beliefs. Rational updating based on the same new information may then result in two different beliefs, but none need be more rational than the other (one is certainty closer to the truth than the other, but that is not the point; beliefs do not have to be true to be rational).

The second point also reveals a tension in Elster’s argument. He demands that probabilities should be intersubjectively valid, but he also insists that rationality is a subjective notion. Consider the following quotation:

> It is not possible, however, to give general optimality criteria for the gathering of information. One frequently made proposal—to collect information up to the point where expected marginal value of more evidence equals marginal cost of collecting it—fails because *it does not respect the subjective character of rational choice. (Elster 1986, p. 14, my emphasis)*

The argument here is that an outside observer might be able to asses the value of information, but this does not help the person who tries to act rationally as long as he cannot estimate the value of information. The information has to be available to the person who is making the decisions. This is true, but it also suggests that probability is an inherently subjective notion. As argued, different persons have different information and as such it is possible that they both rationally estimate probabilities that differ. To demand that probabilities be intersubjectively valid (if one by this means that everybody should arrive at the same estimate), is to impose an objective standard on something that is inherently subjective.

A third reply to the argument that subjective probabilities are not "intersubjectively valid" is that objective probabilities are no more intersubjectively valid than subjective probabilities. This is because there is no neutral criterion that determines which cases are "similar enough" to be used as a basis for calculating the objective probability. Some might argue that it was impossible to estimate the probability that the USSR would collapse (no similar events to use as a basis for calculation), others might argue history provided cases of similar empires that could be used to work out the probability of collapse (the Habsburg Empire, the Ottoman empire and even the Roman empire might be used). Or, to use an example from Elster: "The doctor carrying out a medical diagnosis finds himself *many* times in the *same situation*" while "most people are unemployed only once, or, if more than once, under widely differing circumstance." (Elster 1989a, p. 16, emphasis in the original). For this argument to be "intersubjectively valid" we need a criterion of "sameness" and "different circumstances" and there is no such neutral criterion.

5.2.2 Sub-conclusion on the existence of genuine uncertainty

Elster's argument about the non-existence of probability depends quite heavily on the classical view on probability as relative frequency. The argument in favor of this view, and against the subjective view, is (at least) open to discussion. Beyond this I have no strong conclusions on whether the non-existence of probabilities is a serious problem. I tend to believe (rather weakly) that there is often some aspect of the problem that allows us to make some inferences on probabilities. For instance, in the mentioned problem about betting on a decimal of π I would certainly choose alternative A (see p. 50) as long as the percentage of red balls was below 10 since there are only ten decimals to choose from. In many cases it also seems reasonable to translate "I don't know" into "all alternative are equally likely." Yet, I am also aware of the problems with the other proposals and this is the reason for my guarded conclusion.

5.3 Is weak probability a good argument for randomization?

First of all, we must ask in what sense probabilities are weak. Since I want to distinguish between bias and weakness, I shall reserve the label weak for beliefs that are unbiased. Conceptually the distinction is important. For instance, we may form a belief about the color of the balls in a large urn based on a sample of three balls (say we know that the balls are either blue or yellow, but that we do not know how many are blue and how many are yellow). This belief is not very strong, but—if the proper statistical formulas are applied—it is not biased.[43]

As mentioned Elster argues that some beliefs are too weak to justify inclusion in a rational calculation of net expected utility (and that we for this reason should refrain from choosing actions based on such calculations).

> In my ignorance about the first decimal—whether my life will go better as a lawyer or as a forester—I look to the second decimal. Perhaps I opt for law school because that will make it easier for me to visit my parents on weekends. This way of deciding is as good as any—but it is not one that can be underwritten by rational choice as superior to, say, just tossing a coin. *(Elster 1989a, p. 10)*

I think the argument is weak. Assume you have to choose between the following two alternatives:

> Lottery A: $10 000 with an estimated probability of 50.01 (and 0 with probability 49.99)
> Lottery B: $10 000 with an estimated probability of 49.99 (and 0 with probability 50.01)

I would choose lottery A even if the probability estimate is weak (based on little evidence). True, my choice is not accompanied with great conviction that A is much better than B, but why toss coins as long as I have an option that gives a higher expected payoff? Elster might reply that this choice is an example of hyperrationality ("defined as *the failure to recognize the failure* of rational-choice theory to yield unique prescriptions or predictions." [Elster 1989a, p.17]). I agree that it would be irrational to spend much time and money trying to estimate the second decimal if we were ignorant about the first in the case above, but that is not the question. We do not ask whether it is profitable to collect more information, but which choice you should make for a given set of information.

One might argue that the difference is small in the example above, but the true comparison is not simply between the difference in probability. It is the difference in expected utility when the probabilities are multiplied by the payoffs that are important. In the case above the difference is $200, which is a significant sum. The larger the payoff the more significant the small difference in probability is. This argument reveals a tension in Elster's view. In the quotations at the beginning of this chapter he argued that both weak probabilities and large payoffs (importance) pull in the direction of "coin-tossing", but the factors (at least in my example) pull in separate directions.

There is, however, an even more serious problem with Elster's suggestion. In the real world we will encounter *many* choices in which we may rely on probabilities of *varying* reli-

[43] See Gärdenfors and Sahlin (1988, part IV) for alternative conceptions and implications of weak (or unreliable) probabilities.

ability. Sometimes we are very uncertain, sometimes we are more certain. Given this fact, let us compare the following two rules for choosing what to do (decision-guides):

> *The Elster rule*: If your beliefs are very weak, you should (or weaker: might as well) toss a coin to decide the matter; if the beliefs are reliable, you should choose the alternative with the highest expected utility (Elster's strategy).
>
> *The expected utility rule*: Choose the action with the highest expected utility both in situations with weak and strong beliefs (Bayesian strategy).

The fact that we have to make many choices means that the many small differences become large in aggregate. As a Bayesian says in response to why we should choose the expected utility rule:

> ... life is full of uncertainties—in a given week, you may buy insurance, bet on a football game, make a guess on an exam question, and so forth. As you add up the uncertainties of the events, the law of large numbers come into play, and the expected value determine your long-run gains. *(Gelman 1998, p. 168)*.[44]

Another problem with Elster's decision-rule, is the fact that before we make a decision we have to determine whether the situation is one of enough certainty to choose the action that maximizes expected utility, or whether we are so uncertain that we should randomize (or something else, like maximin). Where is the limit, and is it not costly to examine the circumstances in this way every time we have to make a decision?

Of course, we could go all the way and say that all our knowledge is always so weak that we always should toss coins. In this way we could avoid the problem of choosing when using Elster's strategy. Sometimes Elster is attracted to this argument, but at other times he seems to want to "have the cake and eat it." For instance, he is sympathetic to Descartes when he claims that our knowledge is limited in a way that can be compared to being completely lost in a forest (see chapter II.4 in *Ulysses and the Sirens*). Yet, when discussing child custody after a divorce he does not want to go all the way and argue that it might as well *always* be decided using randomization. In some "obvious" cases the court should not toss a coin (Elster 1989a, p. 170-171).[45] But then the court first has to examine whether the case is obvious and this process is costly in the same way (but maybe not to the same extent) that a trial about child-custody would be. In short, either Elster's decision-rule has a problem in terms of decid-

[44] One might claim that the argument is invalid as an argument against Elster because it "changes the goalposts." That is, Elster is discussing the rationality of something in one particular context and to broaden the context (make the principle apply to all decisions) is to apply the theory to something it was not intended to cover.

[45] Among the obvious disqualifying reasons are—according to Elster (1989a, p. 170-171) "physical neglect, physical abuse, sexual abuse, psychic disorders."

ing when to toss a coin, or one has to believe that we are so lost that we might as well always toss coins.

5.4 What is the relevance of biased probabilities?

When discussing subjective beliefs (and beliefs in general) Elster presents convincing arguments to the effect that beliefs often are formed by hot (beliefs influence by what you want to be the case) and cold cognitive mechanisms (wrong beliefs even when you do not have any strong preferences about the truth). The argument is also used when discussing the problems involved in collecting an optimal amount of information. For instance, he argues that the elicitation of subjective beliefs is subject to a mechanism called anchoring; If we start from a low probability (few red balls) in the example of eliciting subjective probabilities, the agent is more likely to end up with a low subjective probability than if we start from a high probability and go down (many red balls). The procedure for measuring the belief affects the belief we find! Surely this is a sign that these subjective probabilities are unreliable and should not be used as inputs in decision-making.

Although I find the topic of hot and cold belief-formation both interesting and important, it is not relevant in the present context. The main question in this paper is whether the principle of rationality yields a determinate answer, not whether peoples' actual behavior conform to the standards of rationality.

We could, however, make a comment about Elster's arguments that applies to all the previous situations and the recommendation that agents should use maximin or randomization in situations of great uncertainty. The problem is that that this prescription (toss coins when you are very unsure) is itself subject to the problem it is meant to avoid. Since Elster admits that we sometimes have reliable probabilities, it follows that we have to decide whether to use maximin/randomization or maximize expected utility. If the argument against the use of MEU is that we cannot rely on our subjective probabilities because we tend to deceive ourselves, then one might also suspect that we deceive ourselves when making the choice about which procedure to use. To say that we sometimes should use maximin because we are biased is not very helpful if the same bias makes us exaggerate the reliability of the probabilities so that we will not choose maximin. This is another instance of the problem, already mentioned, of what happens when you do not go all the way to say that we should always use the maximin strategy.[46]

5.5 How much do you have to know to conduct an optimal search?

Elster often makes the argument that we know too little to conduct a rational search. The argument is accompanied with the following quotation from Leif Johansen—for instance in *Sour Grapes* (p. 18) and Elster (1983a, p. 5).[47]

> It is like going in a big forest to pick mushrooms. One may explore the possibilities in a certain limited region, but at some point one must stop the explorations and start picking because further explorations as to the possibilities of finding more and better mushrooms by walking a little bit further would defeat the purpose of the hike. One must decide to stop the explorations on an intuitive basis, i.e. without actually investigating whether further exploration would have yielded better results. *(Johansen, 1977, p. 144)*

In other words of Elster (1983a, pp. 5-6) the argument shows that "there is a process of search that stops when one has found something that is good enough, or 'satisfactory', without it being assumed that what has been found is in any way 'optimal'." To examine the validity of this argument, I made a few investigations into search theory in economics. The literature starts with Stigler (1961). Another landmark is Rotschild (1974). On a more negative note, Hey (1981) presents a very good critical overview of the field. The key underlying question is this: How much do we have to know to conduct and optimal search (and what kind of search is optimal)?

5.5.1 Stigler: Assuming a known distribution

Stigler (1961) begins by noting that price dispersion is an empirical fact. This, in turn, makes it profitable to search for the lowest price before you buy a product. However searching is costly, so the question is whether there is an optimal amount of search. Stigler shows that if we assume the distribution of prices is known, then it is possible to conduct an optimal search. Assume, for instance, that we know half the stores charge $1 and the other half charge $2 for

[46] It is possible that the problem is less acute when one person is set to design an institution that applies to others—such as a politician suggesting (and deciding) that the courts should toss coins when deciding child-custody cases after a divorce. It may also be that we as individuals have the ability to precommit ourselves using various devices to avoid the problem.

[47] There are several (but relatively minor) inaccuracies in Elster's quotations compared to the original. For instance, the original term "hike" has been replaced with "ride" in Elster 1983c and "outing" in Elster 1983a. Moreover, "little bit" in Johansen becomes "little" in Elster's quotation. Finally "explorations" is replaced with "exploration."

the same product.[48] This implies that the expected price if we visit one store is $1.50. If we sample two stores, the probability of finding one with the lowest price is one minus the probability of sampling two stores with high prices:

Probability of finding at least one store with the minimum price = 1 - (½ * ½)= 0.75

Thus, the expected price if we try two stores is $1.25 [(1 * 0.75) + (2 * 0.25)] and the expected gain of visiting one more store (going from one to two) is $0.25. In the same way we can calculate the expected price after sampling n stores and the gain from going from n-1 to n:

Table 5.2 Expected price after sampling n stores

Number of stores visited (n)	Probability of finding at least one store with the lowest price	Probability of n stores with highest price	Expected price	Expected gain from increasing sample from n-1 to n
1	0.5	0.5	1.5	
2	0.75	0.25	1.25	0.25
3	0.875	0.125	1.125	0.125
4	0.9375	0.0625	1.0625	0.0625
...	...	...	...	...

Based on this one might believe that the choice of an optimal amount of information is simple. The collection of information is optimal when the expected gain is equal (or as close as possible) to the expected cost of further collection. Assuming that the cost of going to one more store (or taking one more phone call or whatever it is that brings information) is known, then this is a simple exercise. In our case, if the cost of increasing the sample size by one is 0.11 (and constant for all n), the optimal search is n=3.

Unfortunately, things are not that simple. There are two main problems. First, the claimed optimal search rule is not optimal at all. Second, to make the example work we had to make some quite strong assumptions, such as the one about knowing the distribution of prices.

Here is an illustration of the first problem: Assume that you find the lowest price in the first store. It is then pointless to search more stores even if the "optimal rule" told you to sample three stores. One might try to avoid this problem by specifying the following rule: "Col-

[48] The existence of different prices clearly implies that there is not perfect competition. Since Stigler does not tell us exactly how the market works, one might argue that the example is unsatisfactory (as Olav Bjerkholt has noted). On the other hand, it is possible to give a plausible story as to how these two prices could exist at the same time—for instance when some stores cater for tourists and some for more knowledgeable locals.

lect n price quotations, but stop searching before n if you find the lowest price." This, however, does not solve our problem. Imagine, for instance, that you decide to collect three price quotations (which was the optimal size of sample), but you happen to draw three stores with the highest price. Should you then resign and buy the commodity at the highest price? In fact, after conducting three searches and having found three high prices, it is still optimal to conduct new searches. Stigler's "fixed sample rule" is not credible since it is not optimal to stop after collecting the "optimal" number of price quotations if these quotations are all high. The expected gain of collecting one more after collecting n high prices is always 0.5 which is higher than the cost of searching. This is a general problem with "fixed sample" search rules, and, to be fair, Stigler is aware of the problem but states that he "must leave for others to explore" the optimal sequential procedure (Stigler 1961, p. 219).[49]

The second problem worth mentioning is the rather stringent assumptions needed to make the example work. Take, for instance, the assumption that the individual knows the probability distribution of prices. First, it seems strange to assume that you know the distribution but do not know anything about which stores are most likely to have the lowest price. Second, we often—Johansen's mushroom problem is one example—do not have a good idea of the true distribution so we have to ask whether it is possible to conduct an optimal search when the distribution is unknown. The first of these is not really a problem since it is at least logically possible to know the general shape of the distribution without having specific knowledge of where to find the lowest price. The second problem has been partially answered by Rotschild (1974).

5.5.2 Rotschild: Optimal search when the distribution is unknown

Unlike Stigler who specified fixed sample search rules for a known distribution, Rotschild (1974) tries to construct and examine the optimal sequential search rule for an unknown distribution. The main purpose of the article is to compare the following five properties of different optimal search-rules (different in the sense that some assume a known distribution and others do not):

1. Does it (the search rule) imply a well-behaved demand function (demand is non-increasing in prices)?
2. Is search behavior a function of the cost of search and the distribution of prices?

[49] In fact, Hey (1981, p. 48) notes that H. Simon had published a formal solution to the sequential problem in 1955, six years before Stigler's article.

3. Does the amount of search decrease when the cost of search increase?
4. Does total cost decrease when prices become more dispersed?
5. Does increased price dispersion also increase search?

I can make no claim to fully understand the proof of the theorems behind his conclusion that in his case "optimal-search rules from unknown distributions have the same qualitative properties as optimal rules from known distributions" (Rotschild 1974, p. 694). I did, however, understand that his views on both the existence of an optimal search rule and its properties rested on some very strict assumptions. The general idea is that an agent starts with his initial beliefs and updates these according to Bayes' rule as he receives new information. Based on these beliefs, in turn, the agent decides whether it pays to go on searching. To make this procedure work, Rotschild must assume (and he is very honest about this) that the prior beliefs have a particular distribution (a Dirichlet distribution—the multinominal equivalent of the Beta distribution), that the agent knows all the possible outcomes (but not their probabilities) *and* that the learning process is localized in the sense that observing a price of 10 does not affect my probability of observing a price of 11 or 1000. This is strange because one might believe that observing several prices in one neighborhood also increases the probability of the other prices in that neighborhood (and reduced the probability of prices far away). On the other hand, one cannot conclude from the statement that "Rotschild's proofs need these assumptions" to "when these assumptions are not met no optimal search rule exist and/or it does not have the same properties as search rules for known distributions." The fact that he did not manage to establish a more general proof does not mean that such a proof does not exist (and he makes some comments to the effect that he considers it likely that the results are more general).

5.5.3 Hey's argument: The impossibility of optimal search?

Rotschild's assumption of localized learning exemplifies a general problem with search theory. In the words of Hey (1981, p. 56): "… in order for the Bayesian approach to work, the searcher must either correctly specify the true environment, or adopt a sufficiently minimal assumption about the true environment." For instance, assume you initially assign the probability zero to an event. For the sake of argument, assume that you initially believe that there is a zero probability of the event that the burglar in a case has blond hair (P [B]=0).[50] Assume, next, that a witness says the burglar in fact had blond hair (define this as A). You then want to work out the probability of the burglar being blond given the information from the witness, $P(B|A)$. Using Bayes' rule this is (- here symbolizes negation [complement]):

$$P(B|A) = P(A|B)P(B) / [P(A|B)P(B) + [P(A|-B)P(-B) \qquad (5.4)$$

But since P(B) = 0 this expression will always be zero; Regardless of how many witnesses tell you that the burglar was blond, your initial certain belief that he was not blond prevents you from placing any weight of these witnesses.

Of course, one could always argue that the example is artificial since people would not attach zero probability to events such as a burglar being blond. In our example this is obvious, but in the more general case it would be foolish to rely on the assumption that the true theory, distribution or event are always within your initial specification of the problem. This leads Hey to explore what he calls reasonable search rules. These are rules about search that—unlike many optimal search rules—are as robust as possible given the possibility of mistakes. It turns out, for instance, that often "sequential rules are considerably less robust than fixed-sample rules…" (Hey 1981, p. 63). This suggests that following a rule of thumb may be a reasonable rule. It should be noted that these suggestions are mainly conjectural, and they are not presented as firm conclusions by Hey.

5.5.4 Sub-conclusion

What are the implications of this short detour into the economic theory of search for the problem of mushroom picking, and the more general argument that it is often impossible to conduct an optimal search? The theory shows that we must assume something in order to do an optimal search. The assumptions need not be very strong. It is enough—for instance—that we specify a very general family of distributions that has the true distributions as a special case. If one by "complete ignorance" means that one knows absolutely nothing relevant to the problem, then this means that an optimal search cannot be conducted. One might argue, however, that "complete ignorance" is difficult to conceive (does it exist at all?) since just asking a question reveals some kind of knowledge about the phenomenon.

[50] The example is based on Melberg (1996, p. 487), which in turn is an extension of Hargreaves Heap (1992, p. 295)

6 Implications

6.1 Introduction

When discussing the problems of collecting information, Jon Elster frequently argues that the argument demonstrates that economists are wrong in focusing on maximization as opposed to theories of satisficing:[51]

> One of his [S. Winter] contributions is of particular interest and importance: the demonstration that the neoclassical notion of maximizing involves an infinite regress and should be replaced by that of satisficing. The argument appears to me unassailable, yet it is not universally accepted among economists, no doubt because it does not lead to uniquely defined behavioural postulates. *(Elster 1983b, p. 139)*

> The Nelson-Winter attack on optimality is therefore a two-pronged one. The argument from satisficing is that firms cannot optimise *ex ante*, since they do not have and cannot get the information that would be required. Specifically, they would need an optimal amount information, but this leads to a new optimisation problem and hence into an infinite regress. On the other hand, we cannot expect firms to maximise *ex post*, since the elimination of the unfit does not operate with the same, speed and accuracy as it does in natural selection. Taken together, these two arguments strike at the root of neo-classical orthodoxy." *(Elster 1983a, p. 6)*

When Winter himself elaborates on the implications of his arguments, he makes a more general claim on behalf of behavioral economics, as opposed to the specific theory of satisficing. Economists, he believes, should focus more on actual behavior as opposed to studying the theoretical equilibrium implications of rationality under conditions of perfect competition. In short, there should be "more explicit reliance on empirically grounded behavioral generalizations at the foundation of economic theory" (Winter 1987, p. 244). For instance, economists have struggled with the question of stability in general equilibrium since the traditional restrictions are not enough to generate well-behaved excess demand functions and convergence to equilibrium (see Fisher 1987 and Kirman 1992). More behavioral assumptions about the agent may reduce this problem.

Thus, the claim is that the infinite regress argument makes it less legitimate to use the assumption of rationality, and justifies an interest in behavioral economics. This claim is often followed by a critical discussion of the argument that natural selection justifies the use of the rationality assumption since only those who maximize survive. As Elster (1989b, p. 81) puts

[51] Satisficing means acting when you find an alternative that is "good enough." The concept was invented by H. Simon (see e.g. Simon 1979 and Simon 1987). The idea was that it is too difficult for agents to follow optimal strategies, so to explain behaviour one has to rely on a different choice mechanism than optimization.

it: "selection processes [in social science] work too slowly to produce behavior that is optimally adapted to a rapidly changing environment."

In short, two claims are made by Elster and Winter about the implications of the impossibility of consciously collecting an optimal amount of information:

(1) Infinite regress and the problem of estimation justify theories of satisficing and the approach of behavioral economics.

(2) The natural selection argument against (1) works only under implausible conditions.

In this chapter, I will examine these views more closely. My arguments are:

1. It follows from the previous chapter that I do not believe the infinite regress problem in the collection of information is a good argument on behalf of behavioral economics. This does not mean that I disregard the conclusion. First, the problems of estimation combined with inherent cognitive limitations also justifies a more behavioral approach. Second, there are problems of infinite regress other than in the collection of information and in rational choice theory in general, that provide further justification for exploring supplements and alternatives to rational choice.
2. A closer specification of the arguments involved in the natural selection argument reveals some possible weak points in Elster's and Winter's views. For instance, some features may be adaptive precisely because the environment is fast changing and the "as-if" argument can be used without relying on the analogy with natural selection. I will, however, not go as far as saying that these problems vindicate the "as-if" argument.
3. There are many other factors involved in the debate over behavioral economics, so the above arguments should only be viewed as one of many. As already mentioned several times, I do *not* claim to have answered the general question of behavioral vs. traditional economics in any way.

I have chosen to anchor these arguments within a broad discussion on the nature of economics and economic methodology since this is what they are i.e. discussions about what economists should do and how they should do it.

6.2 What is economics?

Let me start with a very basic question: What is economics? There are many possible answers, ranging from the enjoyable —"Economics is what economists do" (Jacob Viner cited in Barber 1997, p. 88); the frequently quoted—"Economics is the science which studies human behavior as a relationship between ends and scarce means that have alternative uses" (Lionel Robbins quoted in Stigler 1984, p. 301); And finally the very strongly held conviction exemplified by Gary Becker (1986, p. 110): "The combined assumptions of maximizing behavior, market equilibrium, and stable preferences, used relentlessly and unflinchingly, form the heart of the economic approach…"

A short survey of the field, however, reveals the following two distinctions. First, there is disagreement on what questions economists should try to answer (What does the subject matter include—all behavior or only a subset of "economic" behavior?).[52] Second, there is disagreement about how one should go about answering the questions (Are only rational choice explanations acceptable or are we also willing to include norm-based or psychological explanations). Cross-tabulating these distinctions we have the following (Figure 6.1):

Figure 6.1: What is economics?

		What is the subject matter?	
		Some forms of behavior	*All behavior*
What is the method?	*Rational choice analysis*	Core economics	Imperial economics
	Focus on non-rational mechanisms (psychology, norms)	Behavioral economics	Imperial sociology

Those who try to extend rational choice analysis to other areas (crime, divorce, marriage, ethnic conflict and even suicide) can be labeled imperial economists. A prime example in this tradition is the already mentioned Gary Becker, who in addition to the previously quoted definition of economics as the application of rational choice also writes that "Indeed, I have come to the position that the economic approach is a comprehensive one that is applicable to all human behavior be it … repeated or infrequent decisions, large or minor decisions…" (Becker 1986, p. 112).

[52] I am here ignoring those macroeconomists who believe economics is the analysis of aggregates i.e. not individual behaviour, motivational structure or beliefs.

In contrast to the unrestricted scope claimed by Becker—but in agreement on method—we could cite Frank Hahn (1973, p. 21): "In our present state of knowledge however it is routine behavior and not behavior which we can hope to describe." This defines, for instance, marriage decisions as outside the scope of economics (except, maybe, in Hollywood!). I have labeled this position—restricted scope and emphasis on rational choice analysis—"core economics."

Third, there are those who argue that economist should be much more concerned about sociology. For instance, one might use cultural differences to explain differences in economic growth. One might also, as Solow (1980), explain unemployment as the outcome of a norm about the unfairness of reducing wages even when there is excess supply of labor (see Smelser and Swedberg 1994, for more on sociology and economics).

Finally, we have the behavioral economists who want to use psychology to explain traditional economic problems like why firms behave the way they do (Winter 1986) and economic fluctuations (e.g. animal spirits in Keynes 1936).

It is not my purpose here give a comprehensive answer as to who is right and exactly where economists or different arguments should be located in the grid above. Even if this were the purpose, it may not be a very fruitful task. To ask who is right only gives meaning when there is a conflict *and* we have neutral criteria to determine who is correct. And even if we could agree to some of the general criteria that should be satisfied by all good economic approaches (explanatory success, predictive success, fruitfulness, parsimony), the answer may still be indeterminate—what is fruitful in your opinion need not be fruitful in my opinion. Moreover, there is no neutral weighting between explanatory scope and success. One approach may be very successful in a very restricted domain, but how is this to be compared against a less parsimonious theory that is somewhat less successful in a wider domain?

It is, however, my purpose to discuss how the arguments in the previous chapter relate to the question of behavioral economics. First of all, if I am correct in arguing that there is no logical infinite regress argument in the collection of information, then it is wrong to use *that* infinite regress argument in support of behavioral economics. The argument is particularly important because infinite regress is presented as a *logical* impossibility in the theory of rational choice. Logical arguments of this type are more damaging than empirical arguments of the type "this is not how most people usually behave." Thus, it may seem that my argument in the previous chapter supports the imperial economists against the behavioral economists.

The second point reduces the force of the first conclusion. Even if the infinite regress argument about the collection of information fails, we are still left with the estimation problem. The cognitive abilities that are needed to solve this problem are—as we have seen—very large. This is a well known point which need not be elaborated (see e.g. Arrow, 1987). This, in turn, strengthens the claim that we need to examine the cognitive processes to understand economic outcomes.

Third, it is not only in the collection of information that the infinite regress argument may appear. For instance, Lipman (1991) and Smith (1991) both discuss how we should decide how to decide. Although this problem is clearly related to the problem of how to decide how much information to collect they are conceptually distinct (The problems are related in two senses; First, to decide how to decide we need to collect information—one problem involves the other. Second, they are structurally similar—Vassilakis (1992) tries to give a general solution to all such problems). There is also a potential infinite regress problem in the formation of beliefs for given information. In sum, there may still be infinite regress problems in the theory of rational choice theory that justify an increased interest in the behavioral approach, even if the one particular infinite regress argument discussed in the previous chapter fails.

6.3 "As-if"

After arguing against the realism of the assumption of optimization, Elster follows Winter in arguing against the use of Friedman's (1953) well known "as-if" justification for studying optimizing behavior. The argument, briefly, is that natural selection implies that those who do not optimize (or almost optimize) will be eliminated (e.g. go bankrupt). Thus, even if agents do not consciously optimize, it still makes sense to analyze their behavior "as-if" they were optimizing.

Against this Elster and Winter argue that there is good reason to disbelieve the analogy with natural selection in biology. For instance, if there are economies of scale then a large but unfit firm could eliminate those who are "more fit." Winter makes a long list of these exceptions (see Hodgeson 1994 for a discussion). Elster focuses on another aspect—the implicit assumption that the selection mechanism works faster than the changes in the environment. Once again I am in no position to give a definite answer to this criticism, but I do believe I have found some possible weaknesses in Elster's argument.

6.3.1 Rapidly changing environments and other variables

One possible objection to Elster's argument against Friedman could be that social environments are not fast changing or that the speed of environmental change is not the only relevant variable. Consider the following: There are two types of agents (P1 and P2) with different sets of characteristics (A-F):

P1: A, C, D

P2: A, B, E, F

A change in the environment may make characteristic D more functional i.e. having D is in some way beneficial in the new environment. Furthermore, there might be another structural change making B more functional and C dysfunctional. Does this mean that we are unlikely to see a clear change in the composition of P1 v. P2 in the environment?

The answer depends on the degree of the shift i.e. we need to know how much more "functional" P1 becomes as a result of having B and how much less functional it is to have characteristic C in the new environment. Hence, it is not only the speed of environmental change that matters, but also the quantity of change. As an illustration, consider the claim that the industrial revolution changed the structure of the family from extended to nuclear because a nuclear family was more functional in an industrial society. It would be incorrect to argue that his argument is invalid just because there were many other environmental changes going on at the same time (i.e. that the environment was fast changing). First, this would require the other changes to work in the opposite direction (i.e. in favor of the extended family). Second, it would require that the quantitative effect of these changes outweighed the quantitative effects of the industrial revolution. Hence, speed matters but so does quantity and direction.

6.3.2 Why speed might not destroy the selection mechanism: Flexibility

To argue that some environments change so fast that the social selection mechanism does not have time to eliminate the dysfunctional agents, overlooks the possibility that speed itself may make some individuals functional or dysfunctional. Some individuals have more of those characteristics that are functional in rapidly changing environments. For example, ability to learn quickly might lead to greater flexibility in the face of speedy environmental change. It is important to note here that they are functional precisely because the environment is changing. This implies that agents with less of these skills are less functional in fast changing environments. Thus, speed in itself does not make "as-if" arguments invalid.

I want to emphasize once again that this is not to argue in favor of functional explanations in general, only to point out that there may be some weaknesses in Elster's arguments against the use of the "as-if" assumption based on natural selection.

6.3.3 "As-if" without natural selection

Even if the natural selection argument fails to justify the use of "as-if" reasoning, one might still—and this may be a better interpretation of Friedman's position—argue that it is legitimate to use unrealistic assumptions as long as they deliver useful results.[53] For instance, Friedman's (1953) famous billiard table analogy is about how a complex mathematical model may accurately predict a pool player's shot, although the player himself does not calculate his shot using the same mathematical model. Thus, in addition to the literal "as-if" justification (firms that do not maximize profit go bankrupt), it is possible to defend "as-if" using the methodological argument of instrumentalism.

There is a huge philosophical literature on instrumentalism (see, for instance, Hausmann 1998), and I only want to raise a few points that I find relevant. First, there is the obvious argument that economists do not only want to predict, they also want to explain. If the theory is build on an assumption that is both important to the result (so it is not just a simplification) and false, then the theory cannot be said to explain the phenomena in question. Second, there is a problem in terms of measuring accuracy of predictions. Friedman himself has been criticised for ignoring many econometric problems in his work on the money supply and inflation (e.g. spurious correlation, see Gilbert 1986 [and the references there] for an overview of the methodology of one of Friedman's strongest critics, D. Hendry). Third, there is the already mentioned problem that there is no neutral way of choosing between a theory that is moderately accurate for many phenomena vs. one that is highly accurate for only a small subset of phenomena (i.e. the trade-off between universalism and accuracy). These comments, of course, only scratch the surface of a large and controversial literature.

[53] I should thank Kåre Bævre whose comments made me realize (which I should have done long time ago), that there are at least two different "as-if" justifications. First, the literal analogy with natural selection. Second, the instrumental positions that unrealistic assumptions are not problematic as long as the predictions are (reasonably) accurate.

6.3.4 Sub-conclusion

The section above has considered the argument that the problems in the collection of information should lead us to focus on behavioral economics as opposed to approaches based on the assumption of perfectly rational agents. I first created a rough distinction between some possible conceptions of economics. I then examined three arguments about how the problems of information (regress and estimation) related to the argument about how we should do economics (Why infinite regress in the collection is not a good justification for behavioral economics; Why the problem of estimation is a better justification; And why there, after all, may be logical infinite regress problems at other levels than the collection of information). I then examined the so-called "as-if" argument—arguing that there were some weak points in Elster's dismissal of this. The speed of the environment was only one variable among many and the as-if justification need not rely on natural selection at all.

7 Conclusion

Rational choice theory can be attacked for many reasons. However, after reviewing Elster's arguments I no longer belive it is a significant objection to argue that it is impossible to make a rational decision because there is an infinite regress problem in the collection of information that makes a rational decision logically impossible. The problem of estimation is more significant, but it does not prove the *impossibility* of making a rational choice. One might also question the implications that Elster claims follows from the fact that it is difficult to make accurate estimations (i.e. the recommendation of randomization and maximin). Finally, some of Elster's arguments on the issue are, if not contradictory then at least in tension with each other.

As mentioned in the introduction, I do not want to make any strong conclusions on the larger debate of behavioral vs. traditional economics based on my evaluation of Elster's arguments. I do, however, want to mention briefly one limitation and one lesson that I encountered in the more specific field of information-gathering. Both represent what I believe could be useful areas for research. First, my analysis of information gathering is limited in the sense that there are many important variables that are not explored, for instance the implications of various degrees of irreversibility (see Hirshleifer and Riley 1992, pp. 204-208). Second, Hey's (1981) suggestion about reasonable rules may, if developed further, have useful real-life implications for those who routinely have to make decisions about how much information to collect, be it engineers in the oil industry or health administrators trying to decide on what kind of standard procedures to adopt. Thus, a natural extension of this dissertation would be to study one concrete problem of information gathering (Nelson and Winter's [1964] study of the value of weather information is a good example) *combined with research on how different rules could improve the robustness and/or the optimality of the end-result* (Hey [1982] exemplifies this kind of exploratory research, but the analysis is not tied to a real-life problem of information gathering). Having dealt with some of the preliminary conceptual problems in this dissertation ("it is impossible because of infinite regress in the collection of information") the path is open for combining the strengths of Nelson & Winter and Hey to study, for instance, a specific problem of information-gathering in medicine.

8 References

Allais, M. (1987): "Allais paradox", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave)*, W.W. Norton & Company, New York and London, 3-9.

Arrow, K.J. and L. Hurwicz (1972): "An optimality criterion for decision-making under ignorance", in C.F. Carter and J.L. Ford (eds.): *Uncertainty and expectations in economics*, Angus M. Kelley Publishers, New Jersey, 1-11.

Arrow, K.J. (1987): "Economic theory and the hypothesis of rationality", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave)*, W.W. Norton & Company, New York and London, 25-37.

Arrow, K.J., E. Colombatto, M. Perlman and C. Schmidt (eds.) (1996): *The rational foundations of economic behaviour*, Macmillan Press ltd., Houndmills.

Backhouse, R.E. (1995): *Interpreting macroeconomics: Explorations in the history of macroeconomic thought,* Routledge, London and New York.

Barber, W.J. (1997): "Reconfigurations in American academic economics: A general practitioner's view", *Daedalus*, 126, 87-103.

Baumol, W.J. and R.E. Quandt (1964): "Rules of thumb and optimally imperfect decisions", *American Economic Review*, 54, 23-46.

Becker, G.S. (1986/1976): "The economic approach to human behavior", in J. Elster (ed.): *Rational Choice*, Basil Blackwell, Oxford, 108-122.

Bhattacharyya, G.K. and R. A. Johnson (1977): *Statistical concepts and methods*, John Wiley & Sons, New York.

Conlisk, J. (1996): "Why bounded rationality?", *Journal of Economic Literature*, 34, 669-700.

Connor, W. (1994): *Ethnonationalism: The quest for understanding*, Princeton University Press, Princeton.

Davidson, P. (1991): "Is probability theory relevant for uncertainty? A post keynesian perspective", *Journal of Economic Perspectives*, 5, 129-143.

Dawes, R.M. (1988): *Rational choice in an uncertain world*, Harcourt Brace College Publishers, London.

Demsetz, H. (1997): "The primacy of economics: An explanation of the comparative success of economics in the social sciences", *Economic Inquiry,* 35, 1-11.

Elster, J. (1979/1984 revised edition): *Ulysses and the sirens: Studies in rationality and irrationality,* Cambridge University Press, Cambridge.

Elster, J. (1982): "Rationality", in G. Fløistad (ed.): *Contemporary philosophy. A new survey*, Martinus Nijhoff Publishers, The Hague, 111-131.

Elster, J. (1983a): "The crisis in economic theory (Review of R. Nelson and S.G. Winter [1982)] *An Evolutionary Theory of Economic Change* and J. Roemer [1982] *A General Theory of Exploitation and Class*), *London Review of Books*, 4 (9), 5-7.

Elster, J. (1983b): *Explaining technical change*, Cambridge University Press, Cambridge.

Elster, J. (1983c): *Sour grapes: Studies in the subversion of rationality*. Cambridge University Press, Cambridge.

Elster, J. (1985): "The nature and scope of rational-choice explanations", in E. LePore and B.P. McLaughlin (eds.): *Actions and Events: Perspectives on the philosophy of Donald Davidson*, Blackwell, Oxford, 60-72.

Elster J. (1986): "Introduction", in J. Elster (ed.): *Rational choice*, Basil Blackwell, Oxford.

Elster, J. (1987): "The possibility of rational politics", *Archives Europeennes de Sociologie,* 28, 67-103.

Elster, J. (1989a): *Solomonic judgements: Studies in the limitations of rationality,* Cambridge University Press, Cambridge.

Elster, J. (1989b): *Nuts and bolts for the social sciences,* Cambridge University Press, Cambridge.

Elster, J. (1989c): *Vitenskap og politikk*, Universitetsforlaget, Oslo.

Elster, J. (1993): "Some unresolved problems in the theory of rational behaviour", *Acta Sociologica*, 36, 179-190.

Elster, J. (1999): *Strong feelings: Emotion, addiction and human behavior*, The MIT Press, Cambridge, Massachusetts.

Fisher, F. M. (1987): "Adjustment processes and stability." in J. Eatwell, M. Milgate and P. Newman (eds.): *The new palgrave: A dictionary of economics*, W.W. Norton & Company, New York and London, 26-29.

Friedman, M. (1953/1979): "The methodology of positive economics", in M. Hollis and F. Hahn (eds.) (1979): *Philosophy and economic theory*, Oxford University Press, Oxford, 18-35.

Gärdenfors, P. and N-E. Sahlin (1988): *Decision, probability and utility: Selected readings,* University of Cambridge Press, Cambridge.

Gelman, A. (1998): "Some class-participation demonstrations for decision theory and bayesian statistics", *The American Statistician*, 52, 167-174.

Gilbert, C.L. (1986): "Professor Hendry's econometric methodology", *Oxford Bulletin of Economics and Statistics,* 48, 283-307.

Hacking, I. (1967): "Slightly more realistic personal probability", *Philosophy of Science*, 34, 311-325.

Hahn, F.H. (1973): *On the notion of equilibrium in economics*, Cambridge University Press, Cambridge.

Hampton, J. (1994): "The failure of expected utility theory as a theory of reason", *Economics and Philosophy*, 10, 195-242.

Hardin, R. (1988): *Morality within the limits of reason*, University of Chicago Press, Chicago.

Hardin, R. (1995): *One for all: The logic of group conflict*, Princeton University Press, Princeton.

Haregreaves Heap, S., M. Hollis, B. Lyons, R. Sugden, A. Weale (1992): *The theory of choice: A critical guide*, Blackwell, Oxford.

Hausman, D.M. (1998): "Problems with realism in economics", *Economics and Philosophy*, 14, 185-213.

Hey, J.D. (1981): "Are optimal search rules reasonable and vice versa? (and does it matter anyway?)", *Journal of Economic Behavior and Organization*, 2, 47-70.

Hey, J.D. (1982): "Search for rules for search", *Journal of Economic Behavior and Organization*, 3, 65-81.

Hirshleifer, J. (1985): "The expanding domain of economics", *American Economic Review*, 75, 53-68.

Hirshleifer, J. and J.G. Riley (1992): *The analytics of uncertainty and information*, Cambridge University Press, Cambridge.

Hodgson, G.M. (1994): "Optimization and evolution: Winter's critique of Friedman revisited", *Cambridge Journal of Economics*, 18, 413-430.

Hodgson, G.M. (1997): "The ubiquity of habits and rules", *Cambridge Journal of Economics*, 21, 663-684.

Hoogduin, L.H. and J. Snippe (1987): "Uncertainty in/of macroeconomics: An essay on adequate abstraction", *De Economist*, 135, 429-441.

Iversen, G.R. (1984): *Bayesian statistical inference*, SAGE Publications inc., Beverly Hills, CA. (Series: Quantitative Applications in the Social Sciences, no. 43 [07-043]).

Janssen, M.C.W. (1993): *Microfoundations: A critical inquiry*, Routledge, London and New York.

Johansen, L. (1977): *Lectures on macroeconomic planning*, North-Holland Publishing Company, Amsterdam.

Kelsey, D. and J. Quiggin (1992): "Theories of choice under ignorance and uncertainty", *Journal of Economic Surveys*, 6, 133-153.

Keynes, J.M. (1936), *The general theory of employment, interest, and money*, New York: Harcourt, Brace & World. (My version, reprint from 1997, Amherst, New York: Prometheus Books.)

Kirman, A.P. (1992): "Whom or what does the representative agent represent?", *Journal of Economic Perspectives*, 6, 117-136.

Lawson, T. (1988): "Probability and uncertainty in economic analysis", *Journal of Post Keynesian Economics*, 11, 38-65.

Lipman, B. (1991): "How to decide how to decide how to...: Modeling limited rationality", *Econometrica*, 59, 1105-1125.

Machina, M.J. (1987a): "Expected utility hypothesis", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave*), W.W. Norton & Company, New York and London, 79-95.

Machina, M.J. (1987b): "Choice under uncertainty: Problems solved and unsolved", *Journal of Economic Perspectives*, 1, 121-154.

Melberg, H.O. (1996): "Tre årsaker til urettferdigheter: Uunngåelige begrensninger, utopiske handlinger og ubevisste mekanismer", *Lov og Rett*, 471-499.

Melberg, H.O. (1998): "Et forsvar for antagelsen om rasjonalitet", *Impuls*, 52, 27-37.

Mongin, P. and B. Walliser (1988): "Infinite regressions in the optimizing theory of decision", in B.E. Munier (ed.): *Risk, decision and rationality*, D. Reidel Publishing Company, Dordrecht, 435-457.

Nelson, R.R. and S.G. Winter (1964): "A case study in the economics of information and coordination: The weather forecasting system", *Quarterly Journal of Economics*, 78, 420-441.

Newman, P. (1987): "Frank Plumpton Ramsey", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave*), W.W. Norton & Company, New York and London, 186-197.

Nozick, R. (1981): *Philosophical explanations*, Oxford University Press, Oxford.

Nozick, R. (1993): *The nature of rationality*, Princeton University Press, Princeton.

Pingle, M. (1992): "Costly optimization: An experiment," *Journal of Economic Behavior and Organization*, 17, 3-30.

Radner, R. (1996): "Bounded rationality, indeterminacy, and the theory of the firm", *Economic Journal*, 106, 1360-1373.

Raiffa, H. (1968): *Decision analysis: Introductory lectures on choices under uncertainty*, Addison-Wesley, Reading, MA.

Resnik, M.D. (1987): *Choices: An introduction to decision theory*, University of Minnesota Press, Minneapolis and London.

Rotschild, M. (1974): "Searching for the lowest price when the distribution of prices is unknown", *Journal of Political Economy*, 82, 689-711.

Sargent, T.J. (1991): "Equilibrium with signal extraction from endogenous variables", *Journal of Economic Dynamics and Control*, 15, 245-273.

Savage, L.J. (1967): "Difficulties in the theory of personal probability", *Philosophy of Science*, 34, 305-310.

Schmeidler D. and P. Wakker (1987): "Expected utility and mathematical expectation," in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave)*, W.W. Norton & Company, New York and London, 70-78.

Schoemaker, P.J. (1982): "The expected utility model: Its variants, purposes, evidence and limitations", *Journal of Economic Literature*, 20, 529-563.

Sen, A. K. (1987): "Rational behaviour", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The new palgrave)*, W.W. Norton & Company, New York and London, 198-216.

Shulman, S. (1997): "What's so rational about rational expectations? Hyperrationality and the logical limits to neoclassicism", *Journal of Post Keynesian Economics*, 20, 135-148.

Simon, H.A. (1979): "Rational decision-making in business organizations", *American Economic Review*, 69, 493-513.

Simon, H.A. (1987): "Bounded rationality", in J. Eatwell, M. Milgate and P. Newman (eds.): *Utility and probability (The New Palgrave)*, W.W. Norton & Company, New York and London, 15-18.

Smelser N. and R. Swedberg (eds.) (1994): *The handbook of economic sociology*, Princeton University Press, Princeton, 3-26.

Smith, H. (1991): "Deciding how to decide: Is there a regress problem?" in M. Bacharach and S. Hurley (eds.): *Foundations of decision theory*, Blackwell, Oxford, 194-219.

Solow, R.M. (1980): "On theories of unemployment", *American Economic Review*, 70, 1-11.

Stigler, G.J. (1961): "The economics of information", *Journal of Political Economy*, 69, 213-225.

Stigler, G.J. (1984): "Economics—The imperial science", *Scandinavian Journal of Economics*, 86, 301-313.

Sugden, R. (1991): "Rational choice: A survey of contributions from economics and philosophy", *Economic Journal*, 101, 751-785.

Vassilakis, S. (1992): "Some economic applications of Scott domains", *Mathematical Social Sciences*, 24, 173-208.

Winter, S.G. (1964): "Economic 'natural selection' and the theory of the firm", *Yale Economic Essays*, 4, 225-272.

Winter, S.G. (1971): "Satisficing, selection, and the innovating remnant", *Quarterly Journal of Economics*, 85, 237-261.

Winter, S.G. (1975): "Optimization and Evolution in the theory of the firm", in R.H. Day and T. Groves (eds.): *Adaptive economic models*, Academic Press, New York, 73-118.

Winter, S.G. (1986): "The research program of the behavioral theory of the firm: Orthodox critique and evolutionary perspective", in B. Gilad and S. Kaish (eds.): *Handbook of behavioral economics* (volume A), JAI Press Inc., Greenwich, Connecticut, 151-188.

Winter, S.G. (1987): "Comments on Arrow and Lucas", in R.M. Hogarth and M.W. Reder (eds.): *Rational choice: The contrast between economics and psychology*, The University of Chicago Press, Chicago and London, 243-250.