

A Theory of Case-Based Decisions

by

Itzhak Gilboa

and

David Schmeidler

2

To our families

Acknowledgments

We thank many people for conversations, comments, and references that have influenced our thinking and shaped this work. In particular, we wish to thank Yoav Binyamini, Steven Brams, Didier Dubois, Ilan Eshel, Eva Gilboa-Schechtman, Ehud Kalai, Edi Karni, Kimberly Katz, Daniel Lehman, Akihiko Matsui, Sujoy Mukerji, Roger Myerson, Andrew Postlewaite, Ariel Rubinstein, Dov Samet, Peter Wakker, Bernard Walliser, and Peyton Young. Special thanks go to Benjamin Polak and to Enriqueta Aragonés for many specific comments on earlier drafts of several chapters.

“In reality, all arguments from experience are founded on the similarity which we discover among natural objects, and by which we are induced to expect effects similar to those which we have found to follow from such objects. And though none but a fool or madman will ever pretend to dispute the authority of experience, or to reject that great guide of human life, it may surely be allowed a philosopher to have so much curiosity at least as to examine the principle of human nature, which gives this mighty authority to experience, and makes us draw advantage from that similarity which nature has placed among different objects. From causes which appear *similar* we expect similar effects. This is the sum of all our experimental conclusions.”
(Hume 1748, Section IV)

Contents

1	Prologue	9
1	The Scope of This Book	9
2	Meta-Theoretical Vocabulary	12
2.1	Theories and conceptual frameworks	13
2.2	Descriptive and normative theories	16
2.3	Axiomatizations	20
2.4	Behaviorist, behavioral, and cognitive theories	24
2.5	Rationality	25
2.6	Deviations from rationality	27
2.7	Subjective and objective terms	27
3	Meta-Theoretical Prejudices	29
3.1	Preliminary remark on the philosophy of science	29
3.2	Utility and expected utility “theories” as conceptual frameworks and as theories	31
3.3	On the validity of purely behavioral economic theory	33
3.4	What does all this have to do with CBDT?	35
2	Decision Rules	37
4	Elementary Formula and Interpretations	37
4.1	Motivating examples	38
4.2	Model	42
4.3	Aspirations and satisficing	47

4.4	Comparison with EUT	50
4.5	Comments	54
5	Variations and Generalizations	55
5.1	Average similarity	55
5.2	Act similarity	57
5.3	Case similarity	59
6	CBDT as a Behaviorist Theory	61
6.1	<i>W</i> -Maximization	61
6.2	Cognitive Specification: EUT	63
6.3	Cognitive Specification: CBDT	64
6.4	Comparing the cognitive specifications	65
7	Case-Based Prediction	67
3	Axiomatic Derivation	71
8	Highlights	71
9	Model and Result	74
9.1	Axioms	74
9.2	Basic result	76
9.3	Learning new cases	77
9.4	Equivalent cases	78
9.5	<i>U</i> -Maximization	80
10	Discussion of the Axioms	82
11	Proofs	86
4	Conceptual Foundations	99
12	CBDT and Expected Utility Theory	99
12.1	Reduction of theories	99
12.2	Hypothetical reasoning	101
12.3	Observability of data	103
12.4	The primacy of similarity	104
12.5	Bounded rationality?	105

13	CBDT and Rule-Based Systems	106
13.1	What can be known?	106
13.2	Deriving case-based decision theory	108
13.3	Implicit knowledge of rules	113
13.4	Two roles of rules	115
5	Planning	117
14	Representation and Evaluation of Plans	117
14.1	Dissection, selection, and recombination	117
14.2	Representing uncertainty	120
14.3	Plan evaluation	123
14.4	Discussion	126
15	Axiomatic Derivation	128
15.1	Set-up	128
15.2	Axioms and result	130
15.3	Proof	131
6	Repeated Choice	135
16	Cumulative Utility Maximization	135
16.1	Memory-dependent preferences	135
16.2	Related literature	137
16.3	Model and results	139
16.4	Comments	143
16.5	Proofs	144
17	The Potential	147
17.1	Definition	147
17.2	Normalized potential and neo-classical utility	149
17.3	Substitution and Complementarity	151
7	Learning and Induction	157
18	Learning to Maximize Expected Payoff	157

18.1	Aspiration level adjustment	157
18.2	Realism and ambitiousness	159
18.3	Highlights	161
18.4	Model	164
18.5	Results	167
18.6	Comments	169
18.7	Proofs	172
19	Learning the Similarity Function	186
19.1	Examples	186
19.2	Counter-example to U -maximization	189
19.3	Learning and expertise	193
20	Two Views of Induction: CBDT and Simplicism	195
20.1	Wittgenstein and Hume	195
20.2	Examples	196
	Bibliography	200
	Index	211

Chapter 1

Prologue

1 The Scope of This Book

The focus of this book is formal modeling of decision making by a single person who is aware of the uncertainty she is facing. Some of the models and results we propose may be applicable to other situations. For instance, the decision maker may be an organization or a computer program. Alternatively, the decision maker may not be aware of the uncertainty involved or of the very fact that a decision is being made. Yet, our main interest is in descriptive and normative models of conscious decisions made by humans.

There are two main paradigms for formal modeling of human reasoning, which have also been applied to decision making under uncertainty. One involves probabilistic and statistical reasoning. In particular, the Bayesian model coupled with expected utility maximization is the most prominent paradigm for formal models of decision making under uncertainty. The other employs rule-based deductive systems. Each of these paradigms provides a conceptual framework and a set of guidelines for constructing specific models for a wide range of decision problems.

These two paradigms are not the only ways in which people's reasoning may be, or has been, described. In particular, the claim that people reason by analogies dates back at least to Hume. However, reasoning by analogies

has not been the subject of formal analysis to the same degree that the other paradigms have. Moreover, there is no general purpose theory we are aware of that links reasoning by analogies to decision making under uncertainty.

Our goal is to fill this gap. That is, we seek a general purpose formal model, comparable to the model of expected utility maximization, that will (i) provide a framework within which a large class of specific problems can be modeled; (ii) be based on data that are, at least in principle, observable; (iii) allow mathematical analysis of qualitative issues, such as asymptotic behavior; and (iv) be based on reasoning by analogies.

We believe that human reasoning typically involves a combination of the three basic techniques, namely, rule-based deduction, probabilistic inference, and analogies. Formal modeling tends to opt for elegance, and to focus on certain aspects of a problem at the expense of others. Indeed, our aim is to provide a model of case- or analogy-based decision making that will be simple enough to highlight main insights. We discuss the various ways in which our model may capture deductive and probabilistic reasoning, but we do not formally model the latter. It should be taken for granted that a realistic model of the human mind would have to include ingredients of all three paradigms, and perhaps several others as well. At this stage we merely attempt to lay the foundations for one paradigm whose absence from the theoretical discussion we find troubling.

The theory we present here does not purport to be more realistic than other theories of human reasoning or of choice. In particular, our goal is *not* to fine-tune expected utility theory as a descriptive theory of decision making in situations described by probabilities or states of the world. Rather, we wish to suggest a framework within which one can analyze choice in situations that do not fit existing formal models very naturally. Our theory is just as idealized as existing theories. We only claim that in many situations it is a more natural conceptualization of reality than are these other theories.

This book does not attempt to provide even sketchy surveys of the estab-

lished paradigms for formal modeling of reasoning, or of the existing literature on case-based reasoning. The reader is referred to standard texts for basic definitions and background.

Many of the ideas and mathematical results in this book have appeared in journal articles and working papers (Gilboa and Schmeidler 1995, 1996a, 1997a,b, 2000a,b,c). This material has been integrated, organized, and interpreted in new ways. Additionally, several sections appear here for the first time.

In writing this book, we made an effort to address readers from different academic disciplines. Whereas several chapters are of common interest, others may address more specific audiences. The following is a brief guide to the book.

We start with two meta-theoretical sections, one devoted to definitions of philosophical terms, and the other to our own views on the way decision theory and economic theory should be conducted. These two sections may be skipped with no great loss to the main substance of the book. Yet, Section 2 may help to clarify the way we use certain terms (such as “rationality”, “normative science”, and the like), and Section 3 explains part of our motivation in developing the theory described in this book.

Chapter 2 of the book presents the main ideas of case based decision theory (CBDT), as well as its formal model. It offers several decision rules, a behaviorist interpretation of CBDT, and a specification of the theory for prediction problems.

Chapter 3 provides the axiomatic foundations for the decision rules in Chapter 2. In line with the tradition in decision theory and in economics, it seeks to relate theoretical concepts to observables and to specify conditions under which the theory might be refuted.

Chapter 4, on the other hand, focuses on the epistemological underpinnings of CBDT. It compares it with the other two paradigms of human reasoning and argues that, from a conceptual viewpoint, analogical reasoning

is primitive, whereas both deductive inference and probabilistic reasoning are derived from it. Whereas Chapter 3 provides the mathematical foundations of our theory, the present chapter offers the conceptual foundations of the theory and of the language within which the mathematical model is formulated.

Chapter 5 deals with planning. It generalizes the CBDT model from a single-stage decision to a multi-stage one, and offers an axiomatic foundation for this generalization.

Chapter 6 focuses on a special case of our general model, in which the same problem is repeated over and over again. It relates to problems of discrete choice in decision theory and in marketing, and it touches upon issues of consumer theory. It also contains some results that are used later in the book.

Chapter 7 addresses questions of learning, dynamic evolution, and induction in our model. We start with an optimality result for the case of a repeated problem, which is based on a rather rudimentary form of learning. We continue to discuss more interesting forms of learning, as well as inductive inference. Unfortunately, we do not offer any profound results about the more interesting issues. Yet, we hope that the formal model we propose may facilitate discussion of these issues.

2 Meta-Theoretical Vocabulary

We devote this section to define the way we use certain terms that are borrowed from philosophy. Definitions of terms and distinctions among concepts tend to be fuzzy and subjective. The following are no exception. These are merely the definitions that we have found to be the most useful for discussing theories of decision making under uncertainty at the present state. While our definitions are geared toward a specific goal, several of them may facilitate discussion of other topics as well.

2.1 Theories and conceptual frameworks

A theory of social science can be viewed as a formal mathematical structure coupled with an informal interpretation. Consider, for example, the economic theory that consumer’s demand is derived from maximizing a utility function under a budget constraint. A possible formal representation of this theory consists of two components, describing two sets, \mathcal{C} and \mathcal{P} . The first set, \mathcal{C} , consists of all conceivable demand functions. A demand function maps a vector of positive prices $p \in \mathbb{R}_{++}^n$ and an income level $I \in \mathbb{R}_+$ to a vector of quantities $d(p, I) \in \mathbb{R}_+^n$, interpreted as the consumer’s desired quantities of consumption under the budget constraint that total expenditure $d(p, I) \cdot p$ does not exceed income I . The second set, \mathcal{P} , is the subset of \mathcal{C} that is consistent with the theory. Specifically, \mathcal{P} consists of the demand functions that can be described as maximizing a utility function.¹ When the theory is descriptive, the set \mathcal{P} is interpreted as all phenomena (in \mathcal{C}) that might actually be observed. When the theory is normative, \mathcal{P} is interpreted as all phenomena (in \mathcal{C}) that the theory recommends. Thus, whether the theory is descriptive or normative is part of the informal interpretation.

The informal interpretation should also specify the intended applications of the theory. This is done at two levels. First, there are “nicknames” attached to mathematical objects. Thus \mathbb{R}_+^n is referred to as a set of “bundles”, \mathbb{R}_{++}^n – as a set of positive “price vectors”, whereas I is supposed to represent “income” and d – “demand”. Second, there are more detailed descriptions that specify whether, say, the set \mathbb{R}_+^n should be viewed as representing physical commodities in an atemporal model, consumption plans over time, or financial assets including contingent claims, whether d denotes the demand of an individual or a household, and so forth.

Generally, the *formal structure* of a theory consists of a description of

¹Standard (neo-classical) consumer theory imposes additional constraints. For instance, homogeneity and continuity are often part of the definition of demand functions, and utility functions are required to be continuous, monotone, and strictly quasi-concave. We omit these details for clarity of exposition.

a set \mathcal{C} and a description of a subset thereof, \mathcal{P} . The set \mathcal{C} is understood to consist of conceivably observable phenomena. It may be referred to as the *scope* of the theory. A theory thus selects a set of phenomena \mathcal{P} out of the set of conceivable phenomena \mathcal{C} , and excludes its complement $\mathcal{C}\setminus\mathcal{P}$. What is being said about this set \mathcal{P} , however, is specified by the informal interpretation: it may be the *prediction* or the *recommendation* of the theory.

Observe that the formal structure of the theory does not consist of the sets \mathcal{C} and \mathcal{P} themselves. Rather, it consists of *formal descriptions of these sets*, $\mathcal{D}_{\mathcal{C}}$ and $\mathcal{D}_{\mathcal{P}}$, respectively. These formal descriptions are strings of characters that define the sets in standard mathematical notation. Thus, theories are not extensional. In particular, two different mathematical descriptions $\mathcal{D}_{\mathcal{P}}$ and $\mathcal{D}'_{\mathcal{P}}$ of the same set \mathcal{P} will give rise to two different theories. It may be a non-trivial mathematical task to discover relationships between sets described by different theories.

It is possible that two theories that differ not only in the formal structure $(\mathcal{D}_{\mathcal{C}}, \mathcal{D}_{\mathcal{P}})$ but also in the sets $(\mathcal{C}, \mathcal{P})$ may coincide in the real world phenomena they describe. For example, consider again the paradigm of utility maximization in consumer theory. We have spelled out above one manifestation of this paradigm in the language of demand functions. But the literature also offers other theories within the same paradigm. For instance, one may define the set of conceivable phenomena to be all binary relations over \mathbb{R}_+^n , with a corresponding definition of the subset of these relations that conform to maximization of a real-valued function.

The informal interpretation of a theory may also be formally defined. For instance, the assignment of *nicknames* to mathematical objects can be viewed as a mapping from the formal descriptions of these objects, appearing in $\mathcal{D}_{\mathcal{C}}$ and in $\mathcal{D}_{\mathcal{P}}$, into a natural language, provided that the latter is a formal mathematical object. Similarly, one may formally define “real world phenomena” and represent the *(intended) applications* of the theory as a collection of mappings from the mathematical entities to this set. Finally,

the *type of interpretation* of the theory, namely, whether it is descriptive or normative, can easily be formalized.² Thus a theory may be described as a quintuple consisting of \mathcal{D}_C , \mathcal{D}_P , the nicknames assignment, the applications, and the type of interpretation.

We refer to the first three components of this quintuple, that is, \mathcal{D}_C , \mathcal{D}_P , and the nicknames assignment, as a *conceptual framework* (or *framework* for short). A conceptual framework thus describes a scope and a description of a prediction or a recommendation, and it points to a type of applications through the assignment of nicknames. But a framework does not completely specify the applications. Thus, frameworks fall short of qualifying as theories, even if the type of interpretation is given.

For instance, Savage’s (1954) model of expected utility theory involves binary relations over functions defined on a measurable space. The mathematical model is consistent with real world interpretations that have nothing to do with choice under uncertainty, such as choice of streams of consumption over time, or of income profiles in a society. The nickname “space of states of the world”, which is attached to the measurable space in Savage’s model, defines a framework that deals with decision under uncertainty. But the conceptual framework of expected utility theory does not specify exactly what the states of the world are, or how they should be constructed. Similarly, the conceptual framework of Nash equilibrium (Nash (1951)) in game theory refers to “players” and to “strategies”, but it does not specify whether the players are individuals, organizations, or states, whether the theory should be applied to repeated or to one-shot situations, to situations involving few or many players, and so forth.

By contrast, the theory of expected utility maximization under risk (von-Neumann and Morgenstern (1944)), as well as prospect theory (Kahneman

²Our formal model allows other interpretations as well. For instance, it may represent a formal theory of aesthetics, where the set \mathcal{P} is interpreted as defining what is beautiful. One may argue that such a theory can still be interpreted as a normative theory, prescribing how aesthetic judgment should be conducted.

and Tversky's (1979)) are conceptual frameworks according to our definition. Still, they may be classified also as theories, because the scope and nicknames they employ almost completely define their applications.

Terminological remark: The discussion above implies that expected utility theory should be termed a framework rather than a theory. Similarly, non-cooperative games coupled with Nash equilibrium constitute a framework. Still, we follow standard usage throughout most of the book and often use “theory” where our vocabulary suggests “framework”.³ However, the term “framework” will be used only for conceptual frameworks that have several substantially distinct applications.

2.2 Descriptive and normative theories

There are many possible meanings to a selection of a set \mathcal{P} out of a set of conceivable phenomena \mathcal{C} . Among them, we find that it is crucial to focus on, and to distinguish between two that are relevant to theories in the social sciences: descriptive and normative.

A descriptive theory attempts to describe, explain, or predict observations. Despite the different intuitive meanings, one may find it challenging to provide formal qualitative distinctions between description and explanation. Moreover, the distinction between these and prediction may not be very fundamental either. We therefore do not dwell on the distinctions among these goals.

A normative theory attempts to provide recommendations regarding what to do. It follows that normative theories are addressed to an audience of people facing decisions who are capable of understanding their recommendations. However, not every recommendation qualifies as normative science. There are recommendations that may be classified as moral, religious, or political preaching. These are characterized by suggesting goals to the decision makers, and, as such, are outside the realm of academic activity. There is an

³We apply the standard usage to the title of this book as well.

additional type of recommendations that we do not consider as normative theories. These are recommendations that belong to the domain of social planning or engineering. They are characterized by recommending tools for achieving pre-specified goals. For instance, the design of allocation mechanisms that yield Pareto optimal outcomes accepts the given goal of Pareto optimality and solves an engineering-like problem of obtaining it. Our use of “normative science” differs from both these types of recommendations.

A normative scientific claim may be viewed as an implicit descriptive statement about decision makers’ preferences. The latter are about conceivable realities that are the subject of descriptive theories. For instance, whereas descriptive theory of choice investigates actual preferences, normative theory of choice analyzes the kind of preferences that the decision maker *would like* to have, that is, preferences over preferences. An axiom such as transitivity of preferences, when normatively interpreted, attempts to describe the way the decision maker would prefer to make choices. Similarly, Harsanyi (1953, 1955) and Rawls (1971) can be interpreted as normative theories for social choice in that they attempt to describe to what society one would like to belong.

According to this definition, normative theories are also descriptive. They attempt to describe a certain reality, namely, the preferences an individual has over the reality she encounters. To avoid confusion, we will reserve the term “descriptive theory” for theories that are not normative. That is, descriptive theories would deal, by definition, with “first-order” reality, whereas normative theories would deal with “second-order” reality, namely, with preferences over first-order reality. First-order reality may be external or objective, whereas second-order reality always has to do with subjective preferences that lie within the mind of an individual. Yet, first-order reality might include actual preferences, in which case the distinction between first-

order and second-order reality may become a relative matter.^{4,5}

Needless to say, these distinctions are sometimes fuzzy and subjective. A scientific essay may belong to several different categories, and it may be differently interpreted by different readers, who may also disagree with the author's interpretation. For instance, the independence axiom of von-Neumann and Morgenstern's expected utility theory may be interpreted as a component of a descriptive theory. Indeed, testing it experimentally presupposes that it has a claim to describe reality. But it may also be interpreted normatively, as a recommendation for decision making under risk. To cite a famous example, Maurice Allais presented his paradox (see Allais (1953)) to several researchers, including the late Leonard Savage. The latter expressed preferences in violation of expected utility theory. Allais argued that expected utility maximization is not a successful descriptive theory. Savage's reply was that his theory should be interpreted normatively, and that it could indeed help a decision maker avoid such mistakes.

Further, even when a theory is interpreted as a recommendation it may involve different types of recommendations. For instance, Shapley axiomatized his value for cooperative transferable utility games (Shapley (1953)). When interpreted normatively, the axioms attempt to capture decision makers' preferences over the way in which, say, cost is allocated in different problems. A related result by Shapley shows that a player's value can be computed as a weighted average of her marginal contributions. This result can be interpreted in two ways. First, one may view it as an engineering recommendation: given the goal of computing a player's value, it suggests a

⁴There is a temptation to consider a hierarchy of preferences, and to ask which are in the realm of descriptive theories. We resist this temptation.

⁵In many cases first-order preferences would be revealed by actual choices, whereas second-order preferences would only be verbally reported. Yet, this distinction is not sharp. First, there may be first-order preferences that cannot be observed in actual choice. Second, one may imagine elaborate choice situations in which second-order preferences might be observed, as in cases where one decides on a decision-making procedure or on a commitment device.

formula for its computation. Second, one may also argue that compensating a player to the extent of her average marginal contribution is ethically appealing, and thus the formula, like the axioms, has a normative flavor.

To conclude, the distinctions between descriptive and normative theories, as well as between the latter and engineering on the one hand and preaching on the other, are not based on mathematical criteria. The same mathematical result can be part of any of these types of scientific or non-scientific activities. These distinctions are based on the author's *intent*, or on her perceived intent. It is precisely the inherently subjective nature of these distinctions that demands that one be explicit about the suggested interpretation of a theory.

It also follows that one has to have a suggested interpretation in mind when attempting to evaluate a theory. Whereas a descriptive theory is evaluated by its conformity with objective reality, a normative theory is not. On the contrary, a normative theory that suggests to people to do what they would anyway do is of questionable value. How should we evaluate a normative theory, then? Since we define normative theories to be descriptive theories dealing with second-order reality, a normative theory should also be tested according to its conformity to reality. But it is second-order reality that a normative theory should be compared to. For instance, a reasonable test of a normative theory might be whether its subjects accept its recommendations.

It is undeniable, however, that the evaluation of normative theories is inherently more problematic than that of descriptive theories. The evidence for normative theories would have to rely on introspection and self-report data much more than would the evidence for descriptive theories. Moreover, these data may be subject to manipulation. To consider a simple example, suppose we are trying to test the normative claim that income should be evenly distributed. We are supposed to find out whether people would like to live in an egalitarian society. Simply asking people might confound their

ethical preferences with their self-interest. Indeed, a person might not be sure whether she subscribes to a philosophical or a political thesis due to pure conviction or to self-serving biases. A gedanken experiment such as putting oneself behind the “veil of ignorance” (Harsanyi (1953, 1955), Rawls (1971)) may assist one in finding one’s own preferences, but may still be of little help in a social context. Further, the reality one tries to model, namely, a person’s ethical preferences over the rules governing society, or her logico-aesthetical preferences over her own preferences are a moving target that changes constantly with actual behavior and social reality. Yet, the way we define normative theories admits a certain notion of empirical testing.

2.3 Axiomatizations

In common usage, the term “axiomatization” refers to a theory. However, most axiomatizations in the literature apply to conceptual frameworks according to our definitions. In fact, the following definition of axiomatizations refers only to a formal structure $(\mathcal{D}_{\mathcal{C}}, \mathcal{D}_{\mathcal{P}})$.

An *axiomatization* of $T = (\mathcal{D}_{\mathcal{C}}, \mathcal{D}_{\mathcal{P}})$ is a mathematical model consisting of: (i) a formal structure $T' = (\mathcal{D}_{\mathcal{C}}, \mathcal{D}_{\mathcal{P}'})$, which shares the description of the set \mathcal{C} with T , but whose set \mathcal{P}' is described only in the language of phenomena that are deemed observable; (ii) mathematical results relating the set \mathcal{P}' of T' to the set \mathcal{P} of T . Ideally, one would like to have conditions on observable phenomena that are necessary and sufficient for \mathcal{P} to hold, namely, to have a structure T' such that $\mathcal{P} = \mathcal{P}'$. In this case, T' is considered to be a “complete” axiomatization of T . The conditions that describe \mathcal{P}' are referred to as “axioms”.⁶

Observe that whether $T' = (\mathcal{D}_{\mathcal{C}}, \mathcal{D}_{\mathcal{P}'})$ is considered to be an axiomatization of T depends on the question of observability of terms in $\mathcal{D}_{\mathcal{P}'}$. Conse-

⁶As opposed to the original meaning of the word, an “axiom” need not be indisputable or self-evident. However, evaluation of axiomatic systems typically prefers axioms that are simple in terms of mathematical formulation and transparent in terms of empirical content.

quently, the definition above will be complete only given a formal definition of “observability”. We do not attempt to provide such a definition here, and we use the term “axiomatization” as if observability were well-defined.⁷ Throughout the rest of this subsection we assume that the applications of the conceptual frameworks are agreed upon. We will therefore stick to standard usage and refer to axiomatizations of theories (rather than of formal structures).

Because human decisions involve inherently subjective phenomena, it is often the case that the formulation of a theory contains existential quantifiers. In this case, a complete axiomatization would also include a result regarding uniqueness. For instance, consider again the theory stating that “there exists a [so-called utility] function such that, in any decision between two choices, the consumer would opt for the one to which the [utility] function attaches a higher value”. An axiomatization of this theory should provide conditions under which the consumer can indeed be viewed as maximizing a utility function in binary decisions. Further, it should address the question of uniqueness of this function: to what extent can we argue that the utility function is defined by the observable data of binary choices?

There are three reasons for which one might be interested in axiomatizations of a theory. First, the *meta-scientific* reason mentioned above: an axiomatization provides a link between the theoretical terms and the (allegedly) observable terms used by the theory. True to the teachings of logical positivism (Carnap (1923), see also Suppe (1974)), one would like to have observable definitions of theoretical terms, in order to render the latter meaningful. Axiomatizations help us identify those theoretical differences that have observable implications, and avoid debates between theories that are observationally equivalent.

One might wish to have an axiomatization of a theory for *descriptive*

⁷Indeed, people who disagree about the definition of observability may consequently disagree whether a certain mathematical result qualifies as an axiomatization.

reasons. Since an axiomatization translates the theory to directly observable claims, it prescribes a way to test the empirical validity of the theory. To the extent that the axioms do rule out certain conceivable observations, they also ascertain that the theory is non-vacuous, that is, falsifiable, as preached by Popper (1934). Note also that there are many situations, especially in the social sciences, where it is impractical to subject a theory to direct empirical tests. In those cases, an axiomatization can help us judge the plausibility of the theory. In this sense, axiomatizations may serve a rhetorical purpose.

Finally, one is often interested in axiomatizations for *normative* reasons. A normative theory is successful to the extent that it convinces decision makers to change the way they make their choices.⁸ A set of axioms, formulated in the language of observable choice, can often convince decision makers that a certain theory has merit more than its mathematical formulation can. Thus, the normative role of axiomatizations is inherently rhetorical.

It is often the case that an axiomatization serves all three purposes. For instance, an axiomatization of utility maximization in consumer theory provides a definition of the term “utility” and shows that binary choices can only define such a utility function up to an arbitrary monotone transformation. This meta-scientific exercise saves us useless debates about the choice between observationally equivalent utility functions.⁹ On a descriptive level, such an axiomatization shows what utility maximization actually entails. This allows one to test the theory that consumers are indeed utility maximizers. Moreover, it renders the theory of utility maximization much more

⁸Changing actual decision making is the ultimate goal of a normative theory. Such changes are often gradual and indirect. For instance, normative theories may first convince scientists before they sift to practitioners and to the general public. Also, a normative theory may change the way people would like to behave even if they fail to implement their stated policies, for, say, reasons of self-discipline. Finally, a normative theory may convince many people that they would like society to make certain choices, but they may not be able to bring them about for political reasons. In all of these cases the normative theories definitely have some success.

⁹This, at least, is the standard microeconomic textbook view. For an opposing view see Beja and Gilboa (1992) and Gilboa and Lapson (1995).

plausible, because it shows that relatively mild consistency assumptions suffice to treat consumers as if they were utility maximizers, even if they are not conscious of any utility function, of maximization, or even of the act of choice. Finally, on a normative level, an axiomatization of utility maximization may convince a person or an organization that, according to their own judgment, they should better adopt a utility function and act so as to maximize it.

Similarly, Savage's (1954) axiomatization of subjective expected utility maximization provides observable definitions of the terms "utility" and "(subjective) probability". Descriptively, it provides us with reasons to believe that there are decision makers who can be described as expected utility maximizers even if they are not aware of it, thus making expected utility maximization a more plausible theory. Finally, from a normative point of view, decision makers who shrug their shoulders when faced with the theory of expected utility maximization may find it more compelling if they are exposed to Savage's axioms.

While our use of the term "axiomatization" highlights the role of providing a link between theoretical concepts and observable data, this term is often used in other meanings. Both in mathematics and in the sciences, an "axiomatization" sometimes refers to a characterization of a certain entity by some of its properties. One is often interested in such axiomatizations as decompositions of the concept that is axiomatized. Studying the "building blocks" of which a concept is made typically enhances our understanding thereof, and allows the study of related concepts. Axiomatizations in the sense used in this book will typically also have the flavor of decomposition of a construct. Thus, on top of the reasons mentioned above, one may also be interested in axiomatizations simply as a way to better understand what characterizes a certain concept and what is the driving force behind certain results. For instance, an axiomatization of utility maximization in consumer theory will often reduce utility theory to more primitive "ingredients", and

will also suggest alternative theories that share only some of these ingredients, such as preferences that are transitive but not necessarily complete.

2.4 Behaviorist, behavioral, and cognitive theories

We distinguish between two types of data. Behavioral data are observations of actions taken by decision makers. By contrast, cognitive data are choice-related observations that are derived from introspection, self-reports, and so forth. We are interested in actions that are open to introspection, even if they are not necessarily preceded by a deliberate decision making process. Thus, what people say about what their choices might be, the reasons they give for actual or hypothetical choices, their recollection of and motivation for such choices are all within the category of cognitive data. In contrast to common psychological usage, we do not distinguish between cognition and emotion. Emotional motives, inasmuch as they can be retrieved by introspection and memory, are cognitive data. Other relevant data, such as certain physiological or neurological activities, will be considered neither behavioral nor cognitive.

Theories of choice can be classified according to the types of data they recognize as valid, as well as by the types of theoretical concepts they resort to. A theory is *behaviorist* if it only admits behavioral data, and if it also makes no use of cognitive theoretical concepts. (See Watson (1913, 1930) and Skinner (1938).) We reserve the term “*behavioral*” to theories that only recognize behavioral data, but that make use of cognitive metaphors. Neo-classical economics and Savage’s (1954) expected utility theory are behavioral in this sense: they only recognize revealed preferences as legitimate data, but they resort to metaphors such as “utility” and “belief”. Finally, *cognitive* theories make use of cognitive data, as well as of behavior data. Typically, they also use cognitive metaphors.

Cognitive and behavioral theories often have behaviorist implications. That is, they may say something about behavior that will be consistent

with some behaviorist theories but not with others. In this case, we refer to the cognitive or the behavioral theory as a *cognitive specification* of the behaviorist theories it corresponds to. One may view a cognitive specification of a behaviorist theory as a description of a mental process that *implements* the theory.

2.5 Rationality

We find that purely behavioral definitions of rationality invariably miss an important ingredient of the intuitive meaning of the term. Indeed, if one adheres to the behavioral definition of rationality embodied in, say, Savage's axioms, one has a foolproof method of making rational decisions: choose *any* prior and *any* utility function, and maximize the corresponding expected utility. Adopting this method, one will never be caught violating Savage's axioms. Yet, few would accept an arbitrary choice of a prior as rational.

It follows that rationality has to do with reasoning as well as with behavior. As a first approximation we suggest the following definition. An action, or a sequence of actions is rational for a decision maker if, when the decision maker is confronted with an analysis of the decisions involved, but with no additional information, she does not regret her choices.¹⁰ This definition of rationality may apply not only to behavior, but also to decision processes leading to behavior. Observe that our definition presupposes a decision making entity capable of understanding the analysis of the problems encountered.

Consider the example of transitivity of binary preferences. Many people, who exhibit cyclical preferences, regret some of their choices when exposed to this fact. For these decision makers, violating transitivity would be considered irrational. Casual observation shows that most people feel embarrassed when it is shown to them that they have fallen prey to framing effects (Tversky and Kahneman (1981)). Hence we would say that, for most people,

¹⁰Alternatively, one can substitute "does not feel embarrassed by" for "does not regret".

rationality dictates that they be immune to framing effects. Observe, however, that regret that follows from unfavorable resolution of uncertainty does not qualify as a test of rationality.

As another example, consider the decision maker's autonomy. Suppose that a decision maker decides on an act, a , and ends up choosing another act, b , because, say, she is emotionally incapable of foregoing b . If she is surprised or embarrassed by her act, it may be considered irrational. But irrationality in this example may be due to the intent to choose a , or to the implicit prediction that she would implement this decision. Indeed, if the decision maker knows that she is incapable of foregoing act b , it would be rational for her to adjust her decisions and predictions to the actual feasible set. That is, if she accepts the fact that she is technologically constrained to choose b , and if she so plans to do, there will be nothing irrational in this decision, and there will be no reason for her to regret not making the choice a , which was never really feasible. Similarly, a decision maker who has limitations in terms of simple mistakes, failing memory, limited computational capacity, and the like, may be rational as long as her decision takes these limitations into account, to the extent that they can be predicted.

Our definition has two properties that we find necessary for any definition of rationality and one that we find useful for our purposes. First, as mentioned above, it relies on cognitive or introspective data, as well as on behavioral data, and it cannot be applied to decision makers who cannot understand the analysis of their decisions. According to this definition it is meaningless to ask whether, say, bees, are rational. Second, it is subjective in nature. A decision maker who, despite all our preaching, insists on making frame-dependent choices, will have to be admitted into the hall of rationality. Indeed, there is too little agreement among researchers in the field to justify the hope for a unified and objective definition of rationality. Finally, our definition of rationality is closely related to the practical question of what should be done about observed violations of classical theories of choice, as we

explain in the sequel. As such, we hope that this definition may go beyond capturing intuition to simplifying the discussion that follows.¹¹

2.6 Deviations from rationality

There is a large body of evidence that people do not always behave as classical decision theory predicts. What should we do about observed deviations from the classical notion of rational choice? Should we refine our descriptive theories or dismiss the contradicting data as exceptions that can only clutter the basic rules? Should we teach our normative theories, or modify them?

We find the definition of rationality given above useful in making these choices. If an observed mode of behavior is irrational for most people, one may suggest a normative recommendation to avoid that mode of behavior. By definition of irrationality, most people would accept this recommendation, rendering it a successful normative theory. By contrast, there is a weaker incentive to incorporate this mode of behavior into descriptive theories: if these theories were known to the decision makers they describe, the decision makers would wish to change their behavior. Differently put, a descriptive theory of irrational choice is a self-refuting prophecy. If, however, an observed mode of behavior is rational for most people, they will stick to it even if theorists preach the opposite. Hence, recommending to avoid it would make a poor normative theory. But then the theorists should include this mode of behavior in their descriptive theories. This would improve the accuracy of these theories even if the theories are known to their subjects.

2.7 Subjective and objective terms

Certain terms, such as “probability”, are sometimes classified as subjective or objective. Some authors argue that all such terms are inherently subjective, and that the term “objective” is but a nickname for subjective terms on

¹¹For other views of rationality, see Arrow (1986), Etzioni (1986), and Simon (1986).

which there happens to be agreement. (See Anscombe and Aumann (1963).) A possible objection is raised by the following example. Five people are standing around a well that they have just found in the field. They all estimate its depth to be more than 100 feet. This is the subjective estimate of each of the five people. Yet, while they all agree on the estimate, they also all agree that there is a more objective way to measure the depth of the well.

Specifically, assume that Judy is one of the five people who have discovered the well, and that she recounts the story to her friend Jerry. Compare two scenarios. In the first scenario, Judy says, “I looked inside, and I saw that it was over 100 feet deep.” In the second scenario she says, “I had dropped a stone into the well and three seconds had passed before I heard the splash.” Jerry is more likely to accept Judy’s estimate in the second scenario than he is in the first. We would also like to argue that Judy’s estimate of the well’s depth in the second scenario is more objective than in the first. This suggests a definition of objectivity that requires more than agreement among subjective judgments: an assessment is *objective* if someone else is likely to agree with it. Generally, one may argue that objective judgments require some justification beyond the perhaps coincidental agreement among subjective ones.

It may be useful to conceptualize the distinction between subjective and objective judgments as quantitative, rather than qualitative. As such, the first definition of objectivity suggests the following measure: an assessment is more objective, the more people happen to share it as their subjective assessment. The definition we propose here offers a different measure: an assessment is more objective, the more likely are people, who have not given thought to it and who have no additional relevant information, to agree with it.

To consider another example, let us contrast classical with Bayesian statistics. If a classical statistician rejects a hypothesis, the belief that it should be rejected is relatively objective. By contrast, a prior of a Bayesian

statistician is more subjective. Indeed, another classical statistician is more likely to agree with the decision to reject the hypothesis than is another Bayesian statistician likely to adopt a prior if it differs from her own. Two Bayesian statisticians may have the same subjective prior, but if they are about to meet a third one, they have no way to convince her to agree with their prior. There is no agreed upon method to adopt a prior, while there are agreed upon methods to test certain hypotheses.¹²

One may argue that our definition of objectivity again resorts to agreement among subjective judgments, only that the latter should include all hypothetical judgments made by various individuals at various times and at various states of the world. But for practical purposes of the following discussion, we find the distinction between objective terms and coincidence of subjective terms quite useful.

3 Meta-Theoretical Prejudices

This section expresses our views on certain issues relating to the state of the art in decision theory as well as to its applications to economics and to other social sciences. These views undoubtedly correlate with our motivation for developing case-based decision theory. Yet, as we explain below, the theory may be suggested and evaluated independently of our prejudices.

3.1 Preliminary remark on the philosophy of science

Analytical philosophy is a science, in the sense that the social sciences are. Namely, it uses formal or formalizable analytical models to describe, explain, and predict certain phenomena (as a descriptive science), or to produce recommendations and prescriptions (as a normative science). Generally, philos-

¹²Two Bayesians will agree on a prior, i.e., will have a common prior, in practice, if this prior conveys minimal information (say, a uniform prior when it is well-defined). This is essentially Laplace's dictum for the case of complete ignorance, which indicates a failure of the Bayesian claim that any prior information may be expressed by a probability function.

ophy is concerned with phenomena that are primarily concentrated in the human mind, such as language, religion, ethics, aesthetics, science, and others. Hence, philosophy is closer to the social or human sciences than it is to the natural ones.

Many social scientists would agree that their theories are no more than suggestive illustrations or rough sketches.¹³ These theories are supposed to provide insights without necessarily being literally true. For instance, the economic model of perfect competition offers important insights regarding the way markets work. It obviously describes a highly idealized reality. That is, it is false. But it is neither worthless nor useless. Naturally, one has to study other models that are more realistic in certain aspects, such as models of imperfect competition, asymmetric information, and so forth. The analysis of these models qualifies and refines the insights of the perfect competition model, but it does not nullify them.

Viewing philosophy as a social science, there is no reason to expect its theories and models to be more accurate or complete than are those of economics or psychology. For example, the logical positivists' Received View offers useful guidelines for developing scientific theories. It is also very insightful when interpreted descriptively, as a model of how science is done. But one expects the Received View to be neither a perfect description of, nor a foolproof prescription for scientific activity. Hanson's theory that observations are theory-laden (Hanson (1958)), for example, also provides profound insights and enhances our understanding of the scientific process. It qualifies, but does not nullify the insights gained from the Received View. Allowing ourselves a mild exaggeration, we argue for humility in evaluating theories in the social sciences, philosophy included: ask not, is there a counterexample to a theory; ask, is there an example thereof.¹⁴¹⁵

¹³An extreme term, which was suggested by Gibbard and Varian (1978), is "caricatures".

¹⁴In this sense, our view is reminiscent of verificationism. For a review, see, for instance, Lewin (1996).

¹⁵It is inevitable that similar qualifications would apply to this very paragraph.

3.2 Utility and expected utility “theories” as conceptual frameworks and as theories

As mentioned in Subsection 2.1, in common usage the term “theory” often defines only a conceptual framework. Consider, for instance, utility theory, suggesting that, given a set of feasible alternatives, decision makers maximize a utility function. When an intended application is explicitly or implicitly given, this is a descriptive theory in the tradition of logical positivism, which would also pass Popper’s falsifiability test. The theoretical term “utility” is axiomatically derived from observations. The axioms define quite precisely the conditions under which the theory is falsified. It can therefore be tested experimentally. However, when the application is neither obvious nor agreed upon, it is not as clear what constitutes a violation of the axioms, or a falsification of the theory. In fact, almost any case in which the theory seems to be refuted can be re-interpreted in a way that conforms to the theory. One may always enrich the model, redefining the alternatives involved, in such a way that apparent refutations are reconciled.

Utility theory is a conceptual framework in the language of Subsection 2.1 above. There is no canon of application that would state precisely what are the real world phenomena to which utility theory may be applied. In certain situations, such as carefully designed laboratory experiments, there is only one reasonable mapping from the formal structure to observed phenomena. But in many real life situations one has a certain degree of freedom in choosing this mapping. It is then not obvious what counts as a refutation of the theory. Thus, utility maximization may not be a theory in the Popperian sense. Rather, it is a conceptual framework, offering a certain way of conceptualizing and organizing data. Working within this conceptual framework, one may suggest specific theories that also define the set of alternatives among which choice is made.

Similarly, expected utility theory seems to be well defined and clearly falsifiable when one accepts a certain application. Indeed, in certain labo-

ratory experiments it can be directly tested. But most empirical data do not present themselves with a clear definition of states of the world or of outcomes. Hence apparent violations of the theory may be accounted for by re-interpreting these terms.

The “theories” of utility maximization and of expected utility maximization therefore serve a dual purpose: they are formulated as testable theories, but they are also used as conceptual frameworks for the generation of more specific theories. We believe that this dual status is useful. On the one hand, treating these conceptual frameworks as falsifiable theories with axiomatic foundations facilitates the scientific discussion and ensures that we do not engage in debates between observationally equivalent theories. On the other hand, decision theory would be very limited in its applications if one were to insist that these conceptual frameworks be used only when there is no ambiguity regarding their interpretation. It is therefore useful to insist that every conceptual framework be grounded in a falsifiable axiomatization, as if it were a specific theory, yet to use the framework also beyond the realm in which falsification is unambiguously defined.¹⁶

Moreover, many of the theoretical successes of utility theory and of expected utility theory involve situations in which one cannot reasonably hope to test their axioms. For instance, the First Welfare Theorem (Arrow and Debreu (1954)) provides a fundamental insight regarding the optimality of competitive equilibria, even though we cannot measure the utility functions of all individuals involved. Similarly, economic and game theoretic analysis that relies on expected utility maximization offers important insights and recommendations regarding financial markets, markets with incomplete information, and so forth, while neither subjective probabilities nor utility functions can practically be measured. Again, we believe that one should make sure that all theoretical terms are in-principle observable, but allow

¹⁶Historically, conceptual frameworks often originate from theories. A formal theory may be suggested with an implicit understanding of its intended applications, but it might later be generalized and evolve into a conceptual framework.

oneself some freedom in using these terms even when actual observations are impossible.

3.3 On the validity of purely behavioral economic theory

The official position of economic theory, as reflected in many graduate (and undergraduate) level courses and textbooks, is that economics is and should be purely behavioral. Economists do not care what people say, they only care what they do. Moreover, it is argued, people are often not aware of the way they make decisions, and their verbal introspective account of the decision process may be more confusing than helpful. As long as they behave *as if* they follow the theory, we should be satisfied.

One can hardly take issue with this approach if the underlying axioms are tested and proved true, or if the theory generates precise quantitative predictions that conform to reality. But economics often uses a theory such as expected utility maximization in cases where its axioms cannot be tested directly, and its predictions are at best qualitative. That is, theories are often used as rhetorical devices. In these cases, one cannot afford to ignore the intuitive appeal and cognitive plausibility of behavioral theories. The more cognitively plausible a theory is, the more effective rhetorical device will it be.

Many economists would argue that the cognitive plausibility of a theory follows from that of its axioms. For instance, Savage's behavioral axioms are very reasonable; thus it is very reasonable that people should and would behave *as if* they were expected utility maximizers. However, we claim that behavioral axioms that appear plausible when formulated in a given language may not be as convincing when this language is unnatural. For instance, Savage's axiom P2 (often referred to as the "sure-thing principle") is very compelling when acts are given as functions from states to outcomes. But in examples in which outcomes and states are not naturally given in the

description of the problem, it is not clear what all the implications of the sure-thing principle are, namely, how preferences among acts are constrained by the sure-thing principle. It is therefore not at all obvious that actual behavior would follow this seemingly very compelling principle. More generally, the predictive validity of behavioral axioms is not divorced from the cognitive plausibility of the language in which they are formulated.

Moreover, when there is no obvious mapping between the reality modeled and the mathematical model, it is no longer clear that the theory is indeed behavioral. For example, Savage's model assumes that choices between acts are observable. This assumption is valid when one chooses between bets on, say, the color of a ball drawn from an urn. In this case it is quite clear that the states of the world correspond to the balls in the urn, and these states are observable by the modeler. But in many problems to which the theory is applied, states of the world need to be defined, and there is more than one way to define them. In these cases one cannot be sure that the states that are supposedly conceived of by the decision maker are those to which the formal model refers. The need to refer to states that exist in the decision maker's mind renders the theory's data cognitive rather than behavioral. As a matter of principle, a behavioral theory has to rely on data that are observable by an outside observer without ambiguity.

In conclusion, the position that cognitive plausibility of economic or choice theories is immaterial seems untenable. A behaviorist theory, treating a decision maker as a "black box", has to be supported by a convincing cognitive specification. On the other hand, a seemingly compelling mental account of a decision making process is also unsatisfactory as long as we do not know what types of behavior it induces. Thus, a satisfactory theory of choice should tell a convincing story both about the cognitive process and about the resulting behavior.

3.4 What does all this have to do with CBDT?

The motivation for the development of case-based decision theory has to do with our dissatisfaction with classical decision theory. Specifically, we find that expected utility theory is not always cognitively plausible, and that the behavioral approach, often quoted as its theoretical foundation, is untenable. Moreover, we believe that expected utility theory is not always successful as a normative theory, because it is often impractical. And in those cases where one cannot follow expected utility theory, it is neither irrational, nor even boundedly rational, to deviate from Savage's postulates. Yet, the theory we present in this book is independent of these prejudices. In particular, the following chapters may be read also by people who believe only in purely behavioral theories, or who are only willing to consider CBDT as a descriptive theory of bounded rationality.

Chapter 2

Decision Rules

4 Elementary Formula and Interpretations

Expected utility theory enjoys the status of an almost unrivaled paradigm for decision making in face of uncertainty. Relying on such sound foundations as the classical works of Ramsey (1931), de Finetti (1937), von Neumann and Morgenstern (1944), and Savage (1954), the theory has formidable power and elegance, whether interpreted positively or normatively, for situations of given probabilities (“risk”) or unknown ones (“uncertainty”) alike.

While evidence has been accumulating that the theory is too restrictive (at least from a descriptive viewpoint), its various generalizations only attest to the strength and appeal of the expected utility paradigm. With few exceptions, all suggested alternatives retain the framework of the model, relaxing some of the more demanding axioms while adhering to the more basic ones. (See Machina (1987), Karni and Schmeidler (1991), Camerer and Weber (1992), and Harless and Camerer (1994) for extensive surveys.)

Yet it seems that in many situations of choice under uncertainty, the very language of expected utility models is inappropriate. For instance, states of the world are neither naturally given, nor can they be simply formulated. Furthermore, sometimes even a comprehensive list of all possible outcomes is not readily available or easily imagined. The following examples illustrate

these points.

4.1 Motivating examples

Example 2.1 *As a benchmark, we first consider Savage’s famous omelet problem (Savage 1954, pp. 13-15): Leonard Savage is making an omelet using six eggs. Five of them are already cracked into a bowl. He is holding the sixth, and has to decide whether to crack it directly into the bowl, or into a separate, clean dish to examine its freshness. This is a decision problem under uncertainty, because Leonard Savage does not know whether the egg is fresh or not. Moreover, uncertainty matters: if he knew that the egg is fresh, he would be better off cracking it directly into the bowl, saving the need to wash another dish. On the other hand, a rotten egg would result in losing the five eggs already in the bowl; thus, if he knew that the egg were not fresh, he would prefer to crack it into the clean dish.*

In this example, uncertainty may be fully described by two states of the world: “the egg is fresh” and “the egg isn’t fresh”. Each of these states “resolves all uncertainty” as prescribed by Savage (1954). Not only are there relatively few relevant states of the world in this example, they are also naturally given in the description of the problem. In particular, they can be defined independently of the acts available to the decision maker. Furthermore, the possible outcomes can be easily defined. Thus this example falls neatly into decision making under uncertainty in Savage’s model.

Example 2.2 *A couple has to hire a nanny for their child. The available acts are the various candidates for the job. The decision makers do not know how each candidate would perform if hired. For instance, each candidate may turn out to be negligent or dishonest. Coming to think about it, they realize that other problems may also occur. Some nannies are treating children well, but cannot be trusted with keeping the house in order. Others appear to be*

just perfect on the job, but are not very loyal and may quit the job on short notice.

The couple is facing uncertainty regarding the candidates' performance on several measures. However, there are several difficulties in fitting this problem into the framework of expected utility theory (EUT). First, imagining all possible outcomes is not a trivial task. Second, the states of the world do not naturally suggest themselves in this problem. Furthermore, should the decision makers try to construct them analytically, their number and complexity would be daunting: every state of the world should specify the exact performance of each candidate on each measure.¹

Example 2.3 *President Clinton has to decide on military intervention in Bosnia-Herzegovina². The alternative acts are relatively clear: one may do nothing, impose economic sanctions, use limited military force (say, only air strikes) or opt for a full-blown military intervention. The main problem is to decide what are the likely short-run and long-run outcomes of each act. For instance, it is not exactly clear how strong are the military forces of the warring factions in Bosnia; it is hard to judge how many casualties each military option would involve, and what would be the public opinion response; there is some uncertainty about the reaction of Russia, especially if it goes through a military coup.*

In short, the problem is definitely one of decision under uncertainty. But, again, neither all possible eventualities, nor all possible scenarios are readily

¹It may be easier to assess distributions of utility for each candidate and to then apply EUT to these distributions than to spell out all the states of the world. This proves our point, namely, that the state model is unnatural in this example. It follows that one cannot use de Finetti's or Savage's axiomatic derivations of subjective probability as foundations for the assessment of distributions in such examples.

²USA military involvement in parts of what was Yugoslavia is on the political agenda in the USA since we wrote the first version of our first paper on CBDT (March 1992) till now (September 2000).

available. Any list of outcomes or of states is bound to be incomplete. Furthermore, each state of the world should specify the result of each act at each point of time. Thus, an exhaustive set of the states of the world certainly does not naturally pop up.

In example 2.1, expected utility theory seems a reasonable description of how people think about the decision problem. By contrast, we argue that in examples such as 2.2 and 2.3, EUT does not describe a plausible cognitive process. Should the decision maker attempt to think in the language of EUT, she would have to imagine all possible outcomes and all relevant states. Often the definition of states of the world would involve conditional statements, attaching outcomes to acts. Not only would the number of states be huge, the states themselves would not be defined in an intuitive way.

Moreover, even if the agent managed to imagine all outcomes and states, her task would by no means be done. Next she would have to assess the utility of each outcome, and to form a prior over the state space. It is not clear how the utility and the prior are to be defined, especially since past experience appears to be of limited help in these examples. For instance, what is the probability that a particular candidate for the job in example 2.2 will end up being negligent? Or being both negligent and dishonest? Or, considering example 2.3, what are the chances that a military intervention will develop into a full-blown war, while air strikes will not? What is the probability that a scenario that no expert predicted will eventually materialize?

It seems unlikely that decision makers can answer these questions. Expected utility theory does not describe the way people actually think about such problems. Correspondingly, it is doubtful that EUT is the most useful tool for predicting behavior in decision problems of this nature. A theory that will provide a more faithful description of how people think would have a better chance of predicting what they will do. How, then, do people think about such decision problems? We resort to Hume (1748), who argued that “From causes which appear *similar* we expect similar effects. This is the sum

of all our experimental conclusions.” That is, the main reasoning technique that people use is drawing analogies between past cases and the one at hand.³

Applying this idea to decision making, we suggest that people choose acts based on their performance in similar problems in the past. For instance, in example 2.2, a common, and indeed very reasonable thing to do is to ask each candidate for references. Every recommendation letter provided by a candidate attests to his/her performance (as a nanny) in a different situation or “problem”. In this example, the agents do not rely on their own memory; rather, they draw on the experience of other employers. Each past “case” would be judged for its similarity; for instance, serving as a nanny to a month-old toddler is somewhat different from the same job when a two-year old child is concerned. Similarly, the house, neighborhood and other factors may affect the relevance of past cases to the problem at hand. Thus we expect decision makers to put more weight on the experience of people whose decision problem was more similar to theirs. Furthermore, they may rely more heavily on the experience of people they happen to know, or judge to have similar tastes to their own.

Next consider example 2.3. While military and political experts certainly do try to write down possible scenarios and to assign likelihoods to them, this is by no means the only reasoning technique used. (Nor is it necessarily the most compelling a-priori or the most successful a-posteriori.) Very often people’s reasoning employs analogies to past cases. For instance, proponents of military intervention tend to cite the Gulf War as a “successful” case. They stress the similarity of the two problems, say, as local conflicts in post-

³We were first exposed to this idea as an explicit “theory” in the form of “case-based reasoning” (Schank (1986), Riesbeck and Schank (1989)), to which we owe the epithet “case-based”. (See also Kolodner and Riesbeck (1986) and Kolodner (1988).) Needless to say, our thinking about the problem was partly inspired by case-based reasoning. At this stage, however, there does not seem to be much in common between our theory and case-based reasoning, beyond Hume’s basic idea. It should be mentioned that similar ideas were also expressed in the economic literature by Keynes (1921), Selten (1978), and Cross (1983).

cold-war world. Opponents adduce the Vietnam War as a case in which military intervention is generally considered to have been a mistake. They also point to the similarity of the cases, for instance to the “peace-keeping mission” mentioned in both.

Case-based decision theory (CBDT) attempts to formally capture this mode of reasoning as it applies to decision making. In the next subsection we begin by describing the simplest version of CBDT. This version is admittedly rather restrictive. It will be generalized later on to encompass a wider class of phenomena. At this point we only wish to highlight some features of the general theory, which are best illustrated by the simplest version presented below.

4.2 Model

Assume that a set of problems is given as primitive, and that there is some measure of similarity on it. The problems are to be thought of as descriptions of choice situations, as stories involving decision problems. Generally, a decision maker would remember some of the problems that she and other decision makers encountered in the past. When faced with a new problem, the similarity of the situation brings this memory to mind, and with it the recollection of the choice made and the outcome that resulted. We refer to the combination of these three, the problem, the act, and the result, as a “case”. Thus similar cases are recalled, and based on them each possible decision is evaluated. The specific model we propose here evaluates each act by the sum, over all cases in which it was chosen, of the product of the similarity of the problem to the one at hand and the resulting utility.

Formally, we start with three sets: let P be a set of decision *problems*, A – a set of *acts* that may be chosen at the current problem, and R – a set of possible *outcomes*. A *case* is a triple (q, a, r) where q is a problem, a is an act and r – an outcome. Thus, the set of *conceivable cases* is the set of all such triples:

$$C \equiv P \times A \times R .$$

Observe that C is not the set of cases that actually have occurred. Moreover, it will typically be impossible for all cases in C to co-occur, because different cases in C may attach a different act or a different outcome to a given decision problem. The set of cases that are known to have occurred will thus be a subset of C .

The next two components of the formal model are similarity and utility functions. The *similarity function*

$$s : P \times P \rightarrow [0, 1]$$

is assumed to provide a quantification of similarity judgments between decision problems. The concept of similarity is the main engine of the decision models introduced in this book. Performing similarity judgments is the chief cognitive task of the decision maker we have in mind. Hume (1748) has already suggested similarity as a basis for inductive reasoning. Formal models of similarity date back at least to Quine (1969b) and Tversky (1977),⁴ and much attention has been given to analogical reasoning. (See, for instance, Gick and Holyoak (1980, 1983), and Falkenhainer, Forbus, and Gentner (1989).) However, we are unaware of a formal theory of decision making that is explicitly based on similarity judgments.

As is normally the case with theoretical concepts, the exact meaning of the similarity function will be given by the way in which it is employed. The way we use similarity between problems, to be specified shortly, implies uniqueness only up to multiplication by a positive number.

The term “similarity” should not be taken too literally. Past decision problems affect the decision maker’s choice only if they are recalled. For

⁴The classical notion of similarity, as employed in Euclidean geometry, suggests that similarity be modeled as an equivalence relation. Obviously, a real-valued similarity function can capture such a notion of similarity. But it allows similarity relations that are intransitive or asymmetric, and, importantly, it can also capture gradations of similarity.

instance, assume that a voter named John is asked to voice his view regarding military intervention in Bosnia. It is possible that, *if asked*, John would judge the Korean War to be similar to the situation in Bosnia. But it is also possible that John would not be reminded of the Korean War on his own. For our purposes, John may be described as attaching a zero similarity value to the two decision problems. Thus, while we use the suggestive term “similarity”, we think of this function as representing awareness and probability of recall as well as conscious similarity assessments.

Our formulation assumes that similarity judgments cannot be negative. This may sometimes be too restrictive. For instance, Sarah may know that she and Jim always disagree about movies. If Sarah knows that Jim liked a movie, she will take it as a piece of evidence against watching this movie. This may be thought of as if Sarah’s problem has a negative similarity to a decision problem in which Jim was involved. However, for concreteness we choose to restrict the discussion to non-negative similarity values at this point.⁵

This section restricts similarity judgments to decision problems. We later extend the formal model to include similarity judgments between pairs of problems and acts, and even between entire cases.

The decision maker is also characterized by a *utility function*:

$$u : R \rightarrow \mathbb{R} .$$

The utility function measures the desirability of outcomes. The higher is the value of this function, the more desirable is the outcome considered to be. Moreover, positive u -values can be associated with positive experiences, which the decision maker would like to repeat, whereas negative u -values correspond to negative experiences, which the decision maker would rather

⁵This restriction may pose a problem in our example. It can be solved by assuming that Sarah assigns negative utility values to outcomes that Jim likes. The utility function is introduced below.

avoid. As in the case of the similarity function, the exact measurement of desirability of positive outcomes and of the undesirability of negative ones will be tied to the particular formula in which this function is employed.

We now proceed to specify the decision model. A decision maker is facing a problem $p \in P$. We assume that she knows the possible courses of action she might take. These are denoted by A . We do not address the question of specification of the acts that are, indeed, available in a given decision problem, or of identification of the decision problem in the first place.⁶

The decision maker can base her decision on the cases she knows of. We refer to the set consisting of these cases as the decision maker's memory. Formally, *memory* is a subset of cases:

$$M \subset C .$$

Whereas C is the set of all conceivable cases, M represents those that actually occurred, and that the decision maker was informed of. Cases in which the decision maker was the protagonist would naturally belong to M . But so will cases in which other people were the protagonists, if they were related to our decision maker, reported in the media, and so forth. We assume that each decision problem $q \in P$ may appear in M at most once. This will naturally be the case if the description of a decision problem is detailed enough. For instance, if the identity of the decision maker and the exact time of decision are part of the problem's description, no problem would repeat itself precisely. Adopting this assumption allows us to use set notation and to obviate the need for indices. This notation involves no loss of generality, because a problem that does repeat itself in an identical fashion can be represented as several problems that happen to have the same features, though not the same formal description.

⁶In Chapter 5, however, we discuss the way a decision maker develops plans. As such, it may be viewed as an attempt to model the process by which decision makers generate the set of acts available to them.

The most elementary version of CBDT offers that the decision maker would rank available acts according to the similarity-weighted sum of utilities they have resulted in the past. Formally, a decision maker with memory M , similarity function s , and utility function u , who now faces a new decision problem p , will rank each act $a \in A$ according to

$$U(a) = U_{p,M}(a) = \sum_{(q,a,r) \in M} s(p,q)u(r) , \quad (*)$$

(where the summation over the empty set is taken to yield zero), and will choose an act that maximizes this sum. Observe that, given memory M , the function U only uses the u values of outcomes r that have appeared in M .

For each act $a \in A$ and each case $c = (q, a, r) \in M$ in which act a was indeed chosen, one may view the product $s(p,q)u(r)$ as the effect that case c has on the evaluation of act a in the present problem p . If act a has resulted in case c in a desirable outcome r , namely an outcome such that $u(r) > 0$, having case c in memory would make act a more attractive in the new problem as well. Similarly, recalling case c in which act a has resulted in an undesirable outcome r , that is, an outcome for which $u(r) < 0$, would render act a less attractive in the current decision problem. In both cases, the impact that case c would have on the way act a is viewed depends on the similarity of the problem at hand to the problem in case c . Finally, the overall evaluation of act a is obtained by summing up the products $s(p,q)u(r)$ over all cases of the form (q, a, r) , namely, over all cases in which a was chosen in the past.

Let us consider example 2.3 again. President Clinton has to choose among several acts, such as “send in troops”, “use air strikes”, “do nothing”, and so forth. The act “send in troops” would be evaluated based on its performance in the past. The Vietnam War would probably be a case in which this act yielded an undesirable outcome. The Gulf War is another relevant case, in which the same act resulted in a desirable outcome. One has to assess the

degree of desirability or undesirability of each outcome, as well as the degree to which each past problem is similar to the problem at hand. One then multiplies the desirability values by the corresponding similarity values. The summation of these products will constitute the overall evaluation of the act, which may turn out to be positive or negative. An alternative act, “do nothing” might be judged based on its performance in various cases where no intervention was chosen, and so forth.

The formula (*) may be viewed as modeling learning. It attempts to capture the way in which experience affects decision making, and it suggests a dynamic interpretation: with the addition of new cases, memory grows. As we will see in Chapter 7, adding cases to memory is but the simplest form of learning. Yet, it appears to be the basis for other forms of learning as well.

4.3 Aspirations and satisficing

The presentation of the elementary formula above highlights the distinction between desirable outcomes, whose utility values are positive, and undesirable ones, to which negative utility values are attached. While this distinction is rather intuitive, classical decision theory typically finds it redundant: higher utility values are considered more desirable, or, equivalently, less undesirable. All utility indices are used only in the context of pairwise comparisons, and their values have no meaning on any absolute scale. Indeed, most classical theories of decision allow the utility function to be “shifted” by adding a constant to all values, without changing the theory’s observable implications.

This is not the case with the elementary formula (*). One may easily observe that shifting the utility function by a constant will generally result in a different prediction of choice. Indeed, shifting the function u by a constant c would change $U(a)$ by $c \cdot \sum_{(q,a,r) \in M} s(p, q)$. Since there is no reason that $\sum_{(q,a,r) \in M} s(p, q)$ will be the same for different acts a , such a shift is not guaranteed to preserve the U -ranking of the acts. Specifically, if $c > 0$, such

a shift would favor acts that were chosen in the past in similar problems, whereas a shift by $c < 0$ would favor acts that have not been chosen in similar problems.

It follows that the choice of the reference point zero on the utility scale cannot be arbitrary. This suggests that our intuitive distinction between desirable and undesirable outcomes might have a behavioral meaning as well. Indeed, consider a simple case in which there is no uncertainty. That is, whenever an act $a \in A$ is chosen, a particular outcome $r_a \in R$ results. The decision maker is not assumed to know this fact. Moreover, she certainly does not know what utility value corresponds to which act. All she knows are the cases she has experienced, and she follows U -maximization as in (*).

Starting out with an empty memory, all acts are assigned a zero U -value. At this point the decision maker's choice is arbitrary. Assume that she chose an act a that resulted in an undesirable outcome r_a , that is, an outcome with $u(r_a) < 0$. Let us now consider the same decision maker confronted with the next problem. Suppose that this problem bears some similarity to the first one. In this case, act a will have a negative U -value ($u(r_a)$ multiplied by the similarity of the two problems), whereas all other acts, which have empty histories, still have a U -value of zero. Thus the one act that has been tried, and that has resulted in an undesirable outcome, is the only one the decision maker has some information about, and she tries to veer away from this act. Her choice among the other ones will be arbitrary again. Suppose she chooses act b . If $u(r_b) < 0$, then, in the next similar problem, she will find that both a and b are unattractive acts, and will choose (in an arbitrary fashion) among the other ones, which have not yet been tried.

This process may continue until the decision maker has a negative experience with each and every possible act. In this situation she would have to choose the lesser of these evils and opt for an act whose U -value is highest, that is, least negative. Assume, however, that the decision maker finds an act, say, d , that leads to a desirable outcome, namely, an act such that

$u(r_d) > 0$. In this case, in the next similar problem act d would have a positive U -value, whereas all other acts would have non-positive U -values. Thus d will get to be chosen again. Since act d always leads to the same desirable outcome r_d , its U -value will always remain positive, and d will always be chosen. The decision maker will never attempt any act other than d , despite the fact that many other acts might not have been tried even once.

One might view this mode of behavior as *satisficing* in the sense of Simon (1957) and March and Simon (1958). Taking the value zero on the utility scale to be the decision maker's *aspiration level*, the decision maker would cling to an act that achieves this value without attempting other acts and without experimentation. It is only when the decision maker's current choice is evaluated below the aspiration level that the decision maker is "unsatisfied" and is prodded to experiment with other options.

Aspiration levels are likely to change as a result of past experiences.⁷ One might therefore wish to explicitly introduce notation that keeps track of these changes. Let u_M denote the utility function given memory M , restricted to outcomes that have appeared in M . Assume that (for all M and M') u_M and $u_{M'}$ differ by an additive constant on the intersection of their domains. We may then choose a utility function \hat{u} and define the aspiration level given memory M , H_M , such that $u_M = \hat{u} - H_M$. In this case, we can reformulate (*) as follows. The decision maker is maximizing

$$U(a) = U_{p,M}(a) = \sum_{(q,a,r) \in M} s(p,q)[\hat{u}(r) - H_M] . \quad (*')$$

To see the role of the aspiration level, consider the following example. Suppose that act a was chosen 10 times, yielding a payoff (\hat{u}) of 1 each time. Act b , by contrast, was only tried twice, yielding a payoff of 4 each time.

⁷Indeed, experience may and generally does shape both similarity judgments and desirability evaluations in more general ways. At this point we focus only on shifts of the utility function.

Assume that all past problems are equally similar to the one at hand. If the aspiration level is zero, one may equate the utility u with the payoff \hat{u} , and thus the U -value of a exceeds that of b . Next, assume that, with the same payoff function \hat{u} , the aspiration level is 2 rather than 0. This makes all the outcomes of act a undesirable, whereas those corresponding to b are still desirable. Hence, b will be preferred to a .

Consider now a dynamic process of aspiration level adjustment. Assume, as above, that a decision maker starts out with an aspiration level of 0, and considers all outcomes of both acts a and b to be desirable. Having twice observed that 4 is a possible payoff, however, the decision maker considers the payoff 1 to be less satisfactory than she used to. In other words, her aspiration level rises. As we saw above, a large enough increase in the aspiration level would render b more desirable than a according to (*').

In general, the function U , being cumulative in nature, may give preference to past choices that yielded desirable outcomes simply because they were chosen many times. Among those acts that, overall, yielded desirable outcomes, U -maximization does not opt for the highest average performance. It may even exhibit habit formation. But this conservative nature of U -maximization is mitigated by the process of aspiration level adjustment. We devote Section 18 to this issue, and show that reasonable assumptions on the aspiration level adjustment process lead to optimal choices in situations where the same problem is encountered in exactly the same way.

4.4 Comparison with EUT

In CBDT, as in EUT, acts are ranked by weighted sums of utilities. Indeed, the formula (*) so resembles that of expected utility theory that one may suspect CBDT to be no more than EUT in a different guise. However, despite appearances, the two theories have little in common. First, note some mathematical differences between the formulae. In (*) there is no reason for the coefficients $s(p, \cdot)$ to add up to 1 or to any other constant.

More importantly, while in EUT every act is evaluated at every state, in U -maximization each act is evaluated over a different set of cases. To be precise, if $a \neq b$, the set of elements of M summed over in $U(a)$ is disjoint from that corresponding to $U(b)$. In particular, this set may well be empty for some a 's.

On a more conceptual level, in expected utility theory the set of states is assumed to be an exhaustive list of all possible scenarios. Each state “resolves all uncertainty”, and, in particular, attaches a result to each available act. By contrast, in case-based decision theory the memory contains only those cases that actually happened. Correspondingly, the utility values that are used in (*) are only those that were actually experienced. To apply EUT one needs to engage in hypothetical reasoning, namely to consider all possible states and the outcome that would result from each act in each state. To apply CBDT, no hypothetical reasoning is required.

As opposed to expected utility theory, CBDT does not distinguish between certain and uncertain acts. In hindsight, a decision maker may observe that a particular act always resulted in the same outcome (i.e., that it seems to involve no uncertainty), or that it is “uncertain” in the sense that it resulted in different outcomes in similar problems. But the decision maker is not assumed to know a-priori which acts involve uncertainty and which do not. Indeed, she is not assumed to know anything about the outside world, apart from past cases.

CBDT and EUT also differ in the way they treat new information and evolve over time. In EUT new information is modeled as an event (a subset of states) that has obtained. The model is restricted to this event and the probability is updated according to Bayes' rule. By contrast, in CBDT new information is modeled primarily by adding cases to memory. In the basic model, the similarity function calls for no update in face of new information. Thus, EUT implicitly assumes that the decision maker was born with knowledge of and beliefs over all possible scenarios, and her learning con-

sists of ruling out scenarios that are no longer possible. On the other hand, according to CBDT the decision maker was born completely ignorant, and she learns by expanding her memory.⁸ Roughly, an EUT agent learns by observing what *cannot* happen, whereas a CBDT agent learns by observing what *can*.

We do not consider case-based decision theory to be better than or a substitute for expected utility theory. Rather, we view the two theories as complementary. While these theories are justified by supposedly behavioral axiomatizations, we believe that their scope of applicability may be more accurately delineated if we attempt to judge the psychological plausibility of the various constructs. Two related criteria for classification of decision problems may be relevant. One is the problem’s description, the second is its relative novelty.

If a problem is formulated in terms of probabilities, EUT is certainly a natural choice for analysis and prediction. Similarly, when states of the world are naturally defined, it is likely that they will be used in the decision maker’s reasoning process, even if a (single, additive) prior cannot be easily formed. However, when neither probabilities nor states of the world are salient (or easily accessible) features of the problem, CBDT may be more plausible than EUT.

We may thus refine Knight’s dichotomous distinction between risk and uncertainty (Knight (1921)) by introducing a third category of *structural ignorance*: “risk” refers to situations where probabilities are given; “uncertainty” – to situations in which states are naturally defined, or can be simply constructed, but probabilities are not. Finally, decision under “structural ignorance” refers to decision problems for which states are neither (i) naturally given in the problem; nor (ii) can they be easily constructed by the decision maker. EUT is appropriate for decision making under risk. In face

⁸In Chapter 7 we discuss other forms of learning, including changes of the similarity judgments.

of uncertainty (and in the absence of a subjective prior) one may still use those generalizations of EUT that were developed to deal with this problem specifically, such as non-additive probabilities (Schmeidler (1989)⁹) and multiple priors (Bewley (1986)¹⁰, Gilboa and Schmeidler (1989)). However, in cases of structural ignorance CBDT is a viable alternative to the EUT paradigm.

Classifying problems based on their novelty, one may consider three categories. When a problem is repeated frequently enough, such as whether to stop at a red traffic light, the decision becomes almost automated. Very little thinking is involved in making such a decision, and it may be best described as rule-based. When deliberation is required, but the problem is familiar enough, such as whether to buy insurance, it can be analyzed in isolation. In these situations the history of the same problem suffices for the formulation of states of the world and perhaps even of a prior, and EUT (or some generalization thereof) may be cognitively plausible. Finally, if the problem is unfamiliar, such as whether to get married or to invest in a politically unstable country, it needs to be analyzed in a context- or memory-dependent fashion, and CBDT is a more accurate description of the way decisions are made.

Thus, rule-based systems are the simplest description of decision making when the decision maker is not aware of uncertainty or even of the fact that she is making a decision.¹¹ Expected utility theory offers the best description of decisions in which uncertainty is present and the decision maker has enough data to analyze it. By contrast, case-based decision theory is probably the most natural description of decision making when the decision problem is amorphous and there are insufficient data to analyze it properly.

⁹See also Gilboa (1987) and Wakker (1989).

¹⁰See also Aumann (1962).

¹¹Such decisions may also be viewed as a special type of case-based decisions. Specifically, a rule can be thought of as a summary of many cases, from which it was probably derived in the first place. See Section 13 for a more detailed discussion.

We defer a more detailed discussion of CBDT and EUT until Chapter 4, following the axiomatic derivation of CBDT in Chapter 3.

4.5 Comments

Subjective Similarity As will be shown in Chapter 3, the similarity function in our model may be derived from observed preferences. Since different people may have different preferences, we should expect that the derived similarity functions would also differ across individuals. The similarity function is therefore subjective, as is probability in the works of de Finetti (1937) and Savage (1954). Yet, for some applications one may find that the data suggest an agreed-upon way to measure similarity, which may be viewed as objective.

Cumulative Utility In the description of CBDT above, we advance a certain cognitive interpretation of the function u . We assume that it represents fixed preferences, and that memory may affect choices only by providing information about the u -value of the outcomes that acts yielded in the past. However, one may suggest that memory has a direct effect on preferences. According to this interpretation, the utility function is the aggregate U , while the function u describes the way in which U changes with experience. For instance, if the decision maker has a high aspiration level, corresponding to negative u values, she will like an option less, the more she used it in the past, and will exhibit change seeking behavior. On the other hand, a low aspiration level, corresponding to positive u values, would make her evaluate an option more highly, the more she is familiar with it, and would result in habit formation. Chapter 6 is devoted to this interpretation.

Hypothetical Cases Consider the following example. Jane has to drive to the airport and she can choose between road A and road B. She chooses road A and arrives at the airport on time. On the way to the airport, however, she learns that road B was closed for construction. A week later Jane is faced with the same problem. Regardless of her aspiration level, it seems

obvious that she will choose the road A again. (Road constructions, at least in psychologically plausible models, never end.)

This example shows that relevant cases may also be hypothetical, or counterfactual. More explicitly, Jane’s reasoning probably involves a counterfactual proposition such as “Had I taken road B, I would never have made it.” This may be modeled as a hypothetical case in which she took road B and arrived at the airport with a considerable delay. Hypothetical cases may endow a case-based decision maker with reasoning abilities she would otherwise lack. Moreover, it seems that any knowledge the agent possesses and any conclusions she deduces from it can, inasmuch as they are relevant to the decision at hand, be reflected by hypothetical cases.

5 Variations and Generalizations

5.1 Average similarity

The previous chapter introduced the basic decision criterion of CBDT, namely maximization of the function

$$U(a) = U_{p,M}(a) = \sum_{(q,a,r) \in M} s(p,q)u(r) . \quad (*)$$

This function is cumulative, summing up the impact of past cases. Consequently, the number of times a certain act was chosen in the past affects its perceived desirability. As mentioned above, according to the function U , an act a that was chosen relatively many times, yielding low but positive payoffs, may be preferred to an act b that was chosen less frequently, yielding consistently high payoffs. Moreover, this is true even if both acts were chosen frequently enough to allow statistical estimation of their payoffs means, or even if no uncertainty is present.

We have argued that in presence of high and low payoffs, the aspiration level is likely to take an intermediate value, making the low payoffs appear negative. While this may make a U -maximizer prefer b over a in the

example above, one may still wonder whether the function U is the most reasonable way to evaluate alternatives. An obvious variation would be to use a similarity-based “average” utility function, namely

$$V(a) = \sum_{(q,a,r) \in M} s'(p,q)u(r)$$

where

$$s'(p,q) = \frac{s(p,q)}{\sum_{(q',a,r) \in M} s(p,q')} .$$

if $\sum_{(q',a,r) \in M} s(p,q') > 0$ and 0 otherwise.

V -maximization may be viewed as a way to formalize the idea of frequentist belief formation (insofar as it is reflected in behavior). Although beliefs and probabilities do not explicitly exist in this model, in some cases they may be implicitly inferred from the weights s' . That is, if the decision maker happens to choose the same act in many similar cases, the evaluation function V may be interpreted as gathering statistical data, or as forming a “frequentist” prior. Observe that CBDT does not presuppose any a priori beliefs. Actual cases generate statistics, but no beliefs are assumed in the absence of data.

V -maximization is more plausible than is U -maximization if very similar problems are encountered over and over again. But when memory consists only of remotely related problems, V -maximization may not be as convincing. In particular, it involves discontinuity at zero total similarity. For instance, if there is but one case in memory, (q, a, r) , with problem similarity $s(p, q) = \varepsilon$, then $V(a) = u(r)$ for every $\varepsilon > 0$, but $V(a) = 0$ for $\varepsilon = 0$. One may view both U -maximization and V -maximization as rough approximations, whose appropriateness depends on the range of the similarity function.

In Section 18 we assume that the aspiration level changes over time, mostly as a result of past experiences. We show that, under certain assumptions, U -maximization, but not V -maximization, leads to asymptotically optimal choice in a stochastically repeated decision problem. In view of

this result, we tend to view U -maximization as the primary formula, despite the fact that it seems inappropriate for repeated problems with a constant aspiration level.

5.2 Act similarity

While it stands to reason that past performance of an act would affect the act's evaluation in current problems, it is not necessarily the case that past performance is the only relevant factor in the evaluation process. Specifically, an act's desirability may be affected by the performance of other, *similar* acts. For instance, suppose that Ann is looking for an apartment to rent. One of her options is to rent apartment A. Ann hasn't lived in apartment A before. Thus, she has never chosen the particular act now available to her. But Ann has lived in the past in apartment B, which is located in the same neighborhood as is A. The act "renting apartment A" is similar to the act "renting apartment B" at least in that the two apartments are in the same neighborhood. It seems unavoidable that Ann's experience with the similar act will affect her evaluation of the act she is now considering.

Similarly, suppose that Bob tries to decide whether or not to buy a new product in the supermarket. He has never purchased this product in the past, but he has consumed similar products by the same producer. Thus the particular act that Bob is now evaluating has no history. But Bob's memory contains cases in which he chose other acts that are similar to the one he now considers. Again, we would expect Bob's decision to depend on his experience with similar acts.

A decision maker is often faced with new acts, that is, acts that her memory contains no information about their past performance. According to formula (*), the evaluation index attached to these acts is the default value zero. As in the examples above, this application of CBDT is not very plausible. Correspondingly, it may lead to counter-intuitive behavioral predictions. For instance, it would suggest that Ann will be as likely to

rent the apartment in the neighborhood she knows as an apartment in a neighborhood she does not know. Similarly, Bob will be predicted to buy the new product based on his aspiration level alone, without distinguishing among products by the record of their producers. In reality, however, decision makers are reminded of cases involving similar acts, and expect similar acts, chosen in similar problems, to result in similar outcomes. Thus we need to extend the model to reflect act similarity and not just problem similarity.

Act similarity effects are especially pronounced in economic problems involving a continuous parameter. For instance, the decision whether or not to “Offer to sell at price p ” for a specific value p , will likely be affected by the results of previous sell offers at different but close values of p . Generally, if there are infinitely many acts available to the decision maker, it is always the case that most of them are new to her. However, she will typically infer something about these new acts from the performance of other acts she has actually tried. While a straightforward application of CBDT to economic models with an infinite set of acts may result in counter-intuitive and unrealistic predictions, the introduction of act similarity may improve these predictions.

Observe that act similarity effects are not restricted to the evaluation of new acts. Even if an act was chosen in the past in a similar problem, its evaluation is likely to be colored by the performance of similar acts.

The need for modeling act similarity may sometimes be obviated by re-defining “acts” and “problems”. For instance, Ann’s acts may be simply “To Buy” and “Not to Buy”, where each possible purchase is modeled as a separate decision problem. However, such a model is hardly very intuitive, especially when many acts are considered simultaneously. It is more natural to explicitly model a similarity function between acts. Moreover, in many cases the similarity function is most naturally defined on problem-act pairs. For example, “Driving on the left in New York” may be more similar to “Driving on the right in London” than to “Driving on the left in London”,

“Buying when the price is low” may be more similar to “Selling when the price is high” than to “Selling when the price is low”, and so forth.

In short, we would like to have a model in which the similarity function s is defined on problem-act pairs, and, given a memory M and a decision problem p , each act a is evaluated according to the weighted sum

$$U'(a) = U'_{p,M}(a) = \sum_{(q,b,r) \in M} s((p,a), (q,b))u(r) . \quad (\bullet)$$

Observe that a case (q, b, r) in memory may be viewed as a pair $((q, b), r)$, where the problem-act pair (q, b) is a single entity, describing the circumstances from which the outcome r resulted. That is, when past cases are considered, the distinction between problems and acts is immaterial. Indeed, it may also be fuzzy in the decision maker’s memory. By contrast, when evaluating currently available acts, this distinction is both clearer and more important: the problem refers to the given circumstances, which are not under the decision maker’s control, whereas the various acts describe alternative choices.

5.3 Case similarity

It is sometimes convenient to think of a further generalization, in which the similarity function is defined over entire cases, rather than over problem-act pairs alone. According to this view, the decision maker may realize that a similar act in a similar problem may lead to a *correspondingly similar* (rather than identical) result. For instance, assume that John has used automatic vending machines for soft drinks in the past, and that he now faces a vending machine for sandwiches for the first time. There is evident similarity between pushing a button on this machine and pushing buttons on the soft drink machines he has used. Based on this similarity, John expects that pushing a button will result in receiving the product shown in the picture that is adjacent to the button. That is, he expects a sandwich to drop out

of the sandwich machine in spite of the fact that he has never seen a sandwich dropping out of any machine before. Indeed, he has never seen such a machine either. Given the similarity of the problem and the act to previous problem-act pairs, John expects a correspondingly similar outcome to result. He probably considers such an outcome to be more likely than any of the outcomes he has experienced.

While one may attempt to fit this type of reasoning into the framework of U' -maximization by a re-definition of the results, it is probably most natural to assume a similarity function that is defined over whole cases. Thus, the case (*soft drink machine, push button A, get drink A*) is similar to the case (*sandwich machine, push button B, get sandwich B*) more than to (*sandwich machine, push button B, get drink A*). If we assume that the decision maker can imagine the utility of every outcome (even if it has not been actually experienced in the past), we are naturally led to the following generalization of CBDT:

$$U''(a) = U''_{p,M}(a) = \sum_{r \in R} \sum_{(q,b,t) \in M} s((p,a,r), (q,b,t))u(r) . \quad (\bullet\bullet)$$

In this formula every outcome r is considered as a possible result of act a in problem p , and the weight of the utility of outcome r is the sum of similarity values of the present case, should act a be chosen and outcome r result in it, to past cases.

Observe that case similarity can also capture asymmetric inferences.¹² For instance, consider a seller of a product, who posts a price and observes an outcome out of $\{sale, no-sale\}$. Consider two acts, *offer to sell at \$10* and *offer to sell at \$12*. Should the seller fail to sell at \$10, she can safely assume that asking \$12 would also result in *no-sale*. By contrast, selling at \$10 provides but indirect evidence regarding the result of *offer to sell at \$12*. Denoting two generic decision problems (say, days) by p and q , the similarity between $(p, offer\ to\ sell\ at\ \$10, no-sale)$ and $(q, offer\ to\ sell\ at\ \$12, no-sale)$

¹²This point is due to Ilan Eshel.

is higher than that between $(q, \text{offer to sell at } \$12, \text{ no-sale})$ and $(p, \text{offer to sell at } \$10, \text{ no-sale})$.

6 CBDT as a Behaviorist Theory

6.1 W -Maximization

Our focus so far has been on CBDT as a theory that attempts to describe an actual mental process. It makes explicit reference to cognitive terms such as “utility” and “similarity”, and it has a claim to cognitive plausibility. But the reference to cases, representing factual knowledge on the one hand, and to choices on the other also allows a behaviorist version of CBDT. Such a theory would treat cases that actually happened as stimuli, or input, and decisions as responses, or output, without specifying the mental process by which cases affect decisions.

The behaviorist case-based decision theory we suggest here retains the additivity of the theories presented above, but abstracts away from the structure of cases and the related similarity and utility functions. Specifically, we assume that cases are abstract entities. For each case c and each act a we assume that there is a number $w_p(a, c)$ that summarizes the degree to which case c supports the choice of a in the given problem p . Further, we assume that these support indices are accumulated additively, that is, that given a set of cases M , the decision maker will choose an act that maximizes

$$W(a) = W_{p,M}(a) = \sum_{c \in M} w_p(a, c) . \quad (\circ)$$

over $a \in A$.

In the behaviorist interpretation, M is the set of cases to which the decision maker is exposed when faced with the decision problem p . For instance, it may represent the collection of stories that were published in the media available to her. Memory, the set of cases, is the input, or the stimulus. The act chosen is the output, or the response. The theory of W -maximization

makes no reference to any cognitive constructs such as probability, similarity, or utility. It does not try to capture wants or needs, beliefs or evaluations. It merely lumps together the impact of stimulus c on potential response a by a number $w_p(a, c)$, and argues that the impact of several stimuli is additive.

In contrast to the theories described in previous sections, this theory allows cases to be completely abstract, in the sense that they do not have to comprise of a decision problem, an act, and an outcome. Cases here need not have any formal relationship to the acts among which the decision maker has to choose. Hence the interpretation of a “case” may go beyond an explicit story of a previous decision. For instance, suppose that Jane has to decide whether to go on a trip to the country. A relevant case might be the fact that it rained the day before. This case specifies neither an act that was chosen nor an outcome that was experienced. It is still possible, and indeed likely that it will influence Jane’s choice of an act in the current decision problem. Similarly, a stock market crash in 1929 may be a case that affects an investment decision of a person who was born after this case occurred.¹³

It will be convenient (and essential for the axiomatic derivation) to allow cases in memory to appear more than once.¹⁴ It is then natural to define a *memory* to be a function $I : M \rightarrow \mathbb{Z}_+$ (where \mathbb{Z}_+ stands for the non-negative integers) that counts occurrences of cases in M . In this formulation, CBDT offers that the decision maker will choose an act that maximizes

$$W(a) = W_{p,I}(a) = \sum_{c \in M} I(c)w_p(a, c) . \quad (\diamond)$$

over $a \in A$.

¹³As argued above, all the abstract cases discussed here, to the extent that they might affect decisions, can be translated into the language of problems, acts, and outcomes as hypothetical cases. Such a translation, however, constitutes a cognitive task. The present formulation is consistent with a purely behaviorist approach.

¹⁴Alternatively, one may assume that each case is unique, but that two cases may be equivalent in the eyes of the decision maker, as far as problem p is concerned. One can then require that to each case there be infinitely many other cases that are equivalent to it in this sense.

6.2 Cognitive Specification: EUT

In the next chapter we axiomatize the behaviorist theory of W -maximization. While we believe that the axioms suggested therein may convince the reader that our theory is plausible, both descriptively and normatively, we wish to support it also by a cognitive account. Descriptively, our belief in the theory as a valid description of actual decision making will be enhanced if we can describe a mental process that implements it. From a normative viewpoint, such a process is essential to make the theory a practical recommendation for decision making. In Section 2 we refer to such a mental process as a “cognitive specification” of the theory.

A behaviorist theory may have more than one cognitive specification. While we view W -maximization as the behaviorist manifestation of mental processes involving similarity and utility judgments, it is also true that expected utility theory is a cognitive specification of W -maximization. To see this, assume that the decision maker conceives of a set of states of the world Ω . Assume further that she forms beliefs μ over Ω as follows. Each case c induces a probability measure μ_c over Ω , such that, given memory I , the decision maker forms beliefs ¹⁵

$$\mu = \sum_{c \in M} I(c) \mu_c .$$

Should this decision maker maximize the expected value of a utility function $u : A \times \Omega \rightarrow \mathbb{R}$ with respect to μ , she may be viewed as maximizing W in (\diamond) with

$$w_p(a, c) = \int u(a, \omega) d\mu_c(\omega) .$$

¹⁵Observe that, in order to be a probability measure, μ may have to be normalized (separately for each memory I). The process of generation of subjective beliefs in this way is axiomatized in Gilboa-Schmeidler (2000b).

6.3 Cognitive Specification: CBDT

We have seen that expected utility theory with case-based probabilities is a possible cognitive specification of W -maximization. Naturally, the case-based decision rules of U -, U' -, and U'' -maximization discussed above are also cognitive specifications of the same theory. Specifically, assume, as in the previous sections, that all cases c are triples of the form (q, b, r) , namely, that each case describes an act chosen in a decision problem and an outcome that resulted from it. The theory of U -maximization corresponds to the following definition of the support weights $w_p(a, c)$:

$$w_p(a, c) = w_p(a, (q, b, r)) = \begin{cases} s(q, p)u(r) & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}.$$

That is, U -maximization tells a story about a mental process in which each act is evaluated based only on those cases in which it was chosen, and for such a case $c = (q, a, r)$, $w_p(a, c) = w_p(a, (q, a, r))$ is the product of the similarity of the problems by the utility of the outcome. Similarly, U' -maximization would result from defining

$$w_p(a, c) = w_p(a, (q, b, r)) = s((p, a), (q, b))u(r).$$

In other words, the cognitive account of $w_p(a, c)$ suggested by U' -maximization involves the separation of similarity and utility, where similarity is defined for problem-act pairs. Finally, U'' -maximization may be viewed as W -maximization where one defines

$$w_p(a, c) = w_p(a, (q, b, r)) = \sum_{t \in R} s((p, a, t), (q, b, r))u(r).$$

Note that V -maximization is not a special case of W -maximization. To fit the W function, one may define

$$w_p^M(a, c) = w_p^M(a, (q, b, r)) = \frac{s(p, q)}{\sum_{(q', a, r) \in M} s(p, q')} u(r)$$

if $a = b$ and if the denominator is positive, and $w_p^M(a, c) = 0$ otherwise. But in this case $w_p^M(a, c)$ depends on the entire memory M , and not only on the case c in it. Formally, in V -maximization memory changes the similarity function. This is done in a proportional manner: the similarity function is normalized to add up to 1. Generally, memory may affect similarity judgments in more subtle ways, as will be discussed in Section 19. Whenever the similarity function depends on memory, the resulting decision rule should not be expected to be a special case of W -maximization.

Viewing both EUT and various variants of CBDT as special cases of W -maximization facilitates the comparison between them. While both theories can be described as accumulating evidence from past cases, in EUT the decision maker is capable of hypothetical thinking, as she uses a case that occurred in the past both for the evaluation of the act she has indeed chosen in this case (if there was one), and for the evaluation of acts she has not chosen. Moreover, the same probabilities of states of the world are used for all acts. Thus an expected utility maximizer trusts her hypothetical thinking as much as she trusts her actual experience. By contrast, a case-based decision maker who maximizes U or U' evaluates acts based only on the outcomes that have actually been experienced in the past. She does not attempt to imagine what outcomes might have resulted from other acts in the same situations.

6.4 Comparing the cognitive specifications

Viewed thus, both theories appear extreme. U' -maximization (and, perforce, U -maximization) does not allow the decision maker to engage in hypothetical thinking, unless hypothetical cases are explicitly introduced into memory. By contrast, EUT insists that hypothetical thinking is possible and, moreover, that it is just as influential as actual experience. However, the behaviorist theory of W -maximization allows many intermediate degrees of hypothetical thinking. Indeed, U'' -maximization is a cognitive specification of W -maximization that spans this range. In one extreme, we may define

$s((p, a, t), (q, b, r))$ in $(\bullet\bullet)$ as

$$s((p, a, t), (q, b, r)) = \begin{cases} s((p, a), (q, b)) & \text{if } t = r \\ 0 & \text{otherwise} \end{cases},$$

yielding U' -maximization as a special case, and precluding hypothetical thinking about the outcome that could have resulted in case (q, b, r) had a been chosen rather than b . On the other extreme, we may set

$$s((p, a, t), (q, b, r)) = \mu_c(\{\omega \in \Omega \mid a(\omega) = r\})$$

yielding EUT (with respect to beliefs $\mu = \sum_{c \in M} I(c)\mu_c$) as a special case. But $s((p, a, t), (q, b, r))$ may, in general, be non-zero for $t \neq r$ (as in EUT), allowing the decision maker to imagine that a different act in a different problem might result in a different outcome, while still depending on t , allowing a distinction between actual and hypothetical outcomes. Thus, the behaviorist theory of W -maximization offers a wide range of decision making patterns, as described by its cognitive specification of U'' -maximization.

The theory of W -maximization has non-behaviorist interpretations as well. First, the weights $w_p(a, c)$, measuring the support that a case c provides to act a , may be directly accessible by introspection, even if they do not take one of the forms suggested by the cognitive specifications above. Second, one may apply W -maximization to a decision maker's memory that may not be observable by the modeler directly, provided that this memory is available to the decision maker by introspection. Thus W -maximization may be a behavioral or even a cognitive theory.

In the next chapter we axiomatize CBDT in the general form of W -maximization.¹⁶ But most of the discussion that follows relates to the simplest model, namely, U -maximization. While this model is rather restrictive, it highlights the way in which CBDT differs from EUT. Thus we focus on the model that is most likely to be wrong but also most likely to be insightful.¹⁷

¹⁶As mentioned above, W -maximization does not generalize V -maximization. For an axiomatization of V -maximization, see Gilboa and Schmeidler (1995).

¹⁷An exception is the discussion of planning in Chapter 5, which requires the use of

7 Case-Based Prediction

In a prediction problem a predictor is faced with certain circumstances, and is asked to name one of possible eventualities that might arise. Prediction may be regarded as a special type of decision making under uncertainty: the acts available to the predictor are the possible predictions, and the possible outcomes are success (for a correct prediction) and failure (for a wrong one). In a more general model, one may also rank predictions on a continuous scale, measuring the proximity of the prediction to the eventuality that actually transpires, allow set-valued predictions, probabilistic predictions, and so forth.

As any other type of decision under uncertainty, prediction is based on past knowledge. Indeed, the formal structure of a prediction problem would typically include also a history, or memory, namely, a set of examples, consisting of circumstances that were encountered in the past, coupled with the eventualities that are known to have resulted from them. This structure also encompasses problems referred to in the literature as classification, categorization, or learning.

Viewed as a special case of decision problems, prediction problems are characterized by the following feature: outcomes are assumed to be independent of acts. The predictor, viewed as a decision maker, believes that she is an outside observer who may guess the eventuality that would result from circumstances, but that she cannot affect it in any way. She affects the outcome of her own decision problem by guessing the eventuality correctly or incorrectly, but not through the eventuality itself. It follows that, when considering a past case of prediction, the decision maker knows what the eventuality that actually transpired would have been the same had she predicted differently. Thus she can estimate the success of each possible prediction she could have made with the same ease of imagination and the

U'' -maximization.

same certainty as for the prediction she actually did make. In fact, in formally describing a past case of prediction, one may completely suppress the predictor's choice and focus only on the circumstances, namely, the decision problem, and the eventuality.

Let us consider a simple formal model in which memory M contains *examples* of the form $(q, r) \in P \times R$, i.e., pairs of problems, or circumstances, q and eventualities r . Assume that a similarity function $s : P \times P \rightarrow [0, 1]$ is given as above. Given a problem $p \in P$, it is natural to suggest that possible eventualities be ranked according to

$$W'(r) = W'_{p,M}(r) = \sum_{(q,r) \in M} s(p, q) .$$

That is, an eventuality r is assigned a numerical value corresponding to the sum, over all examples in which r actually transpired, of the similarity values of the problems in the past to the one at hand. It is easily seen that W' is a special case of W defined above. Indeed, consider a decision problem in which $A = R$. That is, the set of possible predictions coincides with the set of eventualities. Let cases be pairs $(q, t) \in P \times R$ and define the support weights $w_p(r, c)$ by:

$$w_p(r, c) = w_p(r, (q, t)) = \begin{cases} s(p, q) & \text{if } r = t \\ 0 & \text{otherwise} \end{cases} .$$

With these definitions, W -maximization in the decision problem coincides with W' -maximization in the prediction problem. Moreover, W -maximization also allows more general prediction forms. That is, knowing that eventuality t transpired in a past case may make another eventuality $r \neq t$ more or less plausible.

The following chapter provides an axiomatic derivation of W -maximization in the context of decision under uncertainty. It can also be interpreted as an axiomatization of a general prediction rule based on abstract cases. With the special structure above in mind, one may easily formulate additional conditions that would characterize W' -maximization. (See Gilboa and Schmeidler

(2000b,c), in which we follow this path.) Observe that W' -maximization generalizes kernel estimates of density functions. (See Akaike (1954), Rosenblatt (1956), and Parzen (1962).)

Chapter 3

Axiomatic Derivation

8 Highlights

We devote this chapter to the axiomatization of W -maximization. In this section we describe the model and our key assumptions. Formal discussion is presented in the following section.

The decision maker is asked to rank acts in a set A by a binary preference relation, based on her memory M . Further, we assume that the decision maker has a well-defined preference relation on A not only for the actual memory M , but also for other, hypothetical memories that can be generated from M by replication of cases in it. While the representation (\diamond) in Section 6.1 allows repetitions in memory, we now *require* that preferences be defined for any such vector of repetitions.

In many situations it is not hard to imagine cases appearing in memory repeatedly. For instance, in the recommendation letters example one may imagine several letters relating practically identical stories about a candidate. Cases that are available to a physician in treating a specific patient will often be identical to the best of the physician's knowledge. Moreover, cases that are known to be different may still be considered equivalent as far as the present decision is concerned. Patterns in the history of the weather or of the stock market can also be imagined to have occurred in history more or

less times than they actually have.

Yet, the assumption that cases can be repeated is restrictive. Consider the example of war in Bosnia again. We quoted the Vietnam war as a relevant case. But it is not entirely clear what is meant by assuming that the Vietnam war occurred, say, five times. In what order do these repeated wars supposed to occur? Were the relevant decision makers aware of the previous occurrences? And, on a more philosophical level, can anything ever repeat itself in exactly the same way? Will the repetition itself not make it different?

Indeed, no two cases are identical. But many cases can be identical to the best of the decision maker's knowledge, or to the best of her judgment, at least as far as the present decision problem is concerned. As explained below, one might derive the notion of identity of cases from preference relations as well. In this approach, the philosophical problem is resolved, because identity is not assumed primitive. The decision maker is not supposed to believe that cases are identical. Rather, if the decision maker's preferences reveal that two cases are equivalent in her eyes for the problem at hand, we may treat them as if they were identical. Still, one would need to assume that for each case there are infinitely many other cases that are equivalent to it in this sense.

Another approach to the axiomatic derivation of CBDT involves a different type of hypothetical rankings: rather than assume that entire cases repeated themselves in exactly the same way, one may assume that each case occurred only once, but that it has resulted in a different outcome than the one actually experienced. We follow this tack in the axiomatization of U - and of V - maximization with a given utility function in Gilboa and Schmeidler (1995), in the axiomatization of U' -maximization in Gilboa and Schmeidler (1997), and in that of U' -maximization with an endogenously derived utility function in Gilboa, Schmeidler, and Wakker (1999). There are situations where it is easier to imagine an entire case repeating itself with the same outcome, that to imagine the same case resulting in a different outcome, and

there are situations where the converse holds. Observe, however, that both types of hypothetical rankings are needed only for the axiomatization, not for the formulation of the theory itself. They will be required for elicitation of parameters, but not for application of the theory by a decision maker who can estimate these parameters directly.

The repetitions model makes yet another implicit assumption: not only do repetitions make sense, they are also supposed to contain all relevant information for the generation of preferences. Specifically, the order in which cases have supposedly occurred is not reflected in memory. This assumption is not too restrictive, since the description of a case may include a time parameter or certain relevant features of history. Yet, it has a flavor of de Finetti's exchangeability condition.

Our main assumption within this set-up is a combination axiom. Roughly, it states that, if act a is preferred to act b given two disjoint databases of cases, than a should be preferred to b also given their union. To see the basic logic of this condition, imagine that the science of medicine equips physicians with an algorithm to provide recommendation to patients given their personal experience. Imagine that a patient consults two physicians who work in different hospitals, and who therefore have been exposed to disjoint databases of patients. Imagine further that both physicians recommend treatment a over treatment b . Should the patient insist that they get together to consult her case? A negative answer elicits an implicit belief in the combination axiom: applying the same decision making algorithm to the union of the two databases should result in the same conclusion as applying it to each database separately (if these yielded an identical result).

Yet, the combination axiom may seem rather implausible in various cases. We first present the formal model and then discuss this and the other axioms in more detail.

9 Model and Result

The formal model and the results in this section appear in Gilboa and Schmeidler (2000b). In that paper we discuss predictions and inductive inference, where one assumes that rankings of eventualities by their plausibility are given as data. Here we interpret the rankings as preferences over possible acts.

Let M be a finite a non-empty set of *cases* and A – a set of *acts* available at the decision problem p under discussion. We assume that A contains at least two acts. For simplicity of notation we suppress p throughout. Denote $\mathbb{J} = \mathbb{Z}_+^M = \{I|I : M \rightarrow \mathbb{Z}_+\}$ where \mathbb{Z}_+ stands for the non-negative integers. \mathbb{J} is the set of hypothetical memories, or simply *memories*. We assume that for every $I \in \mathbb{J}$ the decision maker has a binary relation “at least as desirable as” on A denoted by \succsim_I (i.e., $\succsim_I \subset A \times A$).¹In other words, we are formalizing here a rule, as opposed to an act, of decision making. The rule generates decisions not only for the available set of cases, but also for hypothetical ones. Some of the hypothetical collections of cases, i.e., vectors in \mathbb{J} , may become actual when the decision maker acquires additional information. (On this see also Subsection 9.3) In conclusion, our structural assumption (axiom 0) is the existence of a binary relation \succsim_I on A for all $I \in \mathbb{J}$.

9.1 Axioms

We will use the four axioms stated below. In their formalization let \succ_I and \approx_I denote the asymmetric and symmetric parts of \succsim_I , as usual. \succsim_I is *complete* if $x \succsim_I y$ or $y \succsim_I x$ for all $x, y \in A$. Finally, algebraic operations

¹This mathematical structure is akin to, and partly inspired by Young (1975) and Myerson (1995), who derive scoring rules in voting theory. Rather than a binary relation \succeq_I on A , they assume a function selecting a subset of A for every $I \in \mathbb{J}$, to be interpreted as the set of winning candidates given a distribution of ballots I . Our model assumes more information, since we require a complete ranking of alternatives in A for each I . In return, we obtain an almost unique representation, and we do not need to assume symmetry among the alternatives.

on \mathbb{J} are performed pointwise.

A1 Ranking: For every $I \in \mathbb{J}$, \succsim_I is complete and transitive on A .

A2 Combination: For every $I, J \in \mathbb{J}$ and every $a, b \in A$, if $a \succsim_I b$ ($a \succ_I b$) and $a \succsim_J b$, then $a \succsim_{I+J} b$ ($a \succ_{I+J} b$).

A3 Archimedean Axiom: For every $I, J \in \mathbb{J}$ and every $a, b \in A$, if $a \succ_I b$, then there exists $l \in N$ such that $a \succ_{lI+J} b$.

Observe that in the presence of Axiom 2, Axiom 3 also implies that for every $I, J \in \mathbb{J}$ and every $a, b \in A$, if $a \succ_I b$, then there exists $l \in N$ such that for all $k \geq l$, $a \succ_{kI+J} b$.

Axiom 1 simply requires that, given any conceivable memory, the decision maker's preference relation over acts be a weak order. Axiom 2 states that, if act a is preferred to act b given two disjoint memories, then a should also be preferred to b given the combination of these memories. In our set-up, combination (or concatenation) of memories takes the form of adding the number of repetitions of each case in the two memories. Axiom 3 is a continuity condition. It states that if, given the memory I , the decision maker strictly prefers act a to act b , then, no matter what is her ranking for another memory, J , there is a number of repetitions of I that is large enough to overwhelm the ranking induced by J .

Finally, we need a diversity axiom that is not necessary for the functional form we would like to derive. While the theorem we present is an equivalence theorem, it characterizes a more restricted class of preferences than those discussed in the introduction. Specifically, we require that for any four acts, there is a memory that will distinguish among all four of them.

A4 Diversity: For every list (a, b, d, e) of distinct elements of A there exists $I \in \mathbb{J}$ such that $a \succ_I b \succ_I d \succ_I e$. If $|A| < 4$, then for any strict ordering of the elements of A there exists $I \in \mathbb{J}$ such that \succ_I is that ordering.

We remind the reader that the next section is devoted to a detailed discussion of the axioms, and proceed to state their implication.

9.2 Basic result

The key result of this chapter can now be formulated.

Theorem 3.1 : *Let there be given A , M , and $\{\succsim_I\}_{I \in \mathbb{J}}$ as above. Then the following two statements are equivalent if $|A| \geq 4$:*

(i) $\{\succsim_I\}_{I \in \mathbb{J}}$ satisfy A1-A4;

(ii) There is a matrix $w : A \times M \rightarrow \mathbb{R}$ such that:

$$(**) \quad \begin{cases} \text{for every } I \in \mathbb{J} \text{ and every } a, b \in A, \\ a \succsim_I b \quad \text{iff} \quad \sum_{c \in M} I(c)w(a, c) \geq \sum_{c \in M} I(c)w(b, c) , \end{cases}$$

and for every list (a, b, d, e) of distinct elements of A , the convex hull of differences of the row-vectors $(w(a, \cdot) - w(b, \cdot))$, $(w(b, \cdot) - w(d, \cdot))$, and $(w(d, \cdot) - w(e, \cdot))$ does not intersect \mathbb{R}_-^M .

Furthermore, in this case the matrix w is unique in the following sense: w and \widehat{w} both satisfy $(**)$ iff there are a scalar $\lambda > 0$ and a matrix $u : A \times M \rightarrow \mathbb{R}$ with identical rows (i.e., with constant columns) such that $\widehat{w} = \lambda w + u$.

Finally, in the case $|A| < 4$, the numerical representation result (as in $(**)$) holds, and uniqueness as above is guaranteed.

The main point of the theorem is that the infinite family of rankings has a linear representation via a real matrix of order $|A| \times |M|$. For $(a, c) \in A \times M$ the number $w(a, c)$ is interpreted as the amount of support that case c lends to act a .

Given any real matrix of order $|A| \times |M|$, one can define for every $I \in \mathbb{J}$ a weak order on A through $(**)$. It is easy to see that these orders would satisfy A2 and A3 but not necessarily A4. For example, A4 will be violated if a row of the matrix dominates another row. It is therefore natural to wonder whether one can obtain a numerical representation using only A1-A3. The answer is given by the following proposition.

Proposition 3.2 *Axioms A1, A2, and A3 do not imply the existence of a matrix w satisfying (**).*

In other words, A1, A2, and A3 are necessary but (jointly) insufficient conditions for the existence of a matrix w satisfying (**).²

9.3 Learning new cases

The decision rule is constructed in anticipation of incorporating new information. Let us assume that the decision maker conceives of cases in a set M and has information represented by $I \in \mathbb{Z}_+^M$. Correspondingly, she ranks acts by \succsim_I . Assume now that the decision maker acquires a new database. If the new database consists of additional repetitions of cases that the decision maker has conceived of, namely, cases in M , it can be represented by a vector $J \in \mathbb{Z}_+^M$, and the decision maker's new preference will be \succsim_{I+J} . But if the new database contains at least one case c that is not in M , it cannot be represented by a member of \mathbb{Z}_+^M . One may choose another set M' and represent the new database by $J \in \mathbb{Z}_+^{M'}$, but then I and J do not belong to the same space and they cannot be summed up. In particular, the combination axiom does not apply and the representation theorem cannot be invoked. True, one may apply the theorem separately to \mathbb{Z}_+^M and to $\mathbb{Z}_+^{M'}$, but then the theorem would yield a matrix w' for memory M' that may be completely unrelated to the matrix w of the memory M .

Specifically, even if case c belongs to both M and M' , the degree of support it lends to an act a in memory M need not be the same as in memory M' . This would imply that the support that a case lends to an act may depend not only on their intrinsic features, but also on other cases in memory. Should one want to rule out this dependence, one should require an additional axiom. Before we introduce it, a piece of notation will be useful:

²A theorem in Ashkenazi and Lehrer (2000) suggests a condition on $\{\succeq_I\}_{I \in \mathbb{J}}$ that is necessary and sufficient for a representation as in (**). This condition is less intuitive than the axioms we use and is therefore omitted.

for two memories M and M' satisfying $M \subset M'$, and for each $I \in \mathbb{Z}_+^M$, let $I' \in \mathbb{Z}_+^{M'}$ be the extension of I to M' defined by $I(c) = 0$ for $c \in M' \setminus M$. We can now state the condition guaranteeing that the similarity function is independent of memory:

A5 Case Independence: If $M \subset M'$, then $\succsim_I = \succsim_{I'}$ for all $I \in \mathbb{Z}_+^M$.

We now state the sought for result. Its proof is quite simple and hence omitted.

Proposition 3.3 : *Let there be given a set of cases \mathbb{C} , a set of alternatives A , and a family \mathcal{M} of finite subsets of \mathbb{C} . Assume that for every $M \in \mathcal{M}$, $\{\succsim_I\}_{I \in \mathbb{Z}_+^M}$, satisfy A1-A4, that $|A| \geq 4$, and that \mathcal{M} is closed under union. Then the following two statements are equivalent:*

- (i) A5 holds on \mathcal{M} ;
- (ii) For every $c \in \mathbb{C}$ and every $a \in A$ there exists a number $w(a, c)$ such that **(**)** of Theorem 3.1 (ii) holds for every $M \in \mathcal{M}$.

9.4 Equivalent cases

Axiom 5 and Proposition 3.3 enable us to resolve an additional modeling difficulty within our framework. From the description of our model it seems that the initial information of the decision maker is M or, equivalently, $1_M \in \mathbb{Z}_+^M$.³ Assume, for instance, that the observed cases are 5 flips of a coin, say $HTHHT$. Are these five separate cases, with $M = \{c_i | i = 1, \dots, 5\}$ and $I = (1, 1, 1, 1, 1)$? Or is the correct model one where $M' = \{H, T\}$ and $I = (3, 2) \in \mathbb{Z}_+^{M'}$? Or should we perhaps opt for a third alternative where the decision maker's knowledge is $J = (3, 2, 0, 0, 0) \in \mathbb{Z}_+^M$? It would be nice to know that these modeling choices do not affect the result.

³ 1_E stands for the indicator vector of a set E in the appropriate space. In our model we identify M with $1_M \in \mathbb{Z}_+^M$.

First we note that, indeed, under Axiom 5 and the proposition, the rankings of A for $(3, 2)$ and for $(3, 2, 0, 0, 0)$ should be identical. Second, if the decision maker believes that $c_1, c_3, c_4 \in M$ are the “same” case, and so are $c_2, c_5 \in M$, we should also expect the rankings for $(1, 1, 1, 1, 1)$ and for $(3, 2, 0, 0, 0)$ to be identical. Moreover, if this is the case also whenever we “replace” occurrences of c_1 by occurrences of, say, c_3 , we may take this as a definition of “sameness” of cases.

Formally, for two cases $c_1, c_2 \in M$, define $c_1 \simeq c_2$ if the following holds: for every $I, J \in \mathbb{J}$ such that $I(c) = J(c)$ for all $c \neq c_1, c_2$, and $I(c_1) + I(c_2) = J(c_1) + J(c_2)$, we have $\succsim_I = \succsim_J$. In conclusion we state the obvious:

Proposition 3.4 *(i) \simeq is an equivalence relation on M .*

*(ii) If (**) holds then: $c_1 \simeq c_2 \Leftrightarrow \exists \beta \in \mathbb{R}$ s.t., $\forall a \in A$, $w(a, c_1) = w(a, c_2) + \beta$*

It follows that one may start out with a large set of distinct cases, and allow the decision maker’s subjective preferences to reveal possible equivalence relations between them. It is then possible to have an equivalent representation of the decision maker’s knowledge, based on the smallest set of different cases. Thus for the minimal representing matrix w , not only are all columns different, but no column is obtained from another by addition of a constant.

This derivation of equivalence relation between cases, however, still assumes that cases may be repeated in an identical manner. As mentioned above, one may find this objectionable on philosophical grounds. In this case, one may prefer a model that does not assume that identity of cases is primitive. Such a model may be adapted from that presented above as follows. Let there be given a set of distinct cases C . Assume that for every finite $M \subset C$ there exists a relation \succsim_M on A . Define cases $i, j \in C$ to be *equivalent* if $\succsim_{M \cup \{i\}} = \succsim_{M \cup \{j\}}$ for all finite $M \subset C$ with $M \cap \{i, j\} = \emptyset$. It is easy to verify that this is indeed an equivalence relation. The structural assumption is then that every equivalence class of this relation is infinite.

9.5 U -Maximization

Consider again the model described in Section 4, where a case consists of a problem, an act, and an outcome. With this additional structure one may distinguish U -maximization from U' - and U'' -maximization by the fact that in U -maximization an act is evaluated based only on its own history. This may be formulated by the following axiom.

A6 Specificity: For all $I, J \in \mathbb{J}$, and all $a, b \in A$, if $I(c) = J(c)$ whenever $c = (q, a, r)$ or $c = (q, b, r)$, then $a \succsim_I b$ iff $a \succsim_J b$.

A6 requires that if two memories agree on the number of occurrences of all cases involving acts a and b , then the two memories induce the same preference order between these two acts. Observe that this axiom should hold if the history of each given act a affects only the evaluation of this very act. In other words, if we knew that the impact of act a 's history were *specific* to the evaluation of act a itself, we would expect A6 to hold. The following result shows that this axiom is also sufficient to obtain the kind of representation we seek. That is, under this additional axiom, the evaluation of each act does not depend on past performance of other acts.

Proposition 3.5 *Assume that $\{\succsim_I\}_{I \in \mathbb{J}}$ satisfy A1-A4 and A6. Then there exists a matrix $w : A \times M \rightarrow \mathbb{R}$ such that*

$$(***) \left\{ \begin{array}{l} \text{for every } I \in \mathbb{J} \text{ and every } a, b \in A, \\ \\ a \succsim_I b \quad \text{iff} \\ \\ \sum_{c \in M} I(c)w(a, c) \geq \sum_{c \in M} I(c)w(b, c) \quad \text{iff} \\ \\ \sum_{c=(q,a,r) \in M} I(c)w(a, c) \geq \sum_{c=(q,b,r) \in M} I(c)w(b, c) \end{array} \right.$$

Further, this matrix w is unique up to multiplication by a positive number and it satisfies $w(a, (q, b, r)) = 0$ whenever $a \neq b$.

A6 implies that the matrix w of Theorem 3.1 has the following property: if acts b, d were not chosen in case c , then $w(b, c) = w(d, c)$. Since every

column in the matrix w can be shifted by a constant, one may set $w(b, c)$ to zero whenever act b was not chosen in case c . Given this choice of w , the last two lines of $(***)$ are obviously equivalent. Moreover, this choice makes w unique up to multiplication by a positive number. We omit the formal proof of the proposition, which follows this reasoning. Observe that the resulting matrix w consists mostly of zeroes: in every column there exists at most one entry that differs from zero.

In case a utility function $u : R \rightarrow \mathbb{R}$ is given, one may write $w(a, (q, a, r)) = s(q)u(r)$ (unless $u(r)$ vanishes and $w(a, (q, a, r))$ does not).⁴ Interpreting $s(q)$ as the similarity of the past problem q to the problem at hand p yields U -maximization. Indeed, with a given utility function one may decompose the function w even without A6 and write $w(a, (q, b, r)) = s((p, a), (q, b))u(r)$ to get U' -maximization. (Again, such a decomposition assumes that $u(r) = 0$ implies $w(a, (q, b, r)) = 0$.)

Observe, however, that in the absence of a given utility function there is no unique decomposition of $w(a, (q, b, r))$ (or of $w(a, (q, a, r))$) to a product of a similarity function and a utility function. Indeed, if cases can only repeat themselves in memory in their entirety, one cannot hope to observe the separate impacts of similarity judgment and of desirability evaluation on behavior.⁵ If a function w can be written as the product of similarity s and utility u , it is also the product of $-s$ and $-u$. Thus even the ordinal ranking of past outcomes cannot be gleaned from the function w alone. For instance, if Jane tends to choose act a in problem p after having chosen a in case (q, a, r) , it is possible that she liked outcome r and that she considers problem p to be similar to problem q (i.e., $u(r), s(p, q) > 0$), but it is also possible that she did not like outcome r , and that she considers problems p

⁴There are many situations in which a utility function may be assumed given. In particular, one may have a cognitive measure of utility that is independent of the behavioral theory under discussion. See, for instance, Alt (1936) for such a derivation.

⁵This problem is reminiscent of the axiomatization of state-dependent utility with subjective probability. See Karni, Schmeidler, and Vind (1983) and Karni and Mongin (2000).

and q to be dissimilar (i.e., $u(r), s(p, q) < 0$).

To obtain a decomposition of the function w into a product of utility and similarity based on behavior data alone, one would need to consider not only repetitions of cases, but also hypothetical cases in which different outcomes resulted from the same problem-act pairs. Such an approach is used by Gilboa and Schmeidler (1995, 1997a) and Gilboa, Schmeidler, and Wakker (1999).

10 Discussion of the Axioms

Axiom 1 states that the decision maker has a weak preference over the set of acts given any memory. This axiom is standard in decision theory, and it is as plausible here as in most models. Once one accepts the structural assumptions, namely, that the decision maker has preferences given *any* conceivable memory and that only the number, and not the order of cases matters, A1 is not too objectionable.⁶

As mentioned above, the combination axiom (A2) is the most fundamental of our assumptions. Whereas its basic logic seems sound, it is by no means universally applicable. In particular, the CBDT rule of V -maximization does not satisfy it. To see this, consider the following version of “Simpson’s Paradox” (see, for instance, DeGroot (1975)).⁷ One has to choose between betting on coin a coming up Head or on coin b coming up Head. In database I coin a resulted in Head 110 out of 1,000 tosses, whereas coin b – 10 out of 100 tosses. Thus, V -maximization, which is a rather reasonable decision rule in this context, would prefer a (with success rate of 11%) to b (with 10%). In database J , a resulted in Head 90 out of 100 tosses, whereas b – in 800 out

⁶It should be mentioned that both completeness and transitivity of preferences, and, in particular, transitivity of indifference, have been challenged in decision theory. (See Luce (1956) and Fishburn (1970, 1985).) The criticism of these axiom applies here as in any other context in which it was raised.

⁷Simpson’s paradox was suggested to us as a counterexample to the combination axiom by Bernard Walliser.

of 1,000. Again, V -maximization would prefer a to b . But in the combined database a resulted in Head 200 times out of 1,100, whereas b – 810 out of 1,100, and V -maximization would prefer b to a .

The implicit assumption that makes V -maximization a reasonable decision rule in this situation is that the coin tosses for both coin a and coin b result from independent sampling from fixed distributions. But this assumption also makes the scenario above unlikely. Observing databases such as I and J probably indicates that this assumption is false, and that not all tosses of the coins are similar. In particular, in database I the percentage of Heads is much lower than in database J , for both coins. It is possible, for instance, that the tossing machine in database J has a Head bias. Introducing the tossing machine as part of the description of a case would resolve the inconsistency. Specifically, coin tosses from database I would differ from tosses of the same coins from database J , and it would no longer make sense to apply V -maximization to the union of the two databases.

The combination axiom may also be violated when small databases are considered. For instance, consider a statistician who has to decide whether to accept or to reject a null hypothesis. Assume that the statistician uses a standard hypothesis testing technique. It is normally the case that one can find two samples, each of which is too small to reach statistical significance, and whose union is large enough for this purpose. Thus, the decision given each database would be “accept”, whereas given their union it would be “reject”. It should be noted, however, that this is due to the inherent asymmetry between acceptance and rejection of a null hypothesis, which often stems from the fact that the null hypothesis enjoys the status of “accepted view” or “common wisdom”. This, in turn, probably relies on past cases that precede the samples under discussion. If we were to consider the entire database on which the statistician relies, we would include these past cases as well as the new samples. In this case, the two databases defined by past cases combined with each of the new samples are not disjoint.

To consider another example, assume that an inspector has to decide whether a roulette wheel is fair. Observing a hundred cases in which the wheel came up on red, she may conclude that the roulette is biased and that legal action is called for. Obviously, she should make the same decision if she observes a hundred cases, all of which involving black outcomes. Yet, the combination of the two databases will not warrant any action.⁸ Observe, however, that this example hinges on the fact that the acts are not described in sufficient detail. If one were to choose among “no act”, “sue for a red bias”, and “sue for a black bias”, the violation of the combination axiom would disappear.⁹

Perhaps the most important violation of the combination axiom occurs when the function $w(a, c)$ itself is learned from experience. We refer to this process as “second-order induction”, since it describes how one learns how to learn from experience. Chapter 7 discusses this issue in detail.

The Archimedean axiom states that conceivable memories are comparable, or that no conceivable memory can be infinitely more weighty than another. It is violated, for instance, by decision systems that are based on nearest neighbor classifications. To consider a concrete example, imagine that a physician has to decide whether to perform a surgery on a patient. One possible algorithm for making this decision is to consider the most similar patient who has undergone the surgical procedure, and to perform the surgery on the current patient if and only if the most similar past case of surgery was successful. For this decision rule, a single most-similar case of, say, a successful operation will outweigh any number of less similar cases of unsuccessful operations. We view this example as a criticism of nearest neighbor approaches more than of the Archimedean axiom. At any rate, one may drop the Archimedean axiom and develop a nonstandard analysis version of our theorem.

⁸This example is based on an example of Daniel Lehman.

⁹As in the example of Simpson’s paradox, the scenario described in this example is highly unlikely under the assumption of stochastic independence (or exchangeability).

In Gilboa and Schmeidler (2000b) we use the combination and the Archimedean axioms as applied to prediction, and we identify several classes of examples in which they are unlikely to hold. Since every prediction problem can be embedded in a decision problem, these classes of examples apply here as well. Yet, it appears that these axioms are rather plausible in a wide range of applications.

We now turn to the diversity axiom. Its main justification is necessity: without it the representation of preferences by W -maximization cannot be derived. We readily admit that it is too restrictive for many purposes. For instance, consider three acts, a = “sell 100 shares”, b = “sell 200 shares”, and d = “sell 300 shares”. One may argue that act b will always be ranked between acts a and d . But this contradicts the diversity axiom. Yet, with a large enough set of conceivable cases M , the diversity axiom is not too objectionable. For instance, there might be circumstances under which obtaining the optimal portfolio would require selling precisely 200 shares. Having strong enough evidence from past cases that the present problem belongs to this category, one may indeed rank, say, b over a and a over d .

The above notwithstanding, it would be useful to have another axiom that, in the presence of A1-3, would yield W -maximization. It will be clear from the proof that one does not need the full strength of A4 to obtain the desired result. It suffices that there be enough quadruples of acts for which the conclusion of the axiom holds, rather than that it hold for *all* quadruples. Such a weakening would allow act b in the example above to be always ranked between a and d , provided that there are other acts, say e and f , such that for every pair of distinct acts $\{x, y\} \subset \{a, b, d\}$, every strict ranking of $\{x, y, e, f\}$ is obtained for some I . For simplicity of exposition we formulate A4 in its full strength, rather than tailor it to its exact usage in the proof. At present we are not aware of any elegant alternative to A4.

11 Proofs

Proof of Theorem 3.1:

Theorem 3.1 is reminiscent of the main result in Gilboa and Schmeidler (1997). In that work, cases are assumed to involve numerical payoffs, and algebraic and topological axioms are formulated in the payoff space. Here, by contrast, cases are not assumed to have any structure, and the algebraic and topological structures are given by the number of repetitions. This fact introduces two main difficulties. First, the space of “contexts” for which preferences are defined is not a Euclidean space, but only integer points thereof. This requires some care with the application of separation theorems. Second, repetitions can only be non-negative. This fact introduces several complications, and, in particular, changes the algebraic implication of the diversity condition.

We present the proof for the case $|A| \geq 4$. The proofs for the cases $|A| = 2$ and $|A| = 3$ will be described as by-products along the way.

We start by proving that (i) implies (ii). We first note that the following homogeneity property holds:

Claim 3.1 *For every $I \in \mathbb{Z}_+^M$ and every $k \in \mathbb{N}$, $\succsim_I = \succsim_{kI}$.*

Proof: Follows from consecutive application of the combination axiom. \square

In view of this claim, we extend the definition of \succsim_I to functions I whose values are non-negative rationals (denoted \mathbb{Q}_+). Given $I \in \mathbb{Q}_+^M$, let $k \in \mathbb{N}$ be such that $kI \in \mathbb{Z}_+^M$ and define $\succsim_I = \succsim_{kI}$. \succsim_I is well-defined in view of Claim 3.1. By definition and Claim 3.1 we also have:

Claim 3.2 (*Homogeneity*) *For every $I \in \mathbb{Q}_+^M$ and every $q \in \mathbb{Q}$, $q > 0$: $\succsim_{qI} = \succsim_I$.*

Claim 3.2, A1, and A2 imply:

Claim 3.3 (The **ranking axiom**) For every $I \in \mathbb{Q}_+^M$, \succsim_I is complete and transitive on A , and (the **combination axiom**) for every $I, J \in \mathbb{Q}_+^M$ and every $a, b \in A$ and $p, q \in \mathbb{Q}$, $p, q > 0$: if $a \succsim_I b$ ($a \succ_I b$) and $a \succsim_J b$, then $a \succsim_{pI+qJ} b$ ($a \succ_{pI+qJ} b$).

Two special cases of the combination axiom are of interest: (i) $p = q = 1$, and (ii) $p + q = 1$. Claims 3.2 and 3.3, and the Archimedean axiom, A3, imply the following version of the axiom for the \mathbb{Q}_+^M case:

Claim 3.4 (The **Archimedean axiom**) For every $I, J \in \mathbb{Q}_+^M$ and every $a, b \in A$, if $a \succ_I b$, then there exists $r \in [0, 1) \cap \mathbb{Q}$ such that $a \succ_{rI+(1-r)J} b$.

It is easy to conclude from Claim 3.4 (and 3.3), that for every $I, J \in \mathbb{Q}_+^M$ and every $a, b \in A$, if $a \succ_I b$, then there exists $r \in [0, 1) \cap \mathbb{Q}$ such that $a \succ_{pI+(1-p)J} b$ for every $p \in (r, 1) \cap \mathbb{Q}$.

The following notation will be convenient for stating the first lemma. For every $a, b \in A$ let

$$Y^{ab} \equiv \{I \in \mathbb{Q}_+^M \mid a \succ_I b\} \text{ and}$$

$$W^{ab} \equiv \{I \in \mathbb{Q}_+^M \mid a \succsim_I b\}.$$

Observe that by definition and A1: $Y^{ab} \subset W^{ab}$, $W^{ab} \cap Y^{ba} = \emptyset$, and $W^{ab} \cup Y^{ba} = \mathbb{Q}_+^M$. The first main step in the proof of the theorem is:

Lemma 3.5 For every distinct $a, b \in A$ there is a vector $w^{ab} \in \mathbb{R}^M$ such that,

- (i) $W^{ab} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \geq 0\}$;
- (ii) $Y^{ab} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I > 0\}$;
- (iii) $W^{ba} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \leq 0\}$;
- (iv) $Y^{ba} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I < 0\}$;
- (v) Neither $w^{ab} \leq 0$ nor $w^{ab} \geq 0$;
- (vi) $-w^{ab} = w^{ba}$.

Moreover, the vector w^{ab} satisfying (i)-(iv), is unique up to multiplication by a positive number.

The lemma states that we can associate with every pair of distinct acts $a, b \in A$ a separating hyperplane defined by $w^{ab} \cdot x = 0$ ($x \in \mathbb{R}^M$), such that $a \succsim_I b$ iff I is on a given side of the plane (i.e., iff $w^{ab} \cdot I \geq 0$). Observe that if there are only two acts, Lemma 3.5 completes the proof of sufficiency: for instance, one may set $w^a = w^{ab}$ and $w^b = 0$. It then follows that $a \succsim_I b$ iff $w^{ab} \cdot I \geq 0$, i.e., iff $w^a \cdot I \geq w^b \cdot I$. More generally, we will show in the following lemmata that one can find a vector w^a for every alternative a , such that, for every $a, b \in A$, w^{ab} is a positive multiple of $(w^a - w^b)$.

Before starting the proof we introduce additional notation: let \widehat{W}^{ab} and \widehat{Y}^{ab} denote the convex hulls (in \mathbb{R}^M) of W^{ab} and Y^{ab} , respectively. For a subset B of \mathbb{R}^M let $\text{int}(B)$ denote the set of interior points of B .

Proof of Lemma 3.5: We break the proof into several claims.

Claim 3.6 For every distinct $a, b \in A$, $Y^{ab} \cap \text{int}(\widehat{Y}^{ab}) \neq \emptyset$.

Proof: By the diversity axiom $Y^{ab} \neq \emptyset$ for all $a, b \in A, a \neq b$. Let $I \in Y^{ab} \cap \mathbb{Z}_+^M$ and let $J \in \mathbb{Z}_+^M$ with $J(m) > 1$ for all $m \in M$. By the Archimedean axiom there is an $l \in \mathbb{N}$ such that $K = lI + J \in Y^{ab}$. Let $(x_j)_{j=1}^{2^{|M|}}$ be the $2^{|M|}$ distinct vectors in \mathbb{R}^M with coordinates 1 and -1 . For j , ($j = 1, \dots, 2^{|M|}$), define $y_j = K + x_j$. Obviously, $y_j \in \mathbb{Q}_+^M$ for all j . By Claim 3.4 there is an $r_j \in [0, 1) \cap \mathbb{Q}$ such that $z_j = r_j K + (1 - r_j)y_j \in Y^{ab}$ (for all j). Clearly, the convex hull of $\{z_j \mid j = 1, \dots, 2^{|M|}\}$, which is included in \widehat{Y}^{ab} , contains an open neighborhood of K . \square

Claim 3.7 For every distinct $a, b \in A$, $\widehat{W}^{ba} \cap \text{int}(\widehat{Y}^{ab}) = \emptyset$.

Proof: Suppose, by way of negation, that for some $x \in \text{int}(\widehat{Y}^{ab})$ there are $(y_i)_{i=1}^k$ and $(\lambda_i)_{i=1}^k$, $k \in \mathbb{N}$ such that for all i , $y_i \in W^{ba}$, $\lambda_i \in [0, 1]$, $\sum_{i=1}^k \lambda_i = 1$, and $x = \sum_{i=1}^k \lambda_i y_i$. Since $x \in \text{int}(\widehat{Y}^{ab})$, there is a ball of radius $\varepsilon > 0$ around x included in \widehat{Y}^{ab} . Let $\delta = \varepsilon / (2 \sum_{i=1}^k \|y_i\|)$ and for each i let $q_i \in \mathbb{Q} \cap [0, 1]$ such

that $|q_i - \lambda_i| < \delta$, and $\sum_{i=1}^k q_i = 1$. Hence, $y = \sum_{i=1}^k q_i y_i \in \mathbb{Q}_+^M$ and $\|y - x\| < \varepsilon$, which, in turn, implies $y \in \widehat{Y}^{ab} \cap \mathbb{Q}_+^M$. Since for all $i : y_i \in W^{ba}$, consecutive application of the combination axiom (Claim 3.3) yields $y = \sum_{i=1}^k q_i y_i \in W^{ba}$. On the other hand, y is a convex combination of points in $Y^{ab} \subset \mathbb{Q}_+^M$ and thus it has a representation with rational coefficients (because the rationals are an algebraic field). Applying Claim 3.3 consecutively as above, we conclude that $y \in Y^{ab}$ – a contradiction. \square

The main step in the proof of Lemma 3.5: The last two claims imply that (for all $a, b \in A, a \neq b$) \widehat{W}^{ab} and \widehat{Y}^{ba} satisfy the conditions of a separating hyperplane theorem. So there is a vector $w^{ab} \neq 0$ and a number c so that

$$\begin{aligned} w^{ab} \cdot I &\geq c && \text{for every } I \in \widehat{W}^{ab} \\ w^{ab} \cdot I &\leq c && \text{for every } I \in \widehat{Y}^{ba} . \end{aligned}$$

Moreover,

$$\begin{aligned} w^{ab} \cdot I &> c && \text{for every } I \in \text{int}(\widehat{W}^{ab}) \\ w^{ab} \cdot I &< c && \text{for every } I \in \text{int}(\widehat{Y}^{ba}) . \end{aligned}$$

By homogeneity (Claim 3.2), $c = 0$. **Parts (i)-(iv) of the lemma** are restated as a claim and proved below.

Claim 3.8 *For all $a, b \in A, a \neq b$: $W^{ab} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \geq 0\}$; $Y^{ab} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I > 0\}$; $W^{ba} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \leq 0\}$; and $Y^{ba} = \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I < 0\}$.*

Proof: (a) $W^{ab} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \geq 0\}$ follows from the separation result and the fact that $c = 0$.

(b) $Y^{ab} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I > 0\}$: assume that $a \succ_I b$, and, by way of negation, $w^{ab} \cdot I \leq 0$. Choose a $J \in Y^{ba} \cap \text{int}(\widehat{Y}^{ba})$. Such a J exists by Claim 3.6. Since $c = 0$, J satisfies $w^{ab} \cdot J < 0$. By Claim 3.4 there exists $r \in [0, 1)$ such that $rI + (1 - r)J \in Y^{ab} \subset W^{ab}$. By (i), $w^{ab} \cdot (rI + (1 - r)J) \geq 0$.

But $w^{ab} \cdot I \leq 0$ and $w^{ab} \cdot J < 0$, a contradiction. Therefore, $Y^{ab} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I > 0\}$.

(c) $Y^{ba} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I < 0\}$: assume that $b \succ_I a$ and, by way of negation, $w^{ab} \cdot I \geq 0$. By Claim 3.6 there is a $J \in Y^{ab}$ with $J \in \text{int}(\widehat{Y}^{ab}) \subset \text{int}(\widehat{W}^{ab})$. The inclusion $J \in \text{int}(\widehat{W}^{ab})$ implies $w^{ab} \cdot J > 0$. Using the Archimedean axiom, there is an $r \in [0, 1)$ such that $rI + (1 - r)J \in Y^{ba}$. The separation theorem implies that $w^{ab} \cdot (rI + (1 - r)J) \leq 0$, which is impossible if $w^{ab} \cdot I \geq 0$ and $w^{ab} \cdot J > 0$. This contradiction proves that $Y^{ba} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I < 0\}$.

(d) $W^{ba} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \leq 0\}$: assume that $b \succ_I a$, and, by way of negation, $w^{ab} \cdot I > 0$. Let J satisfy $b \succ_J a$. By (c), $w^{ab} \cdot J < 0$. Define $r = (w^{ab} \cdot I) / (-w^{ab} \cdot J) > 0$. By homogeneity (Claim 3.2), $b \succ_{rJ} a$. By Claim 3.3, $I + rJ \in Y^{ba}$. Hence, by (c), $w^{ab} \cdot (I + rJ) < 0$. However, direct computation yields $w^{ab} \cdot (I + rJ) = w^{ab} \cdot I + rw^{ab} \cdot J = 0$, a contradiction. It follows that $W^{ba} \subset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \leq 0\}$.

(e) $W^{ab} \supset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \geq 0\}$: follows from completeness and (c).

(f) $Y^{ab} \supset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I > 0\}$: follows from completeness and (d).

(g) $Y^{ba} \supset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I < 0\}$: follows from completeness and (a).

(h) $W^{ba} \supset \{I \in \mathbb{Q}_+^M \mid w^{ab} \cdot I \leq 0\}$: follows from completeness and (b). \square

Completion of the proof of the Lemma.

Part (v) of the Lemma, i.e., $w^{ab} \notin \mathbb{R}_+^M \cup \mathbb{R}_-^M$ for $a \neq b$, follows from the facts that $Y^{ab} \neq \emptyset$ and $Y^{ba} \neq \emptyset$. Before proving part (vi), we prove **uniqueness**.

Assume that both w^{ab} and u^{ab} satisfy (i)-(iv). In this case, $u^{ab} \cdot x \leq 0$ implies $w^{ab} \cdot x \leq 0$ for all $x \in \mathbb{R}_+^M$. (Otherwise, there exists $I \in \mathbb{Q}_+^M$ with $u^{ab} \cdot I \leq 0$ but $w^{ab} \cdot I > 0$, contradicting the fact that both w^{ab} and u^{ab} satisfy (i)-(iv).) Similarly, $u^{ab} \cdot x \geq 0$ implies $w^{ab} \cdot x \geq 0$. Applying the same argument for w^{ab} and u^{ab} , we conclude that $\{x \in \mathbb{R}_+^M \mid w^{ab} \cdot x = 0\} = \{x \in \mathbb{R}_+^M \mid u^{ab} \cdot x = 0\}$. Moreover, since $\text{int}(\widehat{Y}^{ab}) \neq \emptyset$ and $\text{int}(\widehat{Y}^{ba}) \neq \emptyset$, it follows that $\{x \in \mathbb{R}_+^M \mid w^{ab} \cdot x = 0\} \cap \text{int}(\mathbb{R}_+^M) \neq \emptyset$. This implies that

$\{x \in \mathbb{R}^M \mid w^{ab} \cdot x = 0\} = \{x \in \mathbb{R}^M \mid u^{ab} \cdot x = 0\}$, i.e., that w^{ab} and u^{ab} have the same null set and are therefore a multiple of each other. That is, there exists α such that $u^{ab} = \alpha w^{ab}$. Since both satisfy (i)-(iv), $\alpha > 0$.

Finally, we prove **part (vi)**. Observe that both w^{ab} and $-w^{ba}$ satisfy (i)-(iv) (stated for the ordered pair (a, b)). By the uniqueness result, $-w^{ab} = \alpha w^{ba}$ for some positive number α . At this stage we redefine the vectors $\{w^{ab}\}_{a,b \in A}$ from the separation result as follows: for every unordered pair $\{a, b\} \subset A$ one of the two ordered pairs, say (b, a) , is arbitrarily chosen and then w^{ab} is rescaled such that $w^{ab} = -w^{ba}$. (If A is of an uncountable power, the axiom of choice has to be used.) \square

Lemma 3.9 *For every three distinct eventualities, $f, g, h \in A$, and the corresponding vectors w^{fg}, w^{gh}, w^{fh} from Lemma 3.5, there are unique $\alpha, \beta > 0$ such that:*

$$\alpha w^{fg} + \beta w^{gh} = w^{fh} .$$

The key argument in the proof of Lemma 3.9 is that, if w^{fh} is not a linear combination of w^{fg} and w^{gh} , one may find a vector I for which \succ_I is cyclical.

If there are only three alternatives $f, g, h \in A$, Lemma 3.9 allows us to complete the proof as follows: choose an arbitrary vector w^{fh} that separates between f and h . Then choose the multiples of w^{fg} and of w^{gh} defined by the lemma. Proceed to define $w^f = w^{fh}$, $w^g = \beta w^{gh}$, and $w^h = 0$. By construction, $(w^f - w^h)$ is (equal and therefore) proportional to w^{fh} , hence $f \succsim_I h$ iff $w^f \cdot I \geq w^h \cdot I$. Also, $(w^g - w^h)$ is proportional to w^{gh} and it follows that $g \succsim_I h$ iff $w^g \cdot I \geq w^h \cdot I$. The point is, however, that, by Lemma 3.9, we obtain the same result for the last pair: $(w^f - w^g) = (w^{fh} - \beta w^{gh}) = \alpha w^{fg}$ and it follows that $f \succsim_I g$ iff $w^f \cdot I \geq w^g \cdot I$.

Proof of Lemma 3.9:

First note that for every three distinct eventualities, $f, g, h \in A$, if w^{fg} and w^{gh} are colinear, then for all I either $f \succ_I g \Leftrightarrow g \succ_I h$ or $f \succ_I g \Leftrightarrow$

$h \succ_I g$. Both implications contradict diversity. Therefore any two vectors in $\{w^{fg}, w^{gh}, w^{fh}\}$ are linearly independent. This immediately implies the uniqueness claim of the lemma. Next we introduce

Claim 3.10 *For every distinct $f, g, h \in A$, and every $\lambda, \mu \in \mathbb{R}$, if $\lambda w^{fg} + \mu w^{gh} \leq 0$, then $\lambda = \mu = 0$.*

Proof: Observe that Lemma 3.5(v) implies that if one of the numbers λ , and μ is zero, so is the other. Next, suppose, per absurdum, that $\lambda\mu \neq 0$, and consider $\lambda w^{fg} \leq \mu w^{gh}$. If, say, $\lambda, \mu > 0$, then $w^{fg} \cdot I \geq 0$ necessitates $w^{gh} \cdot I \geq 0$. Hence there is no I for which $f \succ_I g \succ_I h$, in contradiction to the diversity axiom. Similarly, $\lambda > 0 > \mu$ precludes $f \succ_I h \succ_I g$; $\mu > 0 > \lambda$ precludes $g \succ_I f \succ_I h$; and $\lambda, \mu < 0$ implies that for no $I \in \mathbb{Q}_+^M$ is it the case that $h \succ_I g \succ_I f$. Hence the diversity axioms holds only if $\lambda = \mu = 0$. \square

We now turn to the main part of the proof. Suppose that w^{fg}, w^{gh} , and w^{hf} are column vectors and consider the $|M| \times 3$ matrix (w^{fg}, w^{gh}, w^{hf}) as a 2-person 0-sum game. If its value is positive, then there is an $x \in \Delta(M)$ such that $w^{fg} \cdot x > 0$, $w^{gh} \cdot x > 0$, and $w^{hf} \cdot x > 0$. Hence there is an $I \in \mathbb{Q}_+^M \cap \Delta(M)$ that satisfies the same inequalities. This, in turn, implies that $f \succ_I g$, $g \succ_I h$, and $h \succ_I f$ - a contradiction.

Therefore the value of the game is zero or negative. In this case there are $\lambda, \mu, \zeta \geq 0$, such that $\lambda w^{fg} + \mu w^{gh} + \zeta w^{hf} \leq 0$ and $\lambda + \mu + \zeta = 1$. The claim above implies that if one of the numbers λ, μ and ζ is zero, so are the other two. Thus $\lambda, \mu, \zeta > 0$. We therefore conclude that there are $\alpha = \lambda/\zeta > 0$ and $\beta = \mu/\zeta > 0$ such that

$$(*) \quad \alpha w^{fg} + \beta w^{gh} \leq w^{fh}$$

Applying the same reasoning to the triple h, g , and f , we conclude that there are $\gamma, \delta > 0$ such that

$$(**) \quad \gamma w^{hg} + \delta w^{gf} \leq w^{hf}.$$

Summation yields

$$(***) \quad (\alpha - \delta)w^{fg} + (\beta - \gamma)w^{gh} \leq 0.$$

Claim 3.10 applied to inequality (***) implies $\alpha = \delta$ and $\beta = \gamma$. Hence inequality (**) may be rewritten as $\alpha w^{fg} + \beta w^{gh} \leq w^{fh}$, which together with (*) yields the desired representation. \square

Lemma 3.9 shows that, if there are more than three alternatives, the preference ranking of every triple of alternatives can be represented as in the theorem. The question that remains is whether these separate representations (for different triples) can be “patched” together in a consistent way.

Lemma 3.11 *There are vectors $\{w^{ab}\}_{a,b \in A, a \neq b}$, as in Lemma 3.5, and for any three distinct acts, $a, b, d \in A$, the Jacobi identity $w^{ab} + w^{bd} = w^{ad}$ holds.*

Proof: The proof is by induction, which is transfinite if A is uncountably infinite. The main idea of the proof is the following. Assume that one has rescaled the vectors w^{ab} for all alternatives a, b in some subset of acts $A' \subset A$, and one now wishes to add another act to this subset, $d \notin A'$. Choose $a \in A'$ and consider the vectors w^{ad}, w^{bd} for $a, b \in A'$. By Lemma 3.9, there are unique positive coefficients α, β such that $w^{ab} = \alpha w^{ad} + \beta w^{db}$. One would like to show that the coefficient α does not depend on the choice of $b \in A'$. Indeed, if it did, one would find that there are $a, b, c \in A'$ such that the vectors w^{ad}, w^{bd}, w^{cd} are linearly dependent, and this contradicts the diversity axiom.

Claim 3.12 *Let $A' \subset A$, $|A'| \geq 3$, $d \in A \setminus A'$. Suppose that there are vectors $\{w^{ab}\}_{a,b \in A', a \neq b}$, as in Lemma 3.5, and for any three distinct acts, $a, b, e \in A'$, $w^{ab} + w^{be} = w^{ae}$ holds. Then there are vectors $\{w^{ab}\}_{a,b \in A' \cup \{d\}, a \neq b}$, as in Lemma 3.5, and for any three distinct acts, $a, b, e \in A' \cup \{d\}$, $w^{ab} + w^{be} = w^{ae}$ holds.*

Proof: Choose distinct $a, b, c \in A'$. Let w^{ad}, w^{bd} , and w^{cd} be the vectors provided by Lemma 3.5 when applied to the pairs (a, d) , (b, d) , and (c, d) , respectively. Consider the triple $\{a, b, d\}$. By Lemma 3.9 there are unique coefficients $\lambda(\{a, d\}, b), \lambda(\{b, d\}, a) > 0$ such that

$$(I) \quad w^{ab} = \lambda(\{a, d\}, b)w^{ad} + \lambda(\{b, d\}, a)w^{db}$$

Applying the same reasoning to the triple $\{a, c, d\}$, we find that there are unique coefficients $\lambda(\{a, d\}, c), \lambda(\{c, d\}, a) > 0$ such that

$$w^{ac} = \lambda(\{a, d\}, c)w^{ad} + \lambda(\{c, d\}, a)w^{dc}.$$

or

$$(II) \quad w^{ca} = \lambda(\{a, d\}, c)w^{da} + \lambda(\{c, d\}, a)w^{cd}.$$

We wish to show that $\lambda(\{a, d\}, b) = \lambda(\{a, d\}, c)$. To see this, we consider also the triple $\{b, c, d\}$ and conclude that there are unique coefficients $\lambda(\{b, d\}, c), \lambda(\{c, d\}, b) > 0$ such that

$$(III) \quad w^{bc} = \lambda(\{b, d\}, c)w^{bd} + \lambda(\{c, d\}, b)w^{dc}.$$

Since $a, b, c \in A'$, we have

$$w^{ab} + w^{bc} + w^{ca} = 0$$

and it follows that the summation of the right-hand sides of (I), (II), and (III) also vanishes:

$$\begin{aligned} & [\lambda(\{a, d\}, b) - \lambda(\{a, d\}, c)]w^{ad} + [\lambda(\{b, d\}, c) - \lambda(\{b, d\}, a)]w^{bd} + \\ & [\lambda(\{c, d\}, a) - \lambda(\{c, d\}, b)]w^{cd} = 0. \end{aligned}$$

If some of the coefficients above are not zero, the vectors $\{w^{ad}, w^{bd}, w^{cd}\}$ are linearly independent, and this contradicts the diversity axiom. For instance, if w^{ad} is a non-negative linear combination of w^{bd} and w^{cd} , for no I will it be the case that $b \succ_I c \succ_I d \succ_I a$.

We therefore obtain $\lambda(\{a, d\}, b) = \lambda(\{a, d\}, c)$ for every $b, c \in A' \setminus \{a\}$. Hence for every $a \in A'$ there exists a unique $\lambda(\{a, d\}) > 0$ such that, for every distinct $a, b \in A'$ $w^{ab} = \lambda(\{a, d\})w^{ad} + \lambda(\{b, d\})w^{db}$. Defining $w^{ad} = \lambda(\{a, d\})w^{ad}$ completes the proof of the claim. \square

The lemma is proved by an inductive application of the claim. In case A is not countable, the induction is transfinite. \square

Note that Lemma 3.11, unlike Lemma 3.9, guarantees the possibility to rescale *simultaneously* all the w^{ab} -s from Lemma 3.5 such that the Jacobi identity will hold on A .

We now complete the proof that (i) implies (ii). Choose an arbitrary act, say, e in A . Define $w^e = 0$, and for any other alternative, a , define $w^a = w^{ae}$, where the w^{ae} -s are from Lemma 3.11.

Given $I \in \mathbb{Q}_+^M$ and $a, b \in A$ we have:

$$\begin{aligned} a \succ_I b &\Leftrightarrow w^{ab} \cdot I \geq 0 \Leftrightarrow (w^{ae} + w^{eb}) \cdot I \geq 0 \Leftrightarrow \\ &(w^{ae} - w^{be}) \cdot I \geq 0 \Leftrightarrow w^a \cdot I - w^b \cdot I \geq 0 \Leftrightarrow w^a \cdot I \geq w^b \cdot I \end{aligned}$$

The first implication follows from Lemma 3.5(i), the second from the Jacobi identity of Lemma 3.11, the third from Lemma 3.5(vi), and the fourth from the definition of the w^a -s. Hence, (**) of the theorem has been proved.

It remains to be shown that the vectors defined above are such that $\text{conv}(\{w^a - w^b, w^b - w^d, w^d - w^e\}) \cap \mathbb{R}_-^M = \emptyset$ for every distinct a, b, d, e . Indeed, in Lemma 3.5(v) we have shown that $w^a - w^b \notin \mathbb{R}_-^M$. To see this, one only uses the diversity axiom for the pair $\{a, b\}$. Lemma 3.9 has shown, among other things, that a non-zero linear combination of $w^a - w^b$ and $w^b - w^c$ cannot be in \mathbb{R}_-^M , using the diversity axiom for triples. Linear independence of all three vectors was established in Lemma 3.11. However, the full im-

plication of the diversity condition will be clarified by the following lemma. Being a complete characterization, we will also use it in proving the converse implication, namely, that part (ii) of the theorem implies part (i). The proof of the lemma below relies on Lemma 3.5. It therefore holds under the assumption that for any distinct $a, b \in A$ there is an I such that $a \succ_I b$.

Lemma 3.13 *For every list (a, b, d, e) of distinct elements of A , there exists $I \in \mathbb{J}$ such that*

$$a \succ_I b \succ_I d \succ_I e \quad \text{iff} \quad \text{conv}(\{w^{ab}, w^{bd}, w^{de}\}) \cap \mathbb{R}_-^M = \emptyset .$$

Proof: There exists $I \in \mathbb{J}$ such that $a \succ_I b \succ_I d \succ_I e$ iff there exists $I \in \mathbb{J}$ such that $w^{ab} \cdot I, w^{bd} \cdot I, w^{de} \cdot I > 0$. This is true iff there exists a probability vector $p \in \Delta(M)$ such that $w^{ab} \cdot p, w^{bd} \cdot p, w^{de} \cdot p > 0$.

Suppose that w^{ab}, w^{bd} , and w^{de} are column vectors and consider the $|M| \times 3$ matrix (w^{ab}, w^{bd}, w^{de}) as a 2-person 0-sum game. The argument above implies that there exists $I \in \mathbb{J}$ such that $a \succ_I b \succ_I d \succ_I e$ iff the maximin in this game is positive. This is equivalent to the minimax being positive, which means that for every mixed strategy of player 2 there exists $d \in M$ that guarantees player 1 a positive payoff. In other words, there exists $I \in \mathbb{J}$ such that $a \succ_I b \succ_I d \succ_I e$ iff for every convex combination of $\{w^{ab}, w^{bd}, w^{de}\}$ at least one entry is positive, i.e., $\text{conv}(\{w^{ab}, w^{bd}, w^{de}\}) \cap \mathbb{R}_-^M = \emptyset$. \square

This completes the proof that (i) implies (ii). \square

Part 2: (ii) implies (i)

It is straightforward to verify that if $\{\succ_I\}_{I \in \mathbb{Q}_+^M}$ are representable by $\{w^a\}_{a \in A}$ as in (**), they have to satisfy Axioms 1-3. To show that Axiom 4 holds, we quote Lemma 3.13 of the previous part. \square

Part 3: Uniqueness

It is obvious that if $\widehat{w}^a = \alpha w^a + u$ for some scalar $\alpha > 0$, a vector $u \in \mathbb{R}^M$, and all $a \in A$, then part (ii) of the theorem holds with the matrix \widehat{w} replacing w .

Suppose that $\{w^a\}_{a \in A}$ and $\{\widehat{w}^a\}_{a \in A}$ both satisfy (**), and we wish to show that there are a scalar $\alpha > 0$ and a vector $u \in \mathbb{R}^M$ such that for all $a \in A$, $\widehat{w}^a = \alpha w^a + u$.

Choose $a \neq e$ ($a, e \in A$, e satisfies $w^e = 0$). From the uniqueness part of Lemma 3.5 there exists a unique $\alpha > 0$ such that $(\widehat{w}^a - \widehat{w}^e) = \alpha(w^a - w^e) = \alpha w^a$. Define $u = \widehat{w}^e$.

We now wish to show that, for any $b \in A$, $\widehat{w}^b = \alpha w^b + u$. By definition of α and of u , this equality holds for $b = e$ and for $b = a$. Assume, then, that $a \neq b \neq e$. Again, from the uniqueness part of Lemma 3.5 there are unique $\gamma, \delta > 0$ such that

$$\begin{aligned}(\widehat{w}^b - \widehat{w}^a) &= \gamma(w^b - w^a) \\(\widehat{w}^e - \widehat{w}^b) &= \delta(w^e - w^b) .\end{aligned}$$

Summing up these two with $(\widehat{w}^a - \widehat{w}^e) = \alpha(w^a - w^e)$, we get

$$\begin{aligned}0 &= \alpha(w^a - w^e) + \gamma(w^b - w^a) + \delta(w^e - w^b) \\ &= \alpha(w^a - w^e) + \gamma(w^b - w^e) + \gamma(w^e - w^a) + \delta(w^e - w^b) .\end{aligned}$$

Thus

$$(\alpha - \gamma)(w^a - w^e) + (\gamma - \delta)(w^b - w^e) = 0 .$$

and, since $w^e = 0$,

$$(\alpha - \gamma)w^a + (\gamma - \delta)w^b = 0 .$$

Since $w^a = (w^a - w^e) \neq 0$, $w^b = (w^b - w^e) = 0$, and $(w^a - w^e) \neq \lambda(w^b - w^e)$ if $0 \neq \lambda \in \mathbb{R}$, we get $\alpha = \gamma = \delta$. Plugging $\alpha = \gamma$ into $(\widehat{w}^b - \widehat{w}^a) = \gamma(w^b - w^a)$ proves that $\widehat{w}^b = \alpha w^b + u$. \square

This completes the proof of the Theorem 3.1.

Proof of Proposition 3.2 – Insufficiency of A1-3:

We show that without the diversity axiom representability is not guaranteed. Let $A = \{ a, b, d, e \}$ and $|M| = 3$. Define the following vectors in \mathbb{R}^3 :

$$\begin{aligned} w^{ab} &= (-1, 1, 0); & w^{ad} &= (0, -1, 1); & w^{ae} &= (1, 0, -1); \\ w^{bd} &= (2, -3, 1); & w^{de} &= (1, 2, -3); & w^{be} &= (3, -1, -2), \end{aligned}$$

and $w^{xy} = -w^{yx}$ and $w^{xx} = 0$ for $x, y \in A$.

For $x, y \in A$ and $I \in \mathbb{J}$ define: $x \succsim_I y$ iff $w^{xy} \cdot I \geq 0$.

It is easy to see that with this definition the axioms of continuity and combination, and the completeness part of the ranking axiom hold. Only transitivity requires a proof. This can be done by direct verification. It suffices to check the four triples (x, y, z) where $x, y, z \in A$ are distinct and in alphabetical order. For example, since $2w^{ab} + w^{bd} = w^{ad}$, $a \succsim_I b$ and $b \succsim_I d$ imply $a \succsim_I d$.

Suppose by way of negation that there are four vectors in \mathbb{R}^3 , w^a, w^b, w^d, w^e that represent \succsim_I for all $I \in \mathbb{J}$ as in Theorem 3.1. By the uniqueness of representations of half-spaces in \mathbb{R}^3 , for every pair $x, y \in A$ there is a positive real number λ^{xy} such that $\lambda^{xy}w^{xy} = (w^x - w^y)$. Further, $\lambda^{xy} = \lambda^{yx}$.

Since $(w^a - w^b) + (w^b - w^d) + (w^d - w^a) = 0$, we have $\lambda^{ab}(-1, 1, 0) + \lambda^{bd}(2, -3, 1) + \lambda^{da}(0, 1, -1) = 0$. So, $\lambda^{bd} = \lambda^{da}$, and $\lambda^{ab} = 2\lambda^{bd}$. Similarly, $(w^a - w^b) + (w^b - w^e) + (w^e - w^a) = 0$ implies $\lambda^{ab}(-1, 1, 0) + \lambda^{be}(3, -1, -2) + \lambda^{ea}(-1, 0, 1) = 0$, which in turn implies $\lambda^{ab} = \lambda^{be}$ and $\lambda^{ea} = 2\lambda^{be}$. Finally, $(w^a - w^d) + (w^d - w^e) + (w^e - w^a) = 0$ implies $\lambda^{ad}(0, -1, 1) + \lambda^{de}(1, 2, -3) + \lambda^{ea}(-1, 0, 1) = 0$. Hence, $\lambda^{ad} = 2\lambda^{de}$ and $\lambda^{ea} = \lambda^{de}$.

Combining the above equalities we get $\lambda^{ad} = 8\lambda^{da}$, a contradiction.

Observe that in this example the diversity axiom does not hold. For explicitness, consider the order (b, d, e, a) . If for some $I \in \mathbb{J}$, say $I = (k, l, m)$, $b \succ_I d$ and $d \succ_I e$, then $2k - 3l + m > 0$ and $k + 2l - 3m > 0$. Hence, $4k - 6l + 2m + 3k + 6l - 9m = 7k - 7m > 0$. But $e \succ_I a$ means $m - k > 0$, a contradiction. \square

Chapter 4

Conceptual Foundations

This chapter deals with the conceptual and epistemological foundations of case-based decision theory. The discussion is most concrete when CBDT is juxtaposed with the two other formal theories of reasoning. We begin with a comparison of CBDT with EUT. We then proceed to compare CBDT with rule-based systems. We argue that, on epistemological grounds, CBDT is more naturally derived than the two other approaches.

12 CBDT and Expected Utility Theory

We devote this section to several remarks on the comparison between CBDT and EUT, and to our views concerning their possible applications and philosophical foundations. For the purposes of this discussion we focus on the most naive, and simplest version of CBDT, namely, U -maximization.

12.1 Reduction of theories

While CBDT is probably a more natural framework in which one may model satisficing behavior, EUT can be used to explain this behavior as well. For instance, assume that we observe a decision maker who never uses certain acts available to her. By ascribing to her an appropriately chosen prior beliefs, we may describe this decision maker as an expected utility maximizer. In

fact, it is probably possible to provide an EUT account of any application in which CBDT can be used, by using a rich enough state space and an elaborate enough prior on it. Conversely, one may also simulate an expected utility maximizer by a U -maximizer whose memory contains a sufficiently rich set of hypothetical cases.¹ Indeed, let there be given a set of states of the world Ω and a set of consequences R . Denote the set of acts by $A = R^\Omega = \{a : \Omega \rightarrow R\}$. Assume that the agent has a utility function $u : R \rightarrow \mathbb{R}$ and a probability measure μ on an algebra of subsets Ω . (For simplicity we may consider a finite Ω .) The corresponding case-based decision maker would have a hypothetical case for each pair of state of the world ω and act a :

$$M = \{((\omega, a), a, a(\omega)) \mid \omega \in \Omega, a \in A\}$$

Letting the similarity of the problem at hand to the problem (ω, a) be $\mu(\omega)$, U -maximization reduces to expected utility maximization. (Naturally, if Ω or R are infinite one would have to extend CBDT to deal with an infinite memory.) Furthermore, Bayes' update of the probability measure may also be reflected in the similarity function: a problem whose description indicates that an event $B \subset \Omega$ has occurred should be set similar to degree zero to any hypothetical problem (ω, a) where $\omega \notin B$.

Since one can mathematically embed CBDT in EUT and vice versa, it is probably impossible to choose between the two on the basis of predicted observable behavior.² Each is a refutable theory *given a description of a decision problem*, where its axioms set the conditions for refutation. But in most applications there is enough freedom in the definition of states or cases, probability or similarity, for each theory to account for the data. Moreover, a problem that is formulated in terms of states has many potential translations

¹As mentioned above, U'' -maximization can simulate expected utility maximization without introducing hypothetical cases explicitly.

²Matsui (2000) formally proves an equivalence result between EUT and CBDT. His construction does not resort to hypothetical cases.

to the language of cases and vice versa. It is therefore hard to offer real world phenomena that conform to one theory but that cannot be explained at all by the other. It is even harder to imagine a clear-cut test that will select the “correct” theory.

To a large extent, EUT and CBDT are not competing theories; they are different *conceptual frameworks*, in which specific theories are formulated.³ Rather than asking which one of them is more accurate, we should ask which one is more convenient. The two conceptual frameworks are equally powerful in terms of the range of phenomena they can describe. But for each phenomenon, they will not necessarily be equally intuitive. Furthermore, the specific theories we develop in these conceptual frameworks need not provide the same predictions given the same observations. For instance, assume that a theorist observes a collection of phenomena that could be classified as satisficing behavior. She can explain all the data observed whether she uses EUT or CBDT as her conceptual framework. But the language of CBDT will be more conducive to observing this pattern and to forming a theory of satisficing for future predictions. Needless to say, there are large classes of phenomena for which the language of EUT will be the more natural, which will also generate more successful predictions. Hence we believe that there is room for both languages and for the two conceptual frameworks that employ these languages.

12.2 Hypothetical reasoning

Judging the cognitive plausibility of EUT and CBDT, one notes a crucial difference between them: CBDT, as opposed to EUT, does not require the decision maker to think in hypothetical or counterfactual terms. In EUT, whether explicitly or implicitly, the decision maker considers states of the world and reasons in propositions of the form “If the state of the world were ω and I chose a , then r would result”. In U -maximization no such

³See the discussion in Section 2.

hypothetical reasoning is assumed.

Similarly, there is a difference between EUT and CBDT in terms of the informational requirements they entail regarding the utility function: to consciously implement EUT, one needs to know the utility function u , that is, its value for any consequence that may result from any act. For U - (or U' -) maximization, on the other hand, it suffices to know the u -values of those outcomes that were actually experienced.

The reader will recall, however, that our axiomatic derivation of CBDT involved preferences between acts given hypothetical memories. It might appear therefore that CBDT depends on hypothetical reasoning just as EUT does. But this conclusion would be misleading. First, one has to distinguish between elicitation of parameters by an outside observer, and application of the theory by the decision maker herself. While elicitation of parameters such as the decision maker's similarity function may involve hypothetical questions, a decision maker who knows her own tastes and similarity judgments need not engage in any hypothetical reasoning in order to apply CBDT. By contrast, hypothetical questions are intrinsic to the application of EUT.

Second, when states of the world are not naturally given, the elicitation of beliefs for EUT also involves inherently hypothetical questions. Classical EUT maintains that no loss of generality is involved in assuming that the states of the world are known, since one may always define the states of the world to be all the functions from available acts to conceivable outcomes. This view is theoretically very appealing, but it undermines the supposedly behavioral foundations of Savage's model. In such a construction, the set of conceivable acts one obtains is much larger than the set of acts from which the decision maker can actually choose. Specifically, let there be given a set of acts A and a set of outcomes X . The states of the world are X^A , that is, the functions from acts to outcomes. The set of conceivable acts will be $\bar{A} \equiv X^{(X^A)}$, that is, all functions from states of the world to outcomes. Hence the cardinality of the set of conceivable acts \bar{A} is by two orders of magnitude

larger than that of the actual ones A . Yet, using a model such as Savage's, one needs to assume a (complete) preference order on \bar{A} , while *in principle* preferences can be observed only between elements of A . Differently put, such a canonical construction of the states of the world gives rise to preferences that are intrinsically hypothetical and is a far cry from the behavioral foundations of Savage's original model.

In summary, when there is no obvious mapping from the formal state space to real-world scenarios, both EUT and CBDT rely on hypothetical questions or on contingency plans for elicitation of parameters. The Savage questionnaire to elicit EUT parameters will typically involve a much larger set of acts than the corresponding one for CBDT. Importantly, the Savage questionnaire contains many acts that are hard to imagine. Furthermore, when it comes to application of the theory, CBDT clearly requires less hypothetical reasoning than does EUT.

12.3 Observability of data

In order to view choice behavior as observable, EUT needs to assume that the states of the world are known to an outside observer. Correspondingly, CBDT needs to assume that cases are observable in this sense. It might appear that states of the world are part of the objective description of the decision problem, whereas cases in the decision maker's memory are entirely subjective and cannot be directly observed. This would indeed be true if one were to use EUT only in problems where states are naturally defined, and to apply CBDT to *all* cases in a person's memory. But, as argued above, EUT is often applied to situations where states are not given. In these situations it is not obvious what are the states conceived of by the decision maker. Further, CBDT can be applied only to cases that are objectively known, say, cases that were published in the media. With such applications, choice behavior in the EUT model would appear to rely on unobservable data, whereas behavior in the CBDT model would be based on sound, objective data. Thus, there

is nothing inherent to states or to cases that makes one concept more easily observable than the other.

12.4 The primacy of similarity

We argue that any assignment of probabilities to events, and therefore any application of EUT, relies on subjective similarity judgments in one form or another. Consider first the “classical” approach, according to which, for example, equal probabilities are assigned to the sides of a die about to be rolled. Evidently, such an assignment, justified by Laplace’s “principle of insufficient reason”, is based on perceived similarity between any two possible outcomes. The “frequentist” approach uses empirical frequencies to determine probabilities. It thus relies on the concept of repetition, namely, on the idea that the *same* experiment is repeated over and over again. But what counts as a repetition of the same experiment is a matter of subjective judgment. Consider the example of medical data. Suppose that one asks one’s doctor what is the probability of success of a certain medical procedure. One would hardly be happy with empirical frequencies that pertain to all patients who have ever undergone the procedure. Rather, one would like to know what are the empirical frequencies in a sub-population of patients with similar characteristics such as age and gender. But if one obtains too specific a definition of “similar patients”, one might be left with an empty database, since no two human bodies are absolutely identical. It follows that there is a subjective element in the judgment of similar cases over which one would like to compute empirical frequencies. In principle, this observation applies to any application of the frequentist method. When using repetitions of a toss of a coin over and over again, one implicitly assumes that the conditions under which the coin is tossed are identical. But no two instances can be completely identical. Rather, one resorts to subjective judgment to determine what conditions are similar enough to be viewed as identical. Finally, using a subjectivist approach of assigning subjective probabilities by

responding to de Finetti's or Savage's questionnaires one again has to assume that different binary decisions can be made under identical conditions.

Indeed, very little can be said about single instances in life. One never knows what would the history of the world have been had different events taken place, or had different choices been made. Our ability to discuss counterfactuals in an intelligent manner, as well as our ability to assign probabilities to events, relies on our subjective similarity judgments. We therefore conclude that the notion of similarity is primitive. It lies at the heart of probability assignments, as well as at the heart of induction (as we argue in the Chapter 7). While we do not attempt to defend the particular way in which CBDT makes use of similarity between cases, we maintain that the language of CBDT is more basic than the language of EUT (or the language of rule-based systems).

12.5 Bounded rationality?

There is a temptation to view EUT as a normative theory, and CBDT – as a descriptive theory of “bounded rationality”. We maintain, however, that normative theories should be constrained by practicality. Theoretically, one may always construct a state space. Further, Savage's axioms appear very compelling when formulated relative to such a space. But they do not give any guidance for the selection of a prior over the state space. Worse still, in many interesting applications the state space becomes too large for Bayesian learning to correct a “wrong” prior. Thus it is far from clear that, in complex problems for which past data are scarce, case-based decision making is any less rational than expected utility maximization.

Using the definition of rationality proposed in Section 2, there may be many situations in which decision makers who do not follow EUT would not wish to change their behavior even after being confronted with its analysis. Consider, for instance, the example of military intervention in Bosnia discussed in Section 4. Suppose that a decision theorist addresses President

Clinton and suggests that he be rational and that he use expected utility theory rather than analogies to past cases. Clinton may say, “Fine, I see your theoretical point. But how should I form a prior over this vast state space?” In absence of a practical way to generate such a prior, Clinton may not be embarrassed by the fact that he bases his decision on analogies to the Vietnam and the Gulf wars. He may well choose to follow the same decision procedure even if he agrees that Savage’s axioms make sense. More generally, we believe that, as long as one employs a definition of rationality that relates to reasoning, and not only to behavior, one may find cases in which CBDT is be more rational than is EUT.

13 CBDT and Rule-Based Systems

In this section we compare CBDT to rule-based systems from an epistemological viewpoint, and proceed to offer a conceptual derivation of CBDT. That is, we argue for a particular epistemological position, and show that the theory presented above is a natural application of this position for the problem of decision making under uncertainty.

13.1 What can be known?

In philosophy and in artificial intelligence (AI) it is common to assume that general rules of the form “For all $x, P(x)$ ” are objects of knowledge. Some rules of this form may be considered analytic propositions, which are true by definition. Yet, many are synthetic propositions, and represent empirical knowledge.⁴ Knowledge of such a proposition is supposedly obtained by explicit induction, that is, by generalization of particular cases that can be viewed as instances of the proposition.

⁴By “synthetic” propositions we refer to non-tautological ones. While this distinction was already rejected by Quine (1953), we still find it useful for the purposes of the present discussion.

Explicit induction is a natural cognitive process. It is particularly common for people to convey knowledge by formulating rules. Yet, as was already pointed out by Hume, explicit induction does not rely on sound logical foundations. Hume (1748, Section IV) writes,

“... The contrary of every matter of fact is still possible; because it can never imply a contradiction, and is conceived by the mind with the same facility and distinctness, as if ever so conformable to reality. *That the sun will not rise to-morrow* is no less intelligible a proposition, and implies no more contradiction than the affirmation, *that it will rise*. We should in vain, therefore, attempt to demonstrate its falsehood.”

Hume claims that synthetic propositions are not necessarily true. Useful rules, which have implications regarding empirical facts in the future, cannot be *known*.⁵ Explicit induction suggests that such rules are nevertheless objects of knowledge. Because unwarranted generalizations may be refuted, explicit induction raises the problem of knowledge revision and update. Much attention has been devoted to this problem in the recent literature in philosophy and AI. (See Levi (1980), McDermott and Doyle (1980), Reiter (1980) and others.) The Humean alternative to this problem will attempt to avoid induction in the first place, rather than performing it and then dealing with the problems it poses. Following this line of thought, knowledge representation should confine itself to those things that can indeed be known. Facts, or cases can be known, while rules can at best be conjectured. In this approach, knowledge revision is unnecessary. While rules can be contradicted by other rules or by evidence, cases cannot contradict each other.

Admittedly, even cases may not be objectively known. Terms such as “empirical knowledge” and “knowledge of a fact” prove elusive to define.

⁵Quine (1969a) writes, “... I do not see that we are further along today than where Hume left us. The Humean predicament is the human predicament”.

(See, for instance, Moser (1986) for an anthology on this subject.) Furthermore, if we take Hanson's (1958) view that observations may be theory-laden, we might find that the very formulation of the cases that we allegedly observe may depend on the rules that we believe to apply. If this is the case, one cannot know or even represent cases as independent of rules. Moreover, even if one could, such a representation would not solve the epistemological problem, since cases cannot be known with certainty any more than rules can.

We are sympathetic to both claims, but we tend to view them as issues of secondary importance. We believe that the theoretical literature on epistemology and knowledge representation may benefit, and indeed has benefited from drawing distinctions between theory and observations, and between the knowledge of a case and that of a rule. As argued in Section 3, philosophy, being a social science, cannot expect models to be more than “approximations”, “metaphors”, or “images”. Namely, we should not expect philosophical models to provide a complete and accurate description of reality. In this spirit, we choose to employ a conceptual framework in which cases can be known, and can be observed independently of rules or theories.

13.2 Deriving case-based decision theory

Assume, then, that cases are the object of knowledge. How should they be described, and how are they used in decision making? In this subsection we argue for the particular model of CBDT presented in Chapter 2.

First, we maintain that the two questions are intrinsically linked. The choice of knowledge representation has to be guided by the use one makes of knowledge. Focusing on our final goal, namely, to provide a theory of decision making, we take the view that knowledge is ultimately used for action. Thus, we should choose a way to represent cases in anticipation of the way they will be used in decision making. Hence the formal structure of a case should reflect its decision making aspect, if such exists.

Assume, first, that all cases tell stories of past decision. When considering such a case, it is natural to set our clocks to the time at which decision occurred. At that point, the **past** was whatever was known to the decision maker at the time when she was called upon to act. This may be referred to as the *decision problem*. The **present** was the *act* chosen by the decision maker. The **future**, at the time of decision, was what resulted from the choice. This is referred to as the *outcome*. Thus, the description of a case as a triple of decision problem, act, and outcome, is a natural way to organize knowledge of cases for the purposes of decision making.

As mentioned in Section 6, there are cases that do not explicitly involve any decision making, yet are relevant to future decisions. For instance, deciding whether or not to go on a trip to the country, and observing clouds in the sky, it is relevant that yesterday cloudy skies were followed by rain. Using the notion of hypothetical cases, one may translate such a case, which does not involve any choice, to a case that does: one might imagine that, had one gone on a trip yesterday, when the sky was cloudy, the trip would have been a disaster. Alternatively, one may suppress the act in the description of the case, and think of the “problem” simply as the “circumstances” that were followed by an outcome. Formally, one may assume that in any decision problem there is at least one act that is available, say, “do nothing”, and thus fit cases that involve no decisions into our framework.

Circumstances are the conditions under which an outcome has been observed. Presumably, a case is relevant only to cases that have similar circumstances. In other words, the circumstances of a past case help us delineate its scope of relevance. But some cases have no such qualifications. For instance, Hume’s sun rises every morning. A case of the sun rising in the past does not involve a set of circumstances. We may thus suppress the decision problem, or represent it by an empty set of circumstances, and find that the resulting case is relevant to all future cases.

It appears that only the outcome, namely, a description of what hap-

pened, is essential for a case to be informative. The problem and the act may or may not be present. When they are, they may be jointly viewed as antecedents, whereas the outcome is the consequent. The distinction between the problem and the act is along the line drawn by control: those circumstances that were not under the decision maker's control are referred to as the problem, and those that were – as the act.

How are the cases used in decision making, then? Again, we resort to Hume (1748) who writes,

“In reality, all arguments from experience are founded on the similarity which we discover among natural objects, and by which we are induced to expect effects similar to those which we have found to follow from such objects. ... From causes which appear *similar* we expect similar effects. This is the sum of all our experimental conclusions.”

As quoted above, Hume argued against explicit induction, by which one formulates a general rule. Instead, he suggested that people engage in *implicit* induction: they learn from the past regarding the future by employing similarity judgments. The “causes” to which Hume refers are the circumstances we mention above, the combination of the decision problem and the act. The “effect” corresponds to our outcomes. Hence our description of a case may be taken as a formalization of Hume's implicit structure, with the further specification allowing a distinction between problems and acts. When observing phenomena over which one has no control, this distinction is trivial, and our notion of a problem coincides with “circumstances” and with Hume's “cause”. But when we focus on decision making and on the way in which one may affect circumstances and thereby – also outcomes, dividing circumstances into problems and acts seems natural.

Moreover, when considering a current decision problem, the distinction between problems and acts is a quintessential of rational choice: one has to

know what is under one's control and what is not, which variables are decision variables and which are given. As mentioned in Section 5, however, when learning from the experience one has with past problems, the distinction between problem and act might be blurred. What seems to be relevant is the circumstances that they jointly constitute, where the question regarding future problems would be whether one can bring about similar circumstances.

We assume that cases are objectively known. By contrast, the similarity function is a matter of subjective judgment. Where does one derive it? How do and how should people measure similarities? We do not offer any general answers to these questions. But one might hope that the questions will be better defined, and the answers easier to obtain, if we know how similarity is used. That is, we should first ask ourselves, what would we do with a similarity function if we had one?

A well-known procedure that employs similarities is the “nearest neighbor” (NN) approach applied to classification problems. (See Fix and Hodges (1951, 1952), Royall (1966), Cover and Hart (1967), and Devroye, Györfi, and Lugosi (1996)). In these problems a “classifier” is faced with an instance that belongs to one of several possible categories. Based on examples of correct categorizations of past instances, the classifier has to guess the classification of the present one. The simplest nearest neighbor algorithm suggests choosing the category to which belongs the most similar instance among the known examples. k -NN algorithms would consider the k most similar cases, and choose the category that is most common amongst them. This corresponds to a majority vote among the k most similar past instances, and it is also generalized to a weighted majority vote, where more similar known examples get a higher weight in the voting scheme.

Viewed merely as a classification system, we argue that one has to use all cases rather than a few nearest neighbors. When only one nearest neighbor is used, the algorithm appears rather extreme: should one most similar case single-handedly outweigh any number of slightly less similar ones? Indeed,

this is the main theoretical reason to consider k -NN approaches. But for any $k > 1$, the weight assigned to a past instance depends not on its intrinsic similarity to the new one, but also on that of other instances. This can be shown to result in violations of the combination axiom of Section 9: it is possible that a k -NN system would provide an identical classification given two disjoint databases of examples, but a different one given their union. It seems logical that each case would be relevant for the new classification to a degree that is independent of other cases. Thus, one may wish to use all neighbors, near or far, in each problem. They might be weighted by their similarities. Indeed, this would result in the prediction rule discussed in Section 7.⁶

But how does one generate a decision making procedure based on these ideas? As mentioned in Section 7, classification and prediction are decision problems of a very special type, in which the decision maker's choices do not affect eventualities. Thus, by definition, each past case provides a correct answer, namely, the eventuality that has actually occurred, and each possible decision can be evaluated in relation to this correct answer. By contrast, in the general problem of decision making under uncertainty, acts affect outcomes. Considering a past case, one knows what resulted from the act that was actually chosen, but one does not know what would have resulted from other acts that could have been chosen. Thus, there is no "correct" choice, and there is no database of examples to learn correct or optimal choices from. If the decision maker is satisfied with the outcome she has experienced in similar cases in the past, she might take the same decision in the new one as well. But what would a nearest neighbor approach prescribe if the decision maker finds her choice in similar cases truly disappointing? She would probably

⁶Using all known examples may prove to be inefficient from a computational point of view. This difficulty may be overcome by using only examples whose similarity to the present instance is above a prespecified threshold. Such a solution is equivalent to rounding off low similarity values to zero, and it may thus render the problem tractable. The key point, however, is that the decision whether to use a given example depends only on *its own* similarity value, and not on the relative ranking of this value.

choose a different act than the one she has chosen. But which one? And how does the decision maker decide whether the most similar cases involved good or bad choices on her part?

It follows that we need to employ some notion of utility, measuring desirability of outcomes. It is the interaction between utility of outcomes and similarity of circumstances that should determine whether one would like to repeat the choice made in similar cases. Combining this notion with the arguments above for using all cases in each problem, one comes up with the CBDT decision rules suggested in Chapter 2. These functions evaluate every act for each problem. Moreover, the impact that a case has on the decision is independent of other cases that may or may not exist in memory.

13.3 Implicit knowledge of rules

Jane and Jim are two decision makers who face the same problem: they are thirsty, and they see an automatic vending machine for soft drinks. Jane is a case-based decision maker, while Jim is a rule-based one. Jim knows the rule, “if coins are put into a vending machine and a button is pushed, a can comes out”. Jane has no knowledge of such general rules. But she has many past cases in memory, in which she has put some coins in a machine, pushed a button, and got a can. Based on this knowledge, Jane chooses to put some coins in and push a button.

Jane and Jim would probably seem identical in their behavior to an outside observer. When Jane lets the coins drop into the appropriate slot she has an air of self-confidence that would make an outside observer believe that she knows what the outcome of her act is going to be. It appears *as if* Jane knows the same rule as does Jim. Jane might not have such a rule explicitly formulated in her mind, and she might not even be able to formulate it should we ask her about it. But she behaves as if she knew such a rule, and, as such, exhibits implicit knowledge of this rule. This implicit knowledge is the result of the process we refer to as “implicit induction”.

Next assume that Jane walks confidently toward the machine, drops the coins in the slot, pushes a button, and observes that nothing happens. No can, no money. Jane is surprised. She evidently had some expectations regarding the outcome of her act, and these were proven wrong. Jane now has to reconsider her choice. She will probably find that, given the new case, the act of dropping coins has a low U -value. Consequently, she is likely to choose another act (the new U -maximizer), such as “try another machine”.

Suppose now that the same unfortunate occurrence happens to Jim. Jim is not only surprised to find that the machine does not deliver. Jim has a problem with his knowledge base. Evidently, one of the rules he believed in is false. Is it that “if coins are put into a vending machine and a button is pushed, a can comes out”? Or perhaps he made a mistake in identifying the machine, and the rule that is wrong is “large colorful boxes with buttons on them are vending machines”? Jim has to decide whether to retract one or more rules, or to refine it, in order to render the rules and the observations in his knowledge base consistent. On top of this non-trivial task, Jim still has to cope with thirst. Thus, both Jane and Jim have to do something about their predicament, and the set of alternative choices they have is identical. But Jim looks for another machine in a state of mental commotion, trying to put his knowledge base in order. Jane looks for another machine with peace of mind.

There are many databases of cases that can be thought of as corresponding to certain rules. These databases can be described either by cases or by rules, where the latter enjoy the advantage of parsimony. But there are other databases in which rules do not provide a good description of the data. In general, rules are but rough approximations of databases, and in many instances they provide unsatisfactory descriptions thereof. In such instances case-based knowledge representation is found to be less pretentious and more flexible than rule-based representation. Specifically, case-based representation does not need to deal with inconsistencies, and it deals in a continuous

and graceful manner with situations that are on the borderline between the realms of applicability of different rules.

13.4 Two roles of rules

Mary tells John, “Sarah is always late”. Thus, she conveys information in the form of a general rule. Let us distinguish between two scenarios: first, assume that John has never met Sarah. In this scenario, the rule that Mary relates to John is a parsimonious description of many cases that are known to Mary but not to John. In the second scenario, John knows Sarah just as well as does Mary. In this situation, the rule that Mary relates does not tell John about cases that were unknown to him. Rather, it may draw his attention to a certain regularity in the cases in his memory.⁷

Case-based decision makers who do not know rules might still use them as convenient tools for the transfer of information. But rules can have two different roles: they can summarize many cases,⁸ and they can point out similarities between cases. As our example suggests, there is nothing in the rule itself that characterizes its specific role. Rather, whether a rule conveys cases or similarities would depend on the context: the state of knowledge of the communicators (and their knowledge thereof). For instance, most people learn that “objects are drawn to the ground” as an explicit rule when they already have a very clear implicit knowledge thereof. Telling this rule to a child of six years of age does not add much to the child’s set of cases that is not already there. But it does help the child see the similarity between many cases and use it for generalizations.

In both roles, rules may contradict each other. As parsimonious descriptions of cases, one rule may summarize one set of cases, while another rule describes another set. Similarly, as cues for similarity, two rules may point out two different patterns that are simultaneously evident in the same set of

⁷The same statement might also insinuate that Mary is upset, and this may be the only new information that Mary is trying to convey. We ignore this possibility here.

⁸Riesbeck and Schank (1989) refer to rules in this capacity as “ossified cases”.

cases. Such rules may be thought of as proverbs. They do not purport to be literally true. Rather, their main goal is to affect similarity judgments. In this capacity, the fact that rules tend to contradict each other poses no theoretical difficulty. Indeed, it is well known that proverbs are often contradictory.⁹

To sum, a CBDT model may use rules. A rule may be a summary of cases or a feature that is common to many cases and that can thus help define the similarity function. In the CBDT model, however, rules are not taken literally and they may well contradict each other, or be inconsistent with some cases.

⁹The notion of a “rule” as a “proverb” also appears in Riesbeck and Schank (1989). They distinguish among “stories”, “paradigmatic cases”, and “ossified cases”, where the latter “look a lot like rules, because they have been abstracted from cases”. Thus, CBR systems would also have “rules” of sorts, or “proverbs”, which may, indeed, lead to contradictions.

Chapter 5

Planning

14 Representation and Evaluation of Plans

14.1 Dissection, selection, and recombination

CBDT describes a decision as a single act that directly leads to an outcome. In many cases of interest, however, one may take an act not for its immediate outcome, but in order to be in a position to take another act following it. In other words, one may plan ahead. In this section we extend CBDT to a theory of case-based planning.

The formal model of CBDT distinguishes between problems, acts, and results. When planning is considered, the distinction between problems and results is blurred. The outcome of today's acts will determine the decision problem faced tomorrow. Thus the formal model of case-based planning will not distinguish between the two. Rather, we employ a unified concept of a "position", which also can be viewed as a set of circumstances. A position might be a starting point for making a decision, that is, a problem, but also as the end result, namely, an outcome. We will therefore endow a position with (i) a set of available acts (in its capacity as a problem) and (ii) a utility valuation (when considered an outcome). Part of the planning process will be the decision, whether a certain position should be a completion of the plan, or a starting point for additional acts.

The basic premise of our model is that planning is case-based, but that it is carried out for each problem separately. That is, as in the single-stage version of CBDT, the decision maker has only her memory to rely on in planning her actions. But she does not have to restrict her attention to those plans that she has used, or learned of, in their entirety. Rather, she can use various parts of plans she knows of and recombine them in a modular way for the problem at hand.

Consider the following example. Mary owns an old car, and she would like to buy a new one. She has a friend, John, who has recently sold a used car, similar to her own, for \$4,000. By the similarity of the cars and the recency of the sale, Mary imagines that she, too, can get \$4,000 for her old car. Naturally, the new car will be more expensive. While its sticker price is rather high, Mary heard that another acquaintance of hers, Jim, managed to get it for \$9,000. Since Mary wants to get a car of the same make and model, she believes that she, too, can get it for the same amount. It is therefore left to figure out how to have that amount. If Mary succeeds in selling her old car, she will need additional \$5,000. But this should not be a problem. Or at least so says her friend Sarah. Indeed, last week Sarah just walked into the bank and walked out with a \$5,000 loan which she used for some home renovation.

With all this information available to her, Mary devises her plan. She will first sell her old car, which should fetch \$4,000, just like John's car did. Then she'll get a \$5,000, as did Sarah. Finally, she should be able to get her desired car for \$9,000, as did Jim.

Figure 5.1 illustrates Mary's plan as a graph. A node in this graph represents a position, that is, a set of circumstances that are relevant to an agent, before or after making a decision. Arcs denote acts that may be chosen by agents, namely, possible decisions. The arc corresponding to an act leads from the position node at which it was taken to the position node it yielded. The arcs on top are part of Mary's memory, while those at the bottom are

part of her plan.

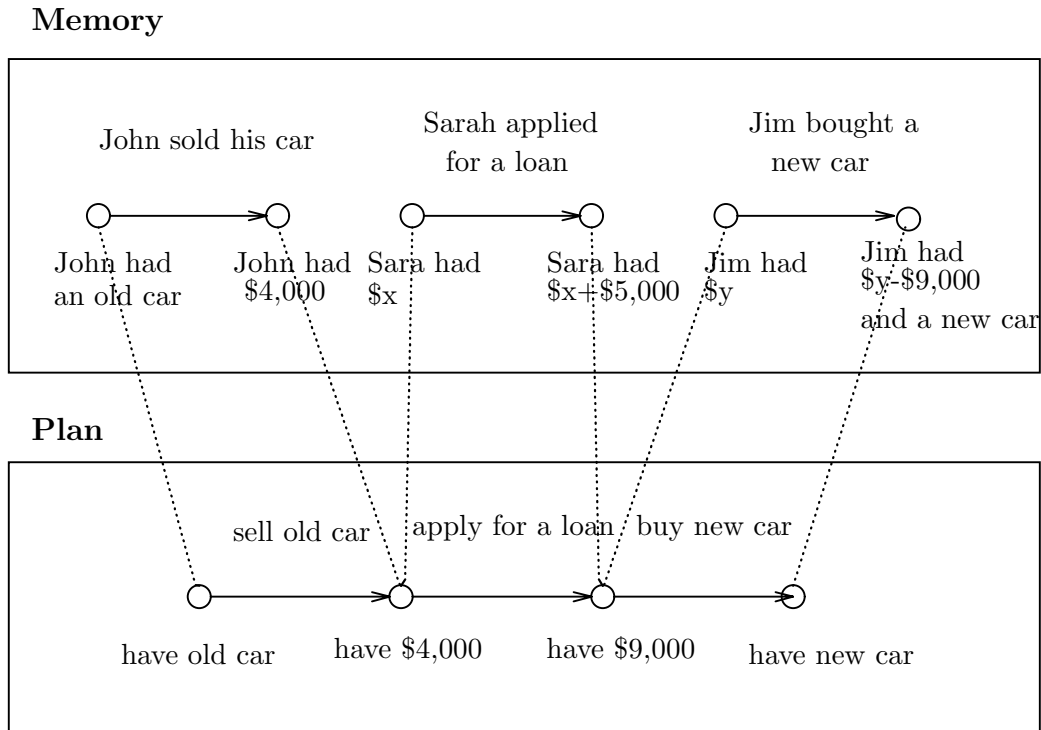


Figure 5.1: Planning as dissection, selection and recombination

Memory arcs need not be isolated. They may be parts of involved and elaborate stories. For instance, John might have sold his car as part of his plan to take a trip around the world. Sarah took the loan to renovate her home, which might have been part of her plan to sell it and move elsewhere, and so forth. The cases that are depicted above as part of memory are therefore assumed to have been selected from entire stories stored in memory.

The point of this example is that, in order for a case-based planner to come up with a plan such as Mary's, she does not need to know of anyone who has done this entire sequence of acts. It suffices that, for each act

in each position, she knows of someone who has done a similar act in a similar position. According to this view, the essence of planning is *dissection, selection, and recombination*. Thinking involves analyzing stories, namely, sequences of cases, and coming up with new combinations of their building blocks. Planning is case-based in the sense that only similarity to known cases can serve as a basis for plan construction and evaluation. But a case-based planner is smarter than a case-based decision maker in that the former can span a wider range of plans than those that were actually tried in the past.

14.2 Representing uncertainty

Next consider the incorporation of uncertainty. While Mary's plan seems very reasonable, it might be wise on her part to consider various eventualities. What happens if she does not succeed to sell her old car for \$4,000? After all, she can *attempt* to sell the car, but she cannot be sure of the outcome. And what will happen if her loan application is denied? In short, uncertainty is present.

The standard approach to describe decision problems of this nature employs decision trees. According to this approach, we are supposed to introduce chance nodes at this point, and allow an act to lead to a chance node, which, in turn, may lead to several position nodes. The distinction between chance nodes and decision nodes, or between states and acts, is fundamental to rational decision making, and may well be the most important cornerstone of decision theory. Confounding acts and states is probably one of the worst sins a theorist can commit, and it leads to such problems as choosing one's beliefs, or to entertaining beliefs over acts. Yet it seems as if we intend to commit this sin.

The reason is that we seek to model decision situations in which decision makers may have very poor knowledge of the potential scenarios that might ensue from each given decision. In such problems, as in the one-stage version

of CBDT, one has to specify how attractive is a choice about which little or no information exists. Hence default values play an important role in our theory. In the plan evaluation rule suggested below, each position will have a utility index, that may be interpreted as the desirability of being in this position indefinitely (and with certainty). Every position will be assigned a weight, and plans will be evaluated by the weighted average of the utility function. Thus, even if the decision maker plans to act at a given position, the utility of this position will play a role in the evaluation of the plan. Modeling this feature in a standard decision tree would make it cyclical: one of the outcomes of an act will be the position at which it was taken. We find it more convenient to suppress chance nodes and retain the acyclic nature of the tree.

We therefore model uncertainty simply by allowing more than one arc to leave a given position for a certain act. The plan tree we get might be thought of as a decision tree in which a chance move was integrated with the decision move preceding it. Thus potential cycles collapse into single nodes, and one retains a simpler mathematical structure.

Going back to our example, what does Mary plan to do if she doesn't manage to sell her car for \$4,000? She might reason as follows: if I get at least \$3,000, I might still proceed with the previous plan, but I will borrow more money from the bank. It stands to reason that, if they were willing to lend \$5,000 to Sarah, they will lend me \$6,000,¹ and the cost of the loan will still not be prohibitive. But if I get something between \$2,000 and \$3,000 for my car, I can't really get the car I want. I will go ahead with the general plan, and borrow up to \$6,000, but I will buy a less expensive car. Finally, if I do not get even \$2,000, I will just drop the whole idea.

We illustrate this plan in Figure 5.2. To simplify the graphical exposition, we assume that the old car can only fetch one of three prices: \$2,000, \$3,000,

¹Mary might also know of another case, in which Carolyn got a loan of \$6,000 to pay her dentist.

and \$4,000. For simplicity we also suppress the role of memory in the graph. However, every arc should be viewed as supported by known cases, as in Figure 5.1.

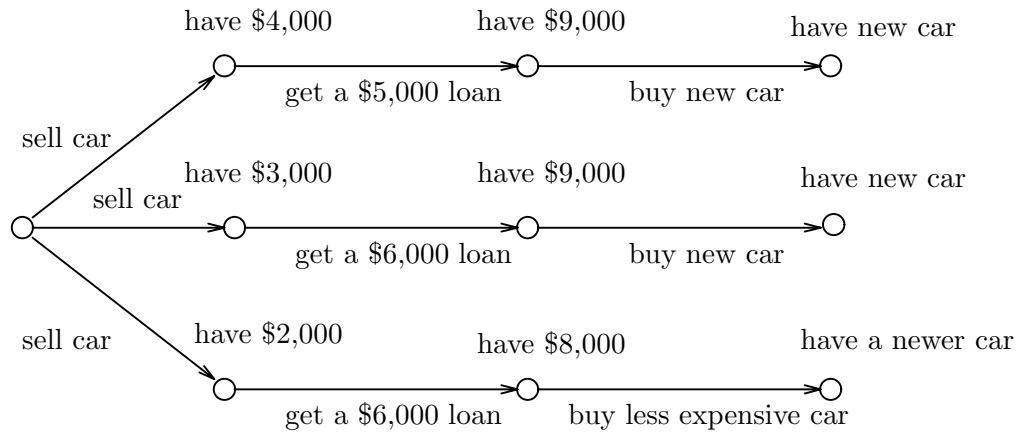


Figure 5.2: Uncertainty represented by multiple arcs

Observe that three arcs correspond to the act “sell car” at the initial position. The plan, however, deals with four contingencies: the three prices mentioned in the graph, and prices that are lower than \$2,000. If only such offers are received, Mary will just forget about the idea of getting a new car. This would mean staying at the initial position. One may therefore assume that at each position and for each act there is also an arc leading from that position to itself. For simplicity we suppress these arcs, and deal with acyclic graphs. But when we come to evaluate a plan we will recall that any position may also be viewed as a terminal node. Observe also that the acts and positions are only sketchily described above. For instance, the act “sell car” should be described as “offer to sell old car for \$4,000”. Finally, one notes that the top two branches of the plan seem to be identical from a certain point onwards. In this example, this is only due to our imprecise description of the positions: in the top branch Mary will be richer by \$1,000,

and this is supposed to be reflected in the description of the position. But in other examples one should expect to find plan paths coinciding, and there is little doubt that in actual planning it is much more efficient to consider general directed acyclic graphs rather than trees. For the purposes of the theoretical study of rules for plan evaluation, however, we restrict attention to trees despite the wastefulness of this representation.

14.3 Plan evaluation

How should Mary evaluate this plan? We assume that she has a payoff assessment for *each* position. This function should reflect, say, the desirability of staying in the initial position with her old car, or of being stuck with no loan and no car, and so forth. Each position will then be assigned a weight, and the plan will be evaluated by the weighted sum of the payoffs. Thus, the weight associated with the initial position should reflect the support that memory lends to the prospect of not being able to sell the old car for \$2,000 or more. These weights naturally depend on the plan. If, for instance, Mary compares the plan above to another plan in which she does not attempt to sell her car at all, the weight of the initial position in the second plan will be higher than in the first.

We now turn to the formal model. Let P be a finite and non-empty set of *positions*. We assume that it is endowed with a binary relation $\succ \subset P \times P$, interpreted as “can follow”. Assume that \succ is a partial order, namely, that it is transitive and irreflexive. Let A denote a finite and non-empty set of *acts*. For $p \in P$, let $A_p \subset A$ be the set of acts *available* at p . We introduce $a_0 \in A$, interpreted as “do nothing”, that is, as the “null act”, and assume that $a_0 \in A_p$ for all $p \in P$. A *case* is a triple (p, a, q) where $p, q \in P$; $a \in A_p$; $q \succ p$. The decision maker is assumed to know of past cases, constituting her memory M , and to imagine future cases that might occur. The set of cases imagined by the decision maker is denoted by C . The decision maker is at an *initial position* $p_0 \in P$. The set of imagined cases C defines a graph

on P , with arcs

$$\hat{C} \equiv \{(p, q) \in P^2 \mid \exists a \in A_p \text{ s.t. } (p, a, q) \in C\} .$$

Assume that this graph is a tree whose root is p_0 . Under this assumption, cases in C can be identified with arcs in this graph. A **plan** is a function $N : P \rightarrow A$ such that $N(p) \in A_p \forall p$. Without loss of generality we assume that a plan does not prescribe any action at a position that is considered terminal. That is, $N(p) = a_0$ whenever p is \succ -maximal. For a plan N , let $P_N \subset P$ be the positions that are consistent with N (not including the initial position):

$$P_N \equiv \left\{ p \in P \mid \begin{array}{l} \exists k \geq 1, \exists p_1, \dots, p_k \in P \text{ s.t.}, \forall 0 < i < k - 1, \\ (p_i, N(p_i), p_{i+1}) \in C, \text{ where } p_k = p \end{array} \right\}$$

A *utility* function is a real-valued function defined on positions. It will prove convenient to normalize it so that the utility of the initial position is zero. Formally, we define a utility function to be $u : P \setminus \{p_0\} \rightarrow \mathbb{R}$.

A *support* function $S : C \rightarrow \mathbb{R}_+$ assigns a weight to each arc, such that, for every $(p', a', p) \in C$ and every $a \in A_p$,

$$\sum_{(p, a, q) \in C} S((p, a, q)) \leq S((p', a', p))$$

Note that only one arc (p', a', p) leads to any position p . The value $S((p, a, q))$ is interpreted as the support memory lends to the future case (p, a, q) if act a is to be taken at position p . (In a Bayesian interpretation, $S((p, a, q))$ would be the probability of the case (p, a, q) conditioning on a plan that prescribes, among other choices, to take a at p . Given the plan, however, S reflects ex-ante probabilities, i.e., $S((p, a, q))$ is *not* conditional on arriving at p .)

A support function S and a plan N induce a weight function $w_S^N : P_N \rightarrow \mathbb{R}_+$ defined by

$$w_S^N(p) = S((p', N(p'), p)) - \sum_{(p, N(p), q) \in C} S((p, N(p), q)) \quad (*)$$

where $(p', N(p'), p)$ is the unique arc leading to p , and $p' \in P_N \cup \{p_0\}$. Observe that, if N, N' are two plans and $p \in P_N \cap P_{N'}$, it follows that

$$w_S^N(p) + \sum_{q \in P_N, q \succ p} w_S^N(q) = w_S^{N'}(p) + \sum_{q \in P_{N'}, q \succ p} w_S^{N'}(q) = S((p', N(p'), p)) .$$

That is, if two plans prescribe the choice of acts that may lead to a given position p , then these two plans induce the same total weight on the sub-tree emanating from p . The plans may, however, differ in the weights they induce for specific positions in this sub-tree.

Conversely, every collection of weight functions $\{w^N : P_N \rightarrow \mathbb{R}_+\}_N$ that satisfies

$$w^N(p) + \sum_{q \in P_N, q \succ p} w^N(q) = w^{N'}(p) + \sum_{q \in P_{N'}, q \succ p} w^{N'}(q) \quad (\bullet)$$

whenever $p \in P_N \cap P_{N'}$, defines a unique support function S satisfying $(*)$ by

$$S((p', N(p'), p)) = w^N(p) + \sum_{q \in P_N, q \succ p} w^N(q) . \quad (**)$$

We assume (and later axiomatize) the following planning rule: given memory M and an initial position $p_0 \in P$, the planner may be ascribed a support function $S \equiv S_{M, p_0}$, such that, given any payoff function u , she ranks plans according to

$$U(N) = U_{M, p_0, u}(N) = \sum_{p \in P_N} w_S^N(p) u(p) .$$

In this formulation, as well as in the axiomatization that follows, M is an abstract index. Thus it can be a set of cases, but also a set of rules, a data base of empirical frequencies, or anything else. If M is indeed a collection of cases, one may further postulate that the support function is derived from similarity to past cases in an additive manner:

$$S((p, a, q)) = \sum_{(p', a', q') \in M} s((p, a, q), (p', a', q'))$$

for some case similarity function s . Moreover, one may derive the additive formula axiomatically in an analogous way to our results from Chapter 3. In this case, the formula above is a straightforward extension of CBDT in a single-stage problem.

Observe that uncertainty is reflected in our model by the fact that a position and an act may evolve into several other positions. But the model also lends itself easily to situations of ignorance, namely, situations where the decision maker has but a vague concept of the possible eventualities. This is captured by the fact that the plan evaluation scheme may assign some weight to the payoff at a position even if one plans to act in it. In this sense, a position is taken to be the default for whatever may happen when an act is taken at it. The default value at the initial position is normalized to zero, and is therefore absent from the formula. But the payoff at any other position appears in the formula with a corresponding weight. The entire formula may be viewed as a linear function of default values.

14.4 Discussion

In what ways do our plans differ from Bayesian decision trees? Superficial differences are the following: (i) Bayesian decision trees distinguish between decision nodes and chance nodes, whereas plan trees allow both the planner and so-called Nature to act at every position; and (ii) In a Bayesian decision tree payoff is only assessed at the terminal nodes, whereas in a plan tree it is assessed at every position.

These differences are mostly notational. One may embed our model in a Bayesian decision tree: one only has to introduce a “chance node” after every act at every position, and allow each position to lead immediately to a corresponding terminal node whose ex-ante probability is that position’s weight. Despite this formal equivalence, we believe that plan trees are closer to the way people conceive of plans, and may therefore be more useful than Bayesian decision trees for descriptive and normative purposes alike. In particular, plan trees are better suited to describe the dynamic process by which a plan evolves. Typically, a decision maker is only aware of certain cases when she plans. As her plan is being executed, she might think further ahead, be reminded of additional cases, or learn new information. The way people think about plans does not admit a clear, static distinction between decision nodes and terminal nodes. Rather, this distinction is dynamic: a set of circumstances that was thought of as terminal yesterday might be the beginning of a new story today.

Is it rational to develop plans in a piecewise manner, rather than thinking them through? Indeed, our model describes people who think ahead several steps, and then treat positions as if they were terminal nodes. But this is often the only way possible. In complex environments, such as a game of chess, one cannot even imagine the entire Bayesian decision tree. Given this constraint, one may not be embarrassed to learn that one has not charted the entire tree. In other words, partial planning may well pass our rationality test of Section 2 on grounds of feasibility.

Planning trees, where weights are assigned to positions, can also be interpreted temporally: consider first a planning tree whose positions are linearly ordered. One may think of $w_S^N(p)$ as the time the planner will be in position p (rather than the probability that the terminal outcome be p). More generally, $w_S^N(p)$ may be the (discounted) expected time of being at p . Our notion of a “position” is also closely related to the notion of a “state” in the literature on Markov chains and dynamic programming, as well as in

the planning literature. Indeed, dynamic programming models also do not distinguish between a “decision node” and an “outcome”. Planning trees, however, are acyclic and do not assume any stationarity.

15 Axiomatic Derivation

15.1 Set-up

We devote this section to an axiomatic derivation of the plan evaluation rule presented in Section 14. We assume that, given *every* payoff function u , there is a preference relation \succsim_u on plans. By referring to $\{\succsim_u\}_u$, we implicitly assume that planners can meaningfully express preferences between pairs of plans given various payoff functions, and not only given their actual function. Thus, the planner is assumed to be able to say, “I plan to sell my house and buy a condo; but if taxes were lower, I would plan not to sell the house”.

This set-up deviates from the revealed preference paradigm. First, we expect planners to express preferences on what they will do if a certain position is reached in the future. Then we also require that they condition these preferences on new, imagined payoffs. We believe that such a deviation from the revealed preferences paradigm is inherent to the formal study of planning. Plans often are not carried out for a variety of reasons. In particular, a planner may find, upon reaching a position, new information that makes her choose a different path than she has previously planned to take. Therefore, one cannot restrict oneself to the planner’s actual choices and has to resort to cognitive and introspective data.

To simplify notation, we will assume that preferences are defined on payoff vectors rather than on plans. That is, given a payoff function u and a plan N , we identify N with the restriction of u to P_N , which is an element of $F_N \equiv \mathbb{R}^{P_N}$. Formally, we impose the following axiom: for every two plans N, \tilde{N} , and every two payoff functions u, u' , if

$$u(p) = u'(p) \text{ for all } p \in P_N \cup P_{\tilde{N}},$$

then

$$N \succsim_u \tilde{N} \Leftrightarrow N \succsim_{u'} \tilde{N}.$$

Denote $F = \bigcup_N F_N$ and define $\succsim' \subset F \times F$ as follows: for $x \in F_N, y \in F_{\tilde{N}}, N \neq \tilde{N}$, $x \succsim' y$ if there exists a (by the axiom above, for all) payoff function u such that

$$u(p) = x(p) \quad \forall p \in P_N, \quad u(p) = y(p) \quad \forall p \in P_{\tilde{N}},$$

and

$$N \succsim_u \tilde{N}.$$

Thus \succsim' ranks any two vectors $x \in F_N, y \in F_{\tilde{N}}, N \neq \tilde{N}$ such that

$$x(p) = y(p) \text{ for all } p \in P_N \cap P_{\tilde{N}}.$$

We will assume without loss of generality that $N \neq \tilde{N}$ implies that $P_N \setminus P_{\tilde{N}} \neq \emptyset$. (In particular, $P_N \neq \emptyset$ for all N .) That is, we assume that the set of imagined cases is rich enough so that each of two distinct plans reaches positions that the other does not. This assumption, however, is only made for notational convenience. Many interesting applications will involve situations in which the decision maker does not have any idea what would be the outcome of a given act in a given position. Rather than dealing with absence of arcs, we model these situations by arcs (imagined cases) with zero support.

We say that $\hat{P} \subset P$ is (\succsim -) *null* if for all $x, x', y, y' \in F$ such that

$$x(p) = x'(p); y(p) = y'(p) \quad \text{for all } p \notin \hat{P},$$

we have $x \succsim y \Leftrightarrow x' \succsim y'$.

15.2 Axioms and result

In the following we start with a relation \succsim' that ranks utility vectors corresponding to different plans. We will define \succsim to be its transitive closure. Thus \succsim may compare two vectors $x, y \in F_N$ that describe the utility derived from the same plan N , if there exists another vector $z \in F_{\tilde{N}}$ for $\tilde{N} \neq N$ that may indirectly compare x and y .

Let $\succsim \subset F \times F$ satisfy the following axioms:

A1 Order: \succsim is reflexive and transitive, and whenever $x \in F_N, y \in F_{\tilde{N}}, N \neq \tilde{N}, x \succsim y \Leftrightarrow x \succsim' y$.

We define $\preceq, \succ, \prec, \approx$ as usual. Implicit in the following axioms is the assumption that a utility function is given, and that the vectors $x, y \in F_N$ are already measured on a utility scale.

A2 Continuity and Comparability: For every two distinct plans N, \tilde{N} , and every $x \in F_N$, the sets $\{y \in F_{\tilde{N}} | y \succ x\}$ and $\{y \in F_{\tilde{N}} | y \prec x\}$ are open in $F_{\tilde{N}}$. If $P_{\tilde{N}} \setminus P_N$ is non-null, they are also non-empty.

A3 Monotonicity: For every two plans N, \tilde{N} , and every $x, z \in F_N, y \in F_{\tilde{N}}$, if $x(p) \geq z(p) \forall p \in P_N$, then $z \succsim y$ implies $x \succsim y$ and $y \succsim x$ implies $y \succsim z$.

A4 Separability: For every two plans N, \tilde{N} , and every $x, z \in F_N, y, w \in F_{\tilde{N}}$, if $x \approx w$ then $x \succsim y$ iff $x + z \succsim y + w$.

A5 Consequentialism: For every two plans N, \tilde{N} , and every $x \in F_N, y \in F_{\tilde{N}}$, if for every $p \in P_N \cap P_{\tilde{N}}$ for which $N(p) \neq \tilde{N}(p)$ it is true that $x(q) = y(\tilde{q})$ for all $q \succ p, q \in P_N$ and all $\tilde{q} \succ p, q \in P_{\tilde{N}}$, as well as $x(p) = y(p)$, then $x \approx y$.

A6 Non-Triviality: There exist two plans N, \tilde{N} such that P_N and $P_{\tilde{N}}$ are non-null, and $N(p_0) \neq \tilde{N}(p_0)$.

Axioms A1, A2, A3, A4, and A6 are rather standard. (See Gilboa and Schmeidler (1995).) A5 is a new axiom that captures the intuition that plans

only matter to the extent that they affect outcomes. If two plans are bound to yield the same utility, they are bound to be equivalent.

Theorem 5.1 *Assume that \succsim satisfies A6. Then the following are equivalent:*

- (i) \succsim satisfies A1-A5;
- (ii) *There is a support function S such that, for every two plans N, \tilde{N} , and every $x \in F_N, y \in F_{\tilde{N}}$,*

$$x \succsim y \quad \text{iff} \quad U(x) \geq U(y)$$

where, for $x \in F_N$

$$U(x) = \sum_{p \in P_N} w_S^N(p)x(p) .$$

Furthermore, in this case S (hence also w_S^N) is unique up to multiplication by a positive number.

Note that w_S^N is uniquely derived from S by (*) of Subsection 14.3.

15.3 Proof

That (ii) implies (i) is obvious. We prove the converse. Consider an act $a \in A_{p_0}$. Let $[N_a]$ be the set of plans that prescribe act a at position p_0 . We claim that P_N is null for some $N \in [N_a]$ iff it is null for all $N \in [N_a]$. Indeed, assume that P_N is null for $N \in [N_a]$ and consider $N' \in [N_a]$. By A6 there exists a plan \tilde{N} such that $P_{\tilde{N}}$ is non-null and $\tilde{N} \notin [N_a]$. By A2 there exists $y \in F_{\tilde{N}}$ such that $y \approx x$ for all $x \in F_N$. Consider an arbitrary $z \in F_{N'}$. Let $z', z'' \in F_{N'}$ be constant vectors that equal the minimal and maximal value of z , respectively, throughout $F_{N'}$. By A5, z', z'' are equivalent to the corresponding constant vectors in F_N and hence also to y . A3 implies that

$z \approx y$. This being true for all $z \in F_{N'}$, we conclude that $P_{N'}$ is null. We can therefore distinguish between sets $[N_a]$ where P_N is null for all $N \in [N_a]$, and sets $[N_a]$ where P_N is null for no $N \in [N_a]$. We refer to the former as sets of *null plans*. Observe that, by A5, all null plans are equivalent.

Next choose a plan $N_a \in [N_a]$ for each set that consists of non-null plans. By A6 there are at least two such sets. Note that $F_{N_a} \cap F_{N_b} = \emptyset$ for $a \neq b$. Restricting attention to $\hat{F} = \bigcup_{N_a} F_{N_a}$, we find that \succsim satisfies all the axioms of the representation theorem of Appendix B in Gilboa and Schmeidler (1995). It follows that there exist non-negative and non-zero vectors of weights s^{N_a} on P_{N_a} such that, for all $x, y \in \hat{F}$, $x \succsim y$ iff $U(x) \geq U(y)$, where, for $x \in F_{N_a}$, $U(x) = \sum_{p \in P_{N_a}} w^{N_a}(p)x(p)$. Further, the vectors are unique up to multiplication by a positive number.

Consider a plan $N'_a \in [N_a]$ where $N'_a \neq N_a$. Applying the same logic to the collection $\{N'_a\} \cup \{N_b\}_{b \neq a}$ we obtain representation by weights $\{w^{N'_a}\} \cup \{\hat{w}^{N_b}\}_{b \neq a}$ for all vectors in $\hat{F}' = \bigcup_{b \neq a} F_{N_b} \cup F_{N'_a}$. By uniqueness, there exists a $\lambda > 0$ such that $\lambda \hat{w}^{N_b} = w^{N_b}$ for $b \neq a$. By transitivity and A6, $\{\lambda w^{N'_a}\} \cup \{w^{N_b}\}$ induces a representation on all $\bigcup_b F_{N_b} \cup F_{N'_a}$. Continuing inductively, we obtain that for every non-null plan N there exists a non-negative and non-zero vector of weights w^N on P_N such that, for all $x, y \in \bigcup_{N \text{ non-null}} F_N$, $x \succsim y$ iff $U(x) \geq U(y)$, where, for $x \in F_N$, $U(x) = \sum_{p \in P_N} w^N(p)x(p)$. Further, the vectors are unique up to multiplication by a positive number.

We wish to show that for every pair of non-null plans, N, N', w^N and $w^{N'}$ satisfy

$$w^N(p) + \sum_{q \in P_N, q \succ p} w^N(q) = w^{N'}(p) + \sum_{q \in P_{N'}, q \succ p} w^{N'}(q) \quad (\bullet)$$

for every $p \in P_N \cap P_{N'}$. Consider $p \in P_N \cap P_{N'}$, and the following utility

function u :

$$\begin{aligned} u(p) &= 1; \\ u(q) &= 1 \quad \text{whenever } q \text{ follows } p \text{ in } N \text{ or in } N'; \\ u(q) &= 0 \quad \text{otherwise.} \end{aligned}$$

Let $x \in F_N$ and $y \in F_{N'}$ be the vectors induced by u . By A5, $x \approx y$. Hence (\bullet) follows.

Finally, we extend the representation to null plans. Define $w(p) = 0$ for all $p \in P_N$ where N is null. We have already established that all $x \in F_N$ are equivalent to a certain $y \in F_{\tilde{N}}$ where \tilde{N} is non-null. Since $x \approx 2x$ we can use A4 to conclude that $y \approx 2y$, and $U(y) = 0$ follows. This completes the proof of Theorem 5.1.

Chapter 6

Repeated Choice

16 Cumulative Utility Maximization

16.1 Memory-dependent preferences

The rule of U -maximization was suggested in Section 4 as a way to cope with uncertainty. But it can also be differently interpreted: rather than viewing the aggregation of past experiences in U as providing information regarding uncertain payoffs, it can be viewed as a dynamic model of changing preferences. To highlight this aspect of U -maximization, let us focus on a case with no uncertainty, and where the decision maker encounters a sequence of problems that are, according to her subjective judgment, identical. That is, for each act $i \in A$ there exists a unique outcome $r_i \in R$ such that only cases of the form (\cdot, i, r_i) may appear in the decision maker's memory with i as their second component. Further, the similarity function between any two problems is constant, say, 1.

Examples of such decisions abound. Every day Mary has to decide where to go for lunch. The problems she is facing on different days are basically identical. Moreover, every choice she makes would lead to a payoff that may be assumed known.¹ John decides every morning whether he should drive to

¹One may assume that the payoff is uncertain, but is drawn from the same distribution every time an act is chosen. As mentioned below, our results extend to the stochastic

work, use public transportation, or walk. In short, many everyday decisions are of this type. Our goal is to analyze the pattern of choices that would result from U -maximization in these problems.

The general model of U -maximization reduces to a simple process governed by few parameters: the decision maker has a cumulative index of satisfaction for each possible act i , denoted U_i , which depends on past choices. In every period, a maximizer of the satisfaction index is chosen. As a result of experiencing the outcome associated with the act, say, consuming the product, the satisfaction index of the act is updated. In the simplest version analyzed here, the update is additive, deterministic, and constant. That is, a number u_i is added to U_i whenever i is chosen (and U_i is unaffected by a choice of $j \neq i$). If it is positive, the alternative appears to be better the more it has been chosen. If negative, the alternative is less desirable as a result of experiencing it. The decision maker is implicitly assumed to be aware of U_i , but not necessarily of u_i .

Consider a decision maker for whom all alternatives appear equally desirable at the outset. As observed in Section 4, if there is at least one alternative i whose payoff exceeds the aspiration level, namely, $u_i > 0$, then, as soon as one such alternative is chosen, it will be chosen forever, and the decision maker may be said to be satisfied in the sense of Simon (1957). If, however, all alternatives have negative u_i 's, the decision maker will never be satisfied. A decision maker with a low aspiration level, namely, one for which $u_i > 0$, would tend to like an alternative more, the more it was chosen in the past, thus explaining habit-formation. By contrast, a high aspiration level implies that the decision maker would appear to be bored with alternatives she recently experienced, and would seek change.

Suppose that Mary can choose between a Pizza place and a Chinese restaurant for lunch. Let us assume that $u_{\text{Pizza}} = -1$ and that $u_{\text{Chinese}} = -2$, and that the initial U value is zero for both alternatives. It is easy to see

version of the model.

that, asymptotically, Mary will choose the Pizza place twice as frequently as she will the Chinese restaurant. Because she has a relatively high aspiration level, she does not stick to one lunch option forever. Moreover, the u -values of the alternatives determine the limit relative frequencies of choice: the latter are inversely proportional to the former. We later prove that this result holds in general.

One may interpret the vector of relative frequencies of choices as a bundle in neoclassical consumer theory, that is, as a vector of quantities of products. Viewed thus, if the aspiration level of the individual is low, she is easily satisfied and will select a corner point as an aggregate choice, as if she had a convex neoclassical utility function over the space of bundles. This corner point may not be optimal in the usual sense, namely, it may not be a maximizer of u_i . If, however, the aspiration level is relatively high, the consumer will keep switching among the various products. The relative frequencies with which the various products are consumed will converge to a limit, which will be an interior point, as if the individual were maximizing a concave neoclassical utility over the bundle space.

According to this interpretation, tastes, and thus decisions, are intrinsically context- or history-dependent. The decision maker does not attempt to maximize the function u . This function is only the rate by which U changes as result of experience. u may be interpreted as an index of instantaneous utility. But it is U , the cumulative satisfaction index, that the decision maker maximizes. Thus, U is the utility function. Whereas it changes as a function of history, it does so at a constant rate, given by u .

16.2 Related literature

Before presenting the formal model, we mention some related literature. Our discussion of small, repeated decisions may bring to mind Herrnstein and Prelec's "melioration theory". (See Herrnstein and Prelec (1991).) Indeed, the small decisions we study are very close to what they call "distributed

choice". Yet, our focus is different from theirs. Melioration theory deals primarily with choices that change intrinsically as a result of experience. In examples such as addictive behavior, or other problems relating to self-discipline (regarding work, savings, and the like), the payoff associated with an act depends on how often it has been chosen in the past. By contrast, we assume that the instantaneous utility of act i , u_i , is constant. Our decision maker may not wish to have the same meal every day; but there is no change in the pleasure derived from it as in the case of addiction, nor are there any long-term effects as in the case of savings. Further, melioration theory compares actual behavior to an optimum, defined by the highest average quality that can be obtained if inter-period interaction is taken into account. We study more commonplace phenomena, in which no such interaction exists. Correspondingly, in our model there is no well defined optimum, and, at least when high aspiration levels are involved, we do not deal with patently sub-optimal choice. Finally, the decision rules in the two models differ. Whereas melioration theory uses average, ours employs sums.

Our model is similar to the learning model of Bush and Mosteller (1955) in that, in every period, the outcome of the most recent choice is compared to an aspiration level, and the choice is reinforced if (and only if) its outcome exceeds the aspiration level. However, there are two main differences between these models. First, Bush and Mosteller are motivated by learning; thus, in their model past experiences serve mostly as sources of information, rather than determinants of preferences. Second, their model is inherently random, where a decision maker's history is summarized by a probability vector, describing the probability of each choice. By contrast, in our model the choice is deterministic. Stochastic choice is useful to summarize hidden variables. But our goal here is to expose the decision process, and we focus on the deterministic implications of our basic assumptions.

Gilboa and Pazgal (1996, 2000) extend the present model by assuming that the rate of change of U_i , namely, u_i , is a random variable, whose dis-

tribution depends only on the act i . They also conduct an empirical study of the random payoff model. Finally, the basic model suggested here was applied to electoral competition in Aragoes (1997). She found that change-seeking behavior on the part of the voters gives rise to ideological behavior of the competing parties.

16.3 Model and results

Let the set of alternatives be $A = \{1, 2, \dots, n\}$. For $i \in A$, denote the rate of the change of the utility function of i by u_i . Since the interesting case will be that of negative rates of change, it will prove convenient to define $a_i = -u_i$. For a sequence $x \in A^\infty$, define the number of appearances of choice $i \in A$ in x up to stage $t \geq 0$ to be

$$F(x, i, t) = \#\{1 \leq j \leq t | x(j) = i\} ,$$

and let $U(x, i, t)$ denote the cumulative satisfaction index of alternative $i \in A$ at that stage:

$$U(x, i, t) = F(x, i, t)u_i .$$

For a vector $u = (u_1, \dots, u_n)$ let $S(u)$ denote the set of all sequences of choices that are stagewise U -maximizing. That is,

$$S(u) = \{x \in A^\infty | \text{for all } t \geq 1, x(t) \in \arg \max_{i \in A} U(x, i, t - 1)\} .$$

Of special interest will be the relative frequencies of the alternatives, denoted

$$f(x, i, t) = \frac{F(x, i, t)}{t}$$

and their limit

$$f(x, i) \equiv \lim_{t \rightarrow \infty} f(x, i, t) .$$

(We will use this notation even if the limit is not guaranteed to exist.) Omission of the index of the alternative will be understood to refer to the corresponding vector:

$$\begin{aligned} f(x, t) &= (f(x, 1, t), \dots, f(x, n, t)) \\ f(x) &= (f(x, 1), \dots, f(x, n)) . \end{aligned}$$

Finally, denote

$$\begin{aligned} Y &= \{u = (u_1, \dots, u_n) \mid \forall i \quad u_i \neq 0\} \\ V &= \{u = (u_1, \dots, u_n) \mid \forall i \quad u_i < 0\} . \end{aligned}$$

We can now formulate our first result:

Proposition 6.1 *Assume that $u \in Y$. Then:*

- (i) *for all $x \in S(u)$, $f(x)$ exists;*
- (ii) *There exists $x \in S(u)$ for which $f(x)$ is one of the extreme points of the $(n-1)$ -dimensional simplex iff this is the case for all $x \in S(u)$, and this holds iff $u_i > 0$ for some $i \in A$;*
- (iii) *$f(x)$ is an interior point of the $(n-1)$ -dimensional simplex iff $u_i < 0$ for all $i \in A$;*
- (iv) *if $u \in V$ (i.e., $u_i > 0$) for all $i \in A$, then for all $x \in S(u)$ $f(x)$ is given by*

$$f(x, i) = \frac{\prod_{j \neq i} a_j}{\sum_{k \in A} \prod_{j \neq k} a_j} ;$$

- (v) *for every interior point y in the $(n-1)$ -dimensional simplex there exists a vector $u \in V$, unique up to a multiplicative scalar, such that $f(x) = y$ for all $x \in S(u)$.*

Remark: In the case that some utility values do vanish ($u \notin Y$), this result does not hold. Consider, for instance, the extreme case where $u_i = 0$ for all $i \in A$. Then we get $S(u) = A^\infty$ and, in particular, $f(x)$ need not exist for every $x \in S(u)$. Furthermore, one may get the entire $(n - 1)$ -dimensional simplex as the range of $f(x)$ when it is well defined.

When we aggregate choices over time, relative frequencies can be interpreted as quantities defining a bundle of products in the language of consumer theory. With this interpretation in mind, this result shows that low aspiration levels (positive u_i 's) are related to extreme solutions of the consumer's aggregate choice problem, whereas high aspiration levels (negative u_i 's) give rise to interior solutions. Similarly, high and low aspiration levels may explain change seeking behavior and habit formation, respectively.

Proposition 6.1 also shows that the interpretation of the instantaneous utility function u as "inverse frequency" holds in general: assume that all utility indices are negative, and consider two alternatives $i, j \in A$. The limit relative frequencies with which they will be chosen (for all $x \in S(u)$) are in inverse proportion to their utility levels:

$$\frac{f(x, i)}{f(x, j)} = \frac{a_j}{a_i} = \frac{u_j}{u_i}.$$

For example, if $n = 3$, $u_1 = -1$, $u_2 = -2$, and $u_3 = -3$, alternative 1 will be chosen twice as frequently as alternative 2, and 3 times as frequently as alternative 3. It follows that the relative frequencies of choice will converge to $(6/11, 3/11, 2/11)$. Thus the instantaneous utility of an alternative, u_i , has a different meaning than the utility value in standard choice theory, or in neoclassical consumer theory. On the one hand, the fact that alternative 2 has a lower utility than alternative 1 does not imply that the former will never be chosen. It will, however, be chosen less often than its substitute. On the other hand, the utility indices do not merely rank the alternatives; they also provide the frequency ratios, and are therefore cardinal.

Proposition 6.1 may be viewed as an axiomatization of the instantaneous

utility function u . Explicitly, for any choice sequence $x \in A^\infty$ the following two statements are equivalent: (i) there is a vector $u \in V$ such that $x \in S(u)$; (ii) $f(x)$ exists, it is strictly positive, and $x \in S(u)$ for any $u \in V$ satisfying $u_i f_i = u_j f_j$ for all $i, j \in A$. Thus the function u can be derived from observed choice, and it is unique up to a multiplicative scalar.

We also find that *any* point in the interior of the simplex may be the aggregate choice of a cumulative utility decision maker for an appropriately chosen instantaneous utility function. While it is comforting to know that the theory is general enough in this sense, one may worry about its meaningfulness. Are there any choices it precludes? The following result shows that in an appropriately defined sense, very few sequences are compatible with it.

Proposition 6.2 *Assume that $u \in Y$. Then:*

(i) *if $u_i > 0$ for some $i \in A$, then $S(u)$ is finite;*

(ii) *if $u_i < 0$ (i.e., $a_i > 0$) for all $i \in A$, then*

$$|S(u)| = \begin{cases} n! & \text{if } a_i/a_j \text{ is irrational for all } i, j \in A \\ \aleph & \text{otherwise} \end{cases}$$

(iii) *denote*

$$S^- = \bigcup_{u \in V} S(u)$$

$$S^+ = \bigcup_{u \in Y \setminus V} S(u) .$$

Let $p = (p_i)_{i \in A}$ be some probability vector on A , and let λ_p be the induced product measure on A^∞ (endowed with the product σ -field). Then S^- is an uncountable subset of a λ_p -null set, whereas S^+ is finite, and if p is not degenerate, S^+ is a λ_p -null set.

Thus there are uncountably many sequences of choices that are U -maximizing. Yet part (iii) of the proposition states that, overall, the set $S \equiv S^- \cup S^+$ is “small” by any of the reasonable measures defined above. Furthermore, it is easy to verify that finitely many observed choices may often suffice to conclude that they cannot be a prefix of any sequence in S .

To sum, the repeated choice theory presented here is non-vacuous. On the contrary, it is too easily refutable. However, Gilboa and Pazgal (1996) show that, if the u_i 's as well as the initial values $U(\cdot, i, 0)$ are random variables, one may choose their distributions so that every finite sequence of choices would have a positive probability. Yet, they also show that, with probability 1, there will be limit frequencies of choice as in the deterministic case, where the limit frequencies are inversely proportional to the expected instantaneous payoffs.

16.4 Comments

Fading Memory In this section we assume that the similarity function is constant. However, there are other ways in which one might model the fact that a problem is repeatedly encountered. For instance, the weight assigned to past cases may decrease exponentially with time, reflecting the fact that past cases may be forgotten, that they are deemed less relevant, or that their impact on future choices is diminishing with time. Incidentally, such a model does not require that the decision maker remember more than one number per alternative. (See Aragoes (1997).)

The Interpretation of Aspirations The notion of “aspiration level” in this model need not be taken literally. Consider, for instance, the choice of a piece of music to listen to. A decision maker may listen to “Don Giovanni” three times as often as to “Turandot”. In our model, the former would have a utility of -1 , while the latter – of -3 . Yet it would be wrong, if not blasphemous, to suggest that these works do not achieve the decision maker’s aspiration level. In this case, zero utility level may better be interpreted as

a “bliss point”. Thus, negative utility values need not conjure up sour faces in one’s mind. They may also describe perfectly content decision makers who merrily and gingerly alternate choices to keep their lives interesting. On the other hand, a satisfied decision maker need not be happy in any intuitive sense. For instance, consider a consumer who often suffers from acute headaches. If she is loyal to a certain brand of medication, she is satisfied according to our model. Yet she is by no means happy.²

16.5 Proofs

Proof of Proposition 6.1: First assume that $u_i > 0$ for some $i \in A$. After at most $(n - 1)$ choices of alternatives with negative utility, one with $u_i > 0$ is chosen. From that point on, that alternative is the unique U -maximizer and is chosen for ever. Hence in this case $f(x)$ exists for all $x \in S(u)$, and it is one of the extreme points of the $(n - 1)$ -dimensional simplex. Note, however, that it need not be the same for all such x ’s.

Let us now consider the case of $u_i < 0$ (i.e., $a_i > 0$) for all $i \in A$. At stage $t \geq 0$ alternative i is weakly preferred to j iff

$$U(x, i, t) = F(x, i, t)u_i \geq F(x, j, t)u_j = U(x, j, t)$$

or, equivalently, $F(x, i, t)a_i \leq F(x, j, t)a_j$.

Hence for *every* stage $t \geq 0$, regardless whether i is to be chosen at it or not, we have

$$F(x, i, t) \leq \frac{a_j}{a_i} F(x, j, t) + 1 .$$

Thus, for $t \geq 1$ we obtain

$$f(x, i, t) \leq \frac{a_j}{a_i} f(x, j, t) + \frac{1}{t} .$$

²Eva Gilboa-Schechtman has pointed out to us that the interpretation of zero utility level is likely to differ between pleasure seeking activities (such as going to a concert) and pain avoidance activities (such as taking a medication).

For $t \geq n$ none of the frequencies vanishes, and it follows that

$$\exists \lim_{t \rightarrow \infty} \frac{f(x, i, t)}{f(x, j, t)} = \frac{a_j}{a_i}.$$

With finitely many alternatives, this also implies that $f(x)$ exists. Furthermore, we find that it is independent of x . Finally, it follows that

$$f(x, i) = \frac{\prod_{j \neq i} a_j}{\sum_{k \in A} \prod_{j \neq k} a_j}.$$

To sum, we have proven that $f(x)$ exists in both cases, hence (i) is proven. If there is an alternative i with $u_i > 0$, only extreme points could be limit frequencies; on the other hand, if $u_i < 0$ for all $i \in A$, only interior points can be chosen. Thus the existence of one $x \in S(u)$ for which $f(x)$ is an extreme point implies that for some alternative i , $u_i > 0$, hence for all $x \in S(u)$ $f(x)$ is an extreme point of the simplex. This concludes the proof of (ii). Claims (iii) and (iv) were explicitly proven above.

We are left with claim (v), i.e., that for every interior point y in the $(n - 1)$ -dimensional simplex there exist negative utility indices $(u_i)_{i \in A}$ such that $f(x) = y$ for all $x \in S(u)$, and that these are unique up to a multiplicative scalar. Let there be given such a point $y = (y_1, \dots, y_n)$. Define $u_i = -a_i = -\prod_{j \neq i} y_j$. For all $i, j \in A$, $y_i/y_j = a_j/a_i$. It follows that $f(x) = y$ for all $x \in S(u)$. Furthermore, this equation shows that the utility $(u_i)_{i \in A}$ is unique up to multiplication by a positive scalar. This completes the proof. \square

Proof of Proposition 6.2: In light of the preceding analysis, (i) is immediate. Consider the case $u_i < 0$ for all $i \in A$. If a_i/a_j is irrational for all $i, j \in A$, after the first n stages, where each alternative is chosen once, U -maximization determines the choice uniquely. Thus there are $n!$ sequences in $S(u)$. If, however, a_i/a_j is rational for some $i, j \in A$, say $a_i/a_j = l/k$ for some integers $l, k \geq 1$, after exactly l choices of alternative j and k choices of i , there will come a stage where both i and j are among the maximizers of U . The choice at this stage is arbitrary, and therefore there are at least

two different continuations that are consistent with U -maximization. Since for every choice made at this point there will be a similar one after $2l$ choices of alternative j and $2k$ choices of i , there are at least 4 such continuations at this stage. By similar reasoning, there are at least

$$|\{0, 1\}^N| = \aleph$$

sequences in $S(u)$. Since \aleph is also an upper bound on $|S(u)|$, (ii) is established.

We need to show (iii). First consider S^- . Denote

$$S_p = \{x \in A^\infty \mid \exists f(x) = p\} .$$

By the strong law of large numbers, $\lambda_p(S_p) = 1$. Hence $\lambda_p(S^- \setminus S_p) = 0$. It therefore suffices to show that $\lambda_p(S^- \cap S_p) = 0$. First consider the case where p_i/p_j is irrational for all $i, j \in A$. Then $S^- \cap S_p$ is finite and of measure zero. Next assume that $p_i/p_j = l/k$ for some integers $l, k \geq 1$. In this case $S^- \cap S_p$ is uncountable, but it is a subset of the event in which for every $m \geq 1$, the $(ml + 1)$ -th appearance of j occurs after the mk -th appearance of i , and the $(mk + 1)$ -th appearance of i occurs after the ml -th appearance of j . Since at every point there is a positive λ_p -probability of a sequence of $(mk + 1)$ consecutive appearances of i , this event is of measure zero.

Finally, the set S^+ is finite. Furthermore, it is λ_p -null unless p is a degenerate probability vector. This completes the proof of the proposition.

□

17 The Potential

17.1 Definition

In this chapter we re-consider the decision process described in Section 16 as a maximization of a function that is defined on the simplex of relative frequencies. This analysis serves two purposes. First, it allows us to compare the asymptotic behavior of the dynamic process to neoclassical consumer theory, which is atemporal. Second, it provides a useful theoretical tool. In particular, we employ the potential to study the effect of act similarity in repeated choice.

For a choice vector $x \in A^\infty$, let $x_i \in A^t$ be the t -prefix of x . Let \cdot denote vector concatenation. Denote the set of all prefixes by $A^* = \bigcup_{t \geq 0} A^t$. For a function $\varphi : A^* \rightarrow \mathbb{R}$, $x_i \in A^t$, and $i \in A$, define

$$\frac{\partial \varphi}{\partial i}(x_i) = \varphi(x_i \cdot (i)) - \varphi(x_i) .$$

That is, $\frac{\partial \varphi}{\partial i}$ is the change in the value of φ that will result from adding i to the choice vector x_i . Denote by $U_i : A^* \rightarrow \mathbb{R}$ the U -value of alternative i . That is, $U_i(x_t) = U(x, i, t)$. Then u_i may be viewed as the derivative of U_i with respect to experiencing alternative i :

$$\frac{\partial U_i}{\partial i}(x_t) = u_i$$

for all $x_t \in A^*$. Similarly, for $j \neq i$,

$$\frac{\partial U_i}{\partial j}(x_t) = 0 .$$

Suppose we wish to measure the well-being of the decision maker at a certain time. One possibility to do so is by the sum of all past experiences. Define $Q : A^* \rightarrow \mathbb{R}$ by

$$Q(x_t) = \sum_{r=1}^t U_{x(\tau)}(x_{\tau-1})$$

for $x_\tau \in A'$. A U -maximizing decision maker may be described as maximizing her Q function at any given time t . Furthermore, Q is uniquely defined (up to a shift by an additive constant) by

$$\frac{\partial Q}{\partial i}(x_t) = U_i(x_t)$$

for all $x_i \in A^*$ and $i \in A$. Hence Q is a single function, such that the utility of alternative i , U_i , is its derivative with respect to i . It thus deserves the title *the potential of the utility*.

In certain ways, the potential is closer to the utility function in neoclassical consumer theory than either U or u are. Both u and U are defined for a single alternative, whereas the neoclassical function is defined for bundles. Correspondingly, if past experiences linger in the decision maker's memory, neither U nor u may qualify as measures of overall well-being, while the neoclassical utility does. By contrast, the potential function Q is defined for bundles (implicit in the vector x_t), and may be interpreted as a measure of well-being.

Furthermore, since

$$\frac{\partial Q}{\partial i} = U_i \quad \text{and} \quad \frac{\partial U_j}{\partial j} = \begin{cases} u_i & i = j \\ 0 & i \neq j \end{cases},$$

we get

$$\frac{\partial^2 Q}{\partial j \partial i} = \begin{cases} u_i & i = j \\ 0 & i \neq j \end{cases}.$$

If we were to define convexity and concavity of Q by its second derivatives, we would find that Q is concave if $u_i < 0$ (for all $i \in A$) and convex if $u_i > 0$. Indeed, we have found earlier (see Proposition 6.1) that negative u values induce limit frequencies of choice that are also predicted by a concave neoclassical utility function, while positive u values correspond to a convex one. Thus the potential parallels the neoclassical utility function also in

terms of the relationship between concavity/convexity and interior/corner solutions.

From a mathematical viewpoint the potential is a rather different creature from the neoclassical utility function. While the former is defined on choice sequences, the latter is defined on bundles. However, since all permutations of a given sequence are equivalent in terms of the behavior they induce, a sequence may be identified with the relative frequencies of the alternatives in it. It follows that, after an appropriate normalization, we may re-define the potential on the simplex of bundles.

17.2 Normalized potential and neo-classical utility

Formally, for $x \in A^\infty$ and $t \geq 0$, recall that $F(x, i, t)$ denotes the number of appearances of i in x up to time t . Let $T(x, i, k)$ stand for the time at which the k -th appearance of i in x occurs, that is,

$$T(x, i, k) = \min\{t \geq 0 \mid F(x, i, t) \geq k\} .$$

This function will be taken to equal ∞ if the set on the right is empty. However, for the case $u_i < 0$, T will be finite. We obtain

$$\begin{aligned} Q(x_t) &= \sum_{\tau=1}^t U(x, x(\tau), \tau - 1) \\ &= \sum_{i=1}^n \sum_{k=1}^{F(x,i,t)} U(x, i, T(x, i, k) - 1) \\ &= \sum_{i=1}^n \sum_{k=1}^{F(x,i,t)} (k - 1)u_i \\ &= \frac{1}{2} \sum_{i=1}^n u_i F(x, i, t) [F(x, i, t) - 1] . \end{aligned}$$

Recall that $f(x, i, t)$ is the relative frequency of i in x up to time t . Thus,

$$\frac{Q(x_t)}{t^2} = \frac{1}{2} \sum_{i=1}^n u_i f(x, i, t) \left[f(x, i, t) - \frac{1}{t} \right] .$$

For a point $f = (f_1, \dots, f_n)$ in the $(n - 1)$ -dimensional simplex, define the *normalized potential* to be

$$q(f) = \frac{1}{2} \sum_{i=1}^n u_i f_i^2 .$$

Then, for large enough t ,

$$\frac{Q(x_t)}{t^2} \approx q(f(x, t)) .$$

At any given time $t \geq 0$, the decision maker has a value for the normalized potential, and behaves as if she were trying to (approximately) maximize it. To be precise, the decision maker chooses an alternative so as to maximize $Q(x_{t+1})$, or, equivalently, to maximize $Q(x_{t+1})/(t + 1)^2$. However, in the long run this is approximately equivalent to maximization of the normalized potential q . Considering the optimization problem

$$\begin{array}{ll} \text{MAX} & q(f) \\ \text{s.t.} & \sum_{i=1}^n f_i = 1 \\ & f_i \geq 0 , \end{array}$$

it is straightforward to check that, should it have an interior solution (relative to the simplex), the solution must satisfy

$$f_i u_i = \text{const} .$$

Indeed, Proposition 6.1 shows that, if $u_i < 0$ (for all $i \in A$), there is an interior solution satisfying the above condition. Furthermore, it shows that the “greedy”, or “hill climbing” algorithm implemented by our decision maker converges to this solution.

At this point it is tempting to identify the (normalized) potential with the neoclassical utility. However, some distinctions should be borne in mind.

First, the neoclassical utility is assumed to be globally maximized by a one-shot decision. The potential is only locally improved at every stage. Should the potential be convex, for instance, our decision maker may be satisfied without optimizing it. Second, the neoclassical utility is assumed to be maximized given the budget constraint. By contrast, the local improvement of the potential is unconstrained. In Gilboa and Schmeidler (1997b) we introduce prices into the model in a way that retains this feature, namely, that lets the consumer follow the gradient of greatest improvement, where prices and income are internalized into the evaluation of possible small changes.

Yet, as long as prices and income are ignored, and if the potential is concave, our model may be viewed as a dynamic derivation of neoclassical utility maximization. Specifically, a neoclassical consumer who happens to have a concave, quadratic, and additively separable utility function over the simplex of frequencies will end up making the same aggregate choice as the corresponding unsatisfied cumulative utility consumer. Our theory thus provides an account of how such a consumer gets to maximize her utility.

Starting with a neoclassical utility, and adopting the identification of quantities with frequencies, one may suggest a hill-climbing algorithm as a reasonable dynamic model of consumer optimization, independently of our theory. From this viewpoint, the cumulative utility model provides a cognitive interpretation of the gradients considered by the optimizing consumer.

17.3 Substitution and Complementarity

Consider a consumer who chooses a daily meal out of $\{beef, chicken, fish\}$. Suppose that she likes them to the same degree, but that she seeks change. Say,

$$u_{beef} = u_{chicken} = u_{fish} = -1 .$$

If the consumer judges beef as closer to chicken than fish is, after having chicken she might prefer fish to beef. To capture these effects, we introduce

a similarity function between products, $s_A : A \times A \rightarrow [0, 1]$, that reflects the fraction of one product's instantaneous utility that is added to another product's cumulative satisfaction index.

Specifically, every time product i is consumed, $s_A(j, i)u_i$ is added to U_j , for every product j , with $s_A(i, i) = 1$. For instance, suppose that

$$s_A(\text{beef}, \text{chicken}) = s_A(\text{chicken}, \text{beef}) = 1$$

where the other off-diagonal s_A values are zero. For these similarity values, when the choice set is $\{\text{chicken}, \text{fish}\}$ or $\{\text{beef}, \text{fish}\}$, the relative frequencies of consumption are $(1/2, 1/2)$. However, for the set $\{\text{beef}, \text{chicken}, \text{fish}\}$, they may be $(1/4, 1/4, 1/2)$ rather than $(1/3, 1/3, 1/3)$, as would be the case in the absence of product similarities. In fact, in this example *beef* and *chicken* are practically identical, and we can only say that the frequencies of $\{\text{beef-or-chicken}, \text{fish}\}$ will converge to $(1/2, 1/2)$.

One interpretation of the product-similarity function is that it measures substitutability: the closer are two products to be perfect substitutes, the higher is the similarity between them, and the greater will be the impact of consuming one on boredom with another. If the similarity between two products is zero, they are substitutes in the sense that both can still be chosen in the same problem, but the consumption of one of them does not affect the desirability of the other.

Next suppose that the recurring choice in our model is a purchase decision, rather than a consumption decision. In each period the consumer chooses a single product, but she also has a supermarket cart, or kitchen shelves, which allow her eventually to consume several products together. Having bought tomatoes yesterday and cucumbers today, the consumer may have a salad.

Following this interpretation, it is only natural to extend the product-similarity function to negative values in order to model complementarities. If the product-similarity between, say, *tea*, and *sugar* is negative, (and so are their utilities), having just purchased the former would make the consumer more likely to purchase the latter next.

Recall that the instantaneous utility function u may be interpreted as the consumption-derivative of U :

$$\frac{\partial U_i}{\partial i} = u_i .$$

In the presence of product-similarity, when product i is chosen, U_j is changed by $s_A(j, i)u(i)$. That is,

$$\frac{\partial U_j}{\partial i} = s_A(j, i)u_i .$$

Combining the above, the similarity function is the ratio of derivatives of the U functions:

$$s_A(j, i) = \frac{\partial U_j / \partial i}{\partial U_i / \partial i} .$$

Using the potential function, we may write

$$\begin{aligned} \frac{\partial U_j}{\partial i} &= \frac{\partial^2 Q}{\partial j \partial i} = s_A(j, i)u_i = \frac{\partial U_j / \partial i}{\partial U_i / \partial i} \cdot \frac{\partial U_i}{\partial i} \\ s_A(j, i) &= \frac{\partial^2 Q / \partial j \partial i}{\partial^2 Q / \partial i^2} . \end{aligned}$$

The neoclassical theory defines the substitution index between two products as the cross derivative of the utility function (with respect to product quantities). By comparison, the product similarity function is the cross derivative of the potential function, normalized by its second derivative with respect to one of the two. Furthermore, if we define substitutability between two products i and j as the impact consumption of i has on the desirability of j , namely $s_A(j, i)u_i$, it precisely coincides with the cross derivative of the potential. This reinforces the analogy between the potential in our model and the neoclassical utility.

In the presence of product similarity, the consumer's choice given a sequence of past choices x_t will not depend on the order of the products in x_t , though the potential generally will. The U -value of product i is

$$U(x, i, t) = \sum_{\tau=1}^t s_A(i, x(\tau))u_{x(\tau)} = \sum_{j=1}^n F(x, j, t)s_A(i, j)u_j$$

which depends only on the number of appearances of each product in x_t . However, the value of the potential is

$$Q(x_t) = \sum_{\tau=1}^t U(x, x(\tau), \tau - 1) = \sum_{\tau=1}^t \sum_{v=1}^{\tau-1} s_a(x(\tau), x(v)) u_{x(\tau)}(v) .$$

It is easy to see that the potential will be invariant with respect to permutations of x_t iff

$$s_A(j, i)u_i = s_A(i, j)u_j \quad \forall i, j \in A$$

that is, iff

$$\frac{\partial^2 Q}{\partial j \partial i} = \frac{\partial^2 Q}{\partial i \partial j} \quad \forall i, j \in A .$$

Note that this is the appropriate notion of symmetry in this model: the product similarity function itself may be symmetric without guaranteeing that the impact of consumption of i on desirability of j equals that of j on i . Rather, we define symmetry by $s_A(j, i)u_i = s_A(i, j)u_j$, that is, by the equality of the cross-derivatives of the potential. Under this assumption,

$$Q(x_t) = \frac{1}{2} \sum_{i=1}^n u_i F(x, i, t) [F(x, i, t) - 1] + \sum_{i=2}^n \sum_{j=1}^{i-1} s_A(i, j) u_j F(x, i, t) F(x, j, t) .$$

Since $s_A(i, i) = 1$ for all $i \in A$ we get

$$\frac{Q(x_t)}{t^2} = \frac{1}{2} \sum_{i, j \in A} s_A(i, j) u_j f(x, i, t) f(x, j, t) - \frac{1}{2t} \sum_{i=1}^n u_i f(x, i, t) .$$

Hence $\frac{Q(x_t)}{t^2}$ can be approximated by the normalized potential defined on the simplex by

$$q(f) = \frac{1}{2} f S f^r ,$$

where $f = (f_1, \dots, f_n)$ is a frequency vector and the matrix S is defined by $S_{ij} = s_A(i, j)u_j$.

Suppose that S is negative definite. In this case the hill-climbing algorithm implemented by a cumulative utility consumer will result in maximization of the normalized potential. Hence, if we start out with a quadratic and concave neoclassical utility u , it defines a matrix S for which the corresponding cumulative utility consumer behaves as if she were locally maximizing u . Furthermore, for any function u that can be locally approximated by a quadratic q , local u -maximization would result in similar choices to those of the cumulative utility consumer characterized by q .

Since S is symmetric, it can be diagonalized by an orthonormal matrix. That is, there exists an $(n \times n)$ matrix P with $P^t = P^{-1}$ such that $P^t S P$ is diagonal, with the eigenvalues of S along its main diagonal. Since the matrix P can be thought of as rotation in the bundle space, we may offer the following interpretation: the consumer is deriving utility from n “basic commodities”, which are the eigenvectors of S , and their u -values are the corresponding eigenvalues. (These would be negative for a negative definite S .) One such commodity may be, for instance, a certain combination of *tea* and *sugar*. However, the consumer can only purchase the products *tea* and *sugar* separately. By choosing the right mix of the products, it is as if the basic commodity was directly consumed.³

According to this interpretation, there is zero similarity between the basic commodities; that is, there are no substitution or complementarity effects between them. These effects among the actual products are a result of the fact that these products are ingredients in the desired basic commodities.

³Note that the analogy to the additively separable case is not perfect. Specifically, the constraint that the sum of frequencies be 1 is also rotated. Hence the negative eigenvalues are not related to frequencies as directly as in Section 16.

Chapter 7

Learning and Induction

This chapter studies several ways in which case-based decision makers learn from their experience. First, we discuss aspiration level adjustment and their impact on behavior. We then proceed to discuss examples in which the similarity function is also learned. Finally, we conclude with some comments on inductive reasoning.

18 Learning to Maximize Expected Payoff

18.1 Aspiration level adjustment

The reader may find decision makers who maximize the function U less rational than expected utility maximizers. The former satisfice rather than optimize. They count on their experience rather than attempt to figure out what will be the outcomes of available choices. They follow what appears to be the best alternative in the short run rather than plan for the long run. They use whatever information they happened to acquire rather than intentionally experiment and learn from a growing experience.

There are many applications for which we find this boundedly rational image of a decision maker quite plausible, and it can even pass the rationality test proposed in Section 2. Especially in novel situations, people often find it hard to specify the space of states of the world in a satisfactory way, let

alone to form a prior over it. Many if not most of the decisions taken by governments, for instance, are made in complex environment that have not been encountered before in precisely the same way.

It is therefore not surprising that the commandments of EUT, appealing as they are, are hard to follow. One does not have enough trials to figure out all the possible eventualities, not to mention their frequencies. Further, there is little value in investing in learning and experimentation, because the knowledge acquired will be obsolete before it gets to be used. Similarly, planning carefully for the long run may prove a futile endeavor.

However, when the environment is more or less fixed, and the situation may be modeled as a repeated choice problem, case-based decision makers do appear to be a little too naive and myopic. While we argue (see Section 12) that CBDT is not designed for these situations to begin with, we show here that the irrational and shortsighted aspect of CBDT is, for the most part, an artifact of the implicit assumption that aspiration levels do not change.

Indeed, the basic formulation of CBDT in Section 4 normalized the utility function in such a way that the utility ascribed to an untried alternative, namely, the aspiration level, was zero. While this normalization can be performed for each choice situation separately, there is no reason to assume that it can be done for all of them simultaneously, that is, that the aspiration level does not change over time, or as a result of past experiences. In this section we propose two properties of aspiration-level adjustment rules, which we find descriptively plausible in general. We show that in the special case of a repeated choice problem, these properties also guarantee optimality. Thus, these properties can also be supported on normative grounds. While there are many rules that may guarantee optimal choice in this special case, we hope to convince the reader that the rules discussed here are also fairly intuitive.

18.2 Realism and ambitiousness

We assume that the aspiration level is adjusted according to a rule that is both realistic and ambitious. Realism means that the aspiration level is set closer to the best average performance so far experienced. Ambitiousness can take one of two forms: it may simply imply that the initial aspiration level is high; alternatively, it may be modeled as suggesting that the aspiration level is increased in certain periods.

More precisely, we model realism by assuming an adaptive rule that sets the new aspiration level at a certain weighted average of the old one and the maximal average performance so far encountered. Thus, if all the acts that were attempted in the past failed to perform up to expectations, the latter will have to be scaled down. Conversely, if the aspiration level is exceeded by the performance of certain acts, it is gradually increased.¹

While we do not provide an axiomatic derivation of this property, we would like to motivate it from several distinct viewpoints. First, suppose that we read the aspiration level as an answer to the question, What can you reasonably hope for in the current problem? If a moderately rational decision maker is to provide the answer, some adaptation of the aspiration level to actual performance seems inevitable. Second, realism appears to be a plausible description of people's emotional reactions: people seem to be able to adapt to circumstances. For simplicity, we do not distinguish here between scaling-up and scaling-down. For our purposes it suffices that over the long run the aspiration level is adjusted. Third, one might argue for the optimality of realism. As we show in this section, adjusting aspiration levels in realistic and ambitious way leads to expected payoff maximizing choices. That is, we show that the combination of realism and ambitiousness is behaviorally optimal. But realism has an emotional justification as well.

¹As will be clear in the sequel, the specific adjustment rule is immaterial; it is crucial, however, that it is gradually pushing the aspiration level towards the actual best average performance.

If we interpret the aspiration level emotionally, when it is set too high the decision maker is bound to be disappointed. Thus, if one were to choose an aspiration level (say, for one's child), one would like to set it at a realistic level.²

The ambitiousness property may have two separate (though compatible) meanings: *static* ambitiousness simply states that the initial aspiration level is relatively high. A high initial aspiration level reflects the fact that our decision maker is aggressive, and entertains great expectations. Whether the decision maker's initial aspiration level is high enough will depend on a variety of psychological, sociological, and perhaps also biological factors. While our optimism assumption may not be universally true, it is not blatantly implausible either. (See Shepperd, Ouellette, and Fernandez (1996) for related empirical evidence.)

The second meaning that the ambitiousness assumption may take is *dynamic*: that is, that the decision maker never quite loses hope. Specifically, we will assume that in certain decision periods, the aspiration level is set to exceed the best average performance by a certain constant. In order to make this compatible with realism, we will allow these decision periods to become more and more infrequent. (As a matter of fact, for the optimality result we will require that the update periods have a limit frequency of zero.) That is, the longer is one's memory, the less does one tend to increase the aspiration level in this somewhat arbitrary manner. However, dynamic ambitiousness requires that these update periods never end. Regardless of the memory's length, a dynamically ambitious decision maker still sometimes stops to ask, Why can't I do better than that?

As in the case of static ambitiousness, the claim of this assumption to

²Continuing this line of reasoning, it would seem that lowering aspirations well below what could realistically be expected may increase happiness even further. But this would result in sub-optimal choices, as will be clear in the sequel. Moreover, this option may not be psychologically feasible in the long run. For a more detailed discussion of these issues, see Gilboa and Schmeidler (1996b).

descriptive validity can be qualified at best. Indeed, we are not trying to claim that all decision makers are realistic but ambitious, just as we do not believe that all people necessarily choose optimal acts in repeated problems. The main point is that the properties of realism and ambitiousness correspond to some general intuition, and they make sense beyond the special case in which a certain problem is encountered over and over again. It is reassuring to know that in this special case these general properties also guarantee optimal choice.

18.3 Highlights

The repeated problem we discuss in this section is akin to the multi-armed bandit problem (see Gittins (1979)): the decision maker is repeatedly faced with finitely many options (“arms”), each of which is guaranteed to yield an independent realization of a certain random variable (with finite expectation and variance). Our results are as follows: first assume that the decision maker is realistic and statically ambitious. Then, given the distributions governing the various arms, and any probability lower than 1, there exists a high enough initial aspiration level such that, with this probability at least, the limit frequency of the expected-utility maximizing acts will be 1. Thus the initial aspiration level depends both on the given distributions and on the desired probability with which this frequency will indeed be 1.

Our second result assumes a decision maker who is realistic and dynamically ambitious. We prove that for any given distributions, and any initial aspiration level, the limit frequency of the optimal acts will be 1 with probability 1. This is a stronger statement than the previous one. It is guaranteed that *almost always* the best acts will be almost-always chosen, and that the same algorithm obtains optimality for *any* given distributions. Thus, dynamic ambitiousness is safer than static ambitiousness. Roughly, it is more important not to lose hope than to have great expectations.

The intuition behind both results can be easily explained in the deter-

ministic case. Suppose that every time an act is chosen, it yields the same outcome. For the first theorem, assume that the decision maker starts out with a very high aspiration level. Thus all options seem unsatisfactory, and the decision maker switches from one to another, as prescribed by CBDT in case of negative utility values. Specifically, in this case the frequencies of choice are inversely proportional to the utility values, as shown in Chapter 6. Hence, a high aspiration level prods the decision maker to experiment with all options with approximately equal frequencies. But the aspiration level cannot remain high for long. As time goes by it is updated towards the best average performance so far encountered. In the deterministic case, the average performance of an act is simply its instantaneous utility value. Thus, in the long run the aspiration level tends to the maximal utility value. Correspondingly, an act that achieves this value is almost satisficing, and the set consisting of all these acts will be chosen with a limit frequency of 1.

Next consider a dynamically-ambitious decision maker in a similar set-up. Such a decision maker may have started out with too low an aspiration level; thus she may be choosing a sub-optimal act, while the aspiration level is being adjusted *upward* towards the utility value of this act. However, if at some point the aspiration level is set above this utility value, this act is no longer satisficing, and the decision maker will try a new one. In the long run all acts will be tried out, and the aspiration level will be realistically adjusted towards the maximal utility value. As opposed to the case of static ambitiousness, the aspiration level does not converge to this value, since it is pushed above it from time to time. Yet, these periods are assumed to have zero limit frequency, and thus the optimality result holds.

The general cases of both results, in which the available acts yield stochastic payoffs, are naturally more involved, but the proofs follow the same basic intuition.

We note here that both realism and ambitiousness are crucial for the optimality results. If our decision maker is realistic but not ambitious, she

may well choose a sub-optimal act forever. In this case her choice is random in the following sense: an act is randomly selected at the first stage, and then it is chosen for ever. On the other hand, if the decision maker is, say, statically ambitious but not realistic, then all choices seem to her almost equally unsatisfactory. In this case the choice is close to random in the sense that all acts will be chosen with approximately the same frequency. (See Chapter 6.) By contrast, the combination of the two guarantees that all acts will be experimented with, but also that in the long run experimentation will give way to optimal choice.

In a sense, our results may be viewed as explaining the evolution of optimal (expected-utility maximizing) choice: a case-based decision maker who is both realistic and ambitious will learn to be an expected-utility maximizer. These results only hold if the decision problem is repeated long enough in the same form. But this is precisely the case in which EUT seems most plausible, that is, when history is long enough to enable the decision maker to figure out what are the states of the world, and to form a prior over them based on their observed frequencies. Furthermore, a case-based decision maker is more open minded than an expected utility maximizer. While the latter may have a-priori beliefs whose support fails to contain the true distribution, the former does not entertain prior beliefs, and thus cannot be wrong about them.

In the context of optimization problems, one may view our results as reinforcing a general principle by which global optimization may be obtained by local optimization coupled with the introduction of noise. The annealing algorithms (Kirkpatrick, Gellatt, and Vecchi (1982)) are probably the most explicit manifestation of this principle. Genetic algorithms (Holland (1975)) are another example, in which the adaptive process leads to a local optimum of sorts, and the cross-over process allows the algorithm to explore new horizons. Yet another example of the same principle may be found in evolutionary models in game theory such as Foster and Young (1990), Kan-

dori, Mailath, and Rob (1993) and Young (1993). In these models, a myopic best-response rule may lead to equilibria that are Pareto dominated (“local optima”), even in pure-coordination games. But the introduction of mutations provides the “noise” that guarantees (in such games) a high probability of a Pareto-dominating equilibrium (a “global optimum”). From this viewpoint, one may interpret our results as follows: the realistic nature of the aspiration-level adjustment rules induces convergence to a local optimum, namely, to a high frequency of choice of the best acts *among those that were tried* often enough. Ambitiousness plays the role of the noise, which prods the decision maker to choose seemingly sub-optimal acts, and, in the long run, to converge to a global optimum.

Annealing algorithms simulate physical phenomena; genetic algorithms and evolutionary game theory models are inspired by biological metaphors; by contrast, our process is motivated by psychological intuition. As mentioned above, we find this intuition valid beyond the specific model at hand.

18.4 Model

We now turn to the formal model. Let $A = \{1, 2, \dots, n\}$ be a set of *acts* ($n \geq 1$). For $i \in A$ let there be given a distribution F_i on \mathbb{R} (endowed with the Borel σ -algebra), to be interpreted as the (conditional) distribution of the utility yielded by act i whenever it is chosen. We assume that F_i has finite expectation and variance, denoted μ_i and σ_i , respectively.

The underlying state space will be a subset of

$$S_0 = (\mathbb{R} \times A \times \mathbb{R})^{\mathbb{N}}$$

where \mathbb{N} denotes the natural numbers. A state $\omega = ((H_1, a_1, x_1), (H_2, a_2, x_2), \dots) \in S_0$ is interpreted as follows: for all $t \geq 1$, in period t , the aspiration level is H_t at the beginning of the period, an act a_t is chosen, and it yields a payoff of x_t . It will be convenient to define, for every $t \geq 1$, the projection functions

$$H_t, x_t : S_0 \rightarrow \mathbb{R} \quad \text{and} \quad a_t : S_0 \rightarrow A$$

with the obvious meaning.

Next we define a function $C : S_0 \times A \times \mathbb{N} \rightarrow 2^{\mathbb{N}}$ to be the set of periods, up to a given period, in which a given act was chosen, at a given state. That is,

$$C(\omega, i, t) = \{j < t \mid a_j(\omega) = i\} .$$

We are also interested in the number of times a certain act was chosen up to a given period. Therefore, we define a function $K : S_0 \times A \times \mathbb{N} \rightarrow \mathbb{N} \cup \{0\}$ by

$$K(\omega, i, t) = \#C(\omega, i, t) .$$

We are mostly interested in the relative frequencies of the decision maker's choices. It will be convenient to define a function $f : S_0 \times A \times \mathbb{N} \rightarrow [0, 1]$ to measure relative frequency up to a given time, that is,

$$f(\omega, i, t) = \frac{K(\omega, i, t)}{t} .$$

Dropping the period index will refer to the limit:

$$f(\omega, i) = \lim_{t \rightarrow \infty} f(\omega, i, t) .$$

Finally, we extend this notation to subsets of A : for $D \subset A$ we define

$$f(\omega, D, t) = \sum_{i \in D} f(\omega, i, t)$$

$$f(\omega, D) = \lim_{t \rightarrow \infty} f(\omega, D, t) .$$

We now turn to define the CBDT functions in this context. Let $U : S_0 \times A \times \mathbb{N} \rightarrow \mathbb{R}$ be defined by

$$U(\omega, i, t) = \sum_{j \in C(\omega, i, t)} [x_j(\omega) - H_t(\omega)] .$$

That is, U is the cumulative payoff of an act, measured relative to the current aspiration level H_t . Observe that past outcomes x_j (for $j < t$) are

re-evaluated in light on the new aspiration level H_t . Since the similarity function is assumed to be identically 1, this definition coincides with the function U of Chapter 2.

We will also use the notation

$$V(\omega, i, t) = \frac{U(\omega, i, t)}{K(\omega, i, t)} .$$

(Thus, “ $V(\omega, i, t)$ is well defined” means “ $K(\omega, i, t)$ is positive.”) As in the case of U , this definition coincides with the function V of Chapter 2 because the similarity of any two problems is assumed to be 1. Since the values of both U and V depend on the aspiration level H_t , it is convenient to have a separate notation for the absolute average utility of each act. Thus we denote

$$\bar{X}(\omega, i, t) = \frac{\sum_{j \in C(\omega, i, t)} x_j(\omega)}{K(\omega, i, t)} .$$

Note that $\bar{X}(\omega, i, t)$ is well defined whenever $V(\omega, i, t)$ is and that

$$\bar{X}(\omega, i, t) = V(\omega, i, t) + H_t(\omega) .$$

We now wish to express the fact that the decision maker considered is a U -maximizer. We do it by restricting the state space as follows: define $S_1 \subset S_0$ by

$$S_1 = \{\omega \in S_0 \mid a_t(\omega) \in \arg \max_{i \in A} U(\omega, i, t) \quad \forall t \geq 1\} .$$

Similarly, we further restrict the state space to reflect the fact that the aspiration level is updated in an adaptive manner. First define, for $t \geq 2$ and $\omega \in S_0$, the relative and absolute maximal average performance to be, respectively,

$$\begin{aligned} \bar{V}(\omega, t) &= \max\{V(\omega, i, t) \mid i \in A, K(\omega, i, t) > 0\} \\ \bar{X}(\omega, t) &= \max\{\bar{X}(\omega, i, t) \mid i \in A, K(\omega, i, t) > 0\} . \end{aligned}$$

Next, for a given $\alpha \in (0, 1)$ and $H_1 \in \mathbb{R}$ we finally define the state space to be

$$\Omega = \Omega(\alpha, H_1) = \left\{ \omega \in S_1 \mid \begin{array}{l} H_1(\omega) = H_1 \text{ and for } t \geq 2 \\ H_t(\omega) = \alpha H_{t-1}(\omega) + (1 - \alpha) \bar{X}(\omega, t) \end{array} \right\} .$$

Endow S_0 with the σ -algebra generated by the Borel σ -algebra on (each copy of) \mathbb{R} and by the algebra 2^A on (each copy of) A . Let $\Sigma = \Sigma(\alpha, H_1)$ be the induced σ -algebra on Ω . Finally, we turn to define the underlying probability measure. Given Ω and Σ , a probability measure P on Σ is *consistent with* $(F_i)_{i \in A}$ if for every $t \geq 1$ and $i \in A$, the conditional distribution of x_t that it induces, given that $a_t = i$, is F_i , and, furthermore, x_t is independent (according to P) of the random variables $H_1, a_1, x_1, \dots, H_{t-1}, a_{t-1}, x_{t-1}, H_t$. Notice that distinct measures on Σ that are consistent with $(F_i)_{i \in A}$ can only disagree regarding the choice of an act a_t where $\arg \max_{i \in A} U(\omega, i, t)$ is not a singleton.

18.5 Results

Our first result is:

Theorem 7.1 *Let there be given $A = \{1, \dots, n\}$, $(F_i)_{i \in A}$ as above, $\alpha \in (0, 1)$ and $\varepsilon > 0$. There exists $H_0 \in \mathbb{R}$ such that for every $H_1 \geq H_0$ and every measure P on $(\Omega(\alpha, H_1), \Sigma(\alpha, H_1))$ that is consistent with $(F_i)_{i \in A}$,*

$$P(\{\omega \in \Omega \mid \exists f(\omega, \arg \max_{i \in A} \mu_i) = 1\}) \geq 1 - \varepsilon .$$

The theorem guarantees that, if we focus on those states ω at which there is a limit choice frequency for the set of expected utility maximizers, and it equals 1, this set is measurable and it has arbitrarily high probability provided that the initial aspiration level is high enough.

Note that Theorem 7.1 cannot guarantee an aspiration level that is uniformly high enough for all given distributions $(F_i)_{i \in A}$. Indeed, it is obvious

that any initial aspiration level may turn out to be too low. By contrast, our second result guarantees optimality for all given distributions, regardless of the initial aspiration level and with probability 1. The assumption that drives this much stronger conclusion is that the aspiration level is “pushed up” every so often. That is, that in a certain set of periods, which is infinite but sparse (i.e., has a limit frequency of zero), the aspiration level is not adjusted by averaging its previous value and the best average performance value; rather, in these periods it is set to be at some level above the best average performance value, regardless of the previous aspiration level.

Formally, we define a new probability space as follows. Let there be given $H_1 \in \mathbb{R}$ and $\alpha \in (0, 1)$ as above. Assume that $N_A \subset \mathbb{N}$ and $h > 0$ are given. N_A is interpreted as the set of periods in which the decision maker is ambitious. The number h should be thought of as the increase in the aspiration level. Define

$$\begin{aligned} \Omega &= \Omega(\alpha, H_1, N_A, h) \\ &= \left\{ \omega \in S_1 \left| \begin{array}{l} H_1(\omega) = H_1 \text{ and for } t \geq 2 \\ H_t(\omega) = \bar{X}(\omega, t) + h \text{ if } t \in N_A \\ H_t(\omega) = \alpha H_{t-1}(\omega) + (1 - \alpha)\bar{X}(\omega, t) \text{ if } t \notin N_A \end{array} \right. \right\}. \end{aligned}$$

Next define $\Sigma = \Sigma(\alpha, H_1, N_A, h)$ to be the corresponding σ -algebra. Similarly, a measure P on Σ is defined to be *consistent with* $(F_i)_{i \in A}$ as above.

We can now state:

Theorem 7.2 *Let there be given $A = \{1, \dots, n\}$, $(F_i)_{i \in A}$ as above, $\alpha \in (0, 1)$, $H_1 \in \mathbb{R}$, $N_A \subset \mathbb{N}$ and $h > 0$. If N_A is infinite but sparse, then for every measure P on $(\Omega(\alpha, H_1, N_A, h), \Sigma(\alpha, H_1, N_A, h))$ that is consistent with $(F_i)_{i \in A}$,*

$$P(\{\omega \in \Omega \mid \exists f(\omega, \arg \max_{i \in A} \mu_i) = 1\}) = 1.$$

18.6 Comments

Hybrid of Summation and Average As briefly mentioned above, the decision rule we use here is a hybrid: our decision makers choose acts by U -maximization; however, when it comes to adjusting their aspiration levels, they use the maximal V value. This apparent inconsistency calls for an explanation.

Following the discussion in Section 4 and Chapter 6, we would like to suggest the following interpretation. In general, memory affects one's decisions in two ways: first, as a source of information, which is especially crucial for decisions under uncertainty; second, as a primary effect in a dynamic choice situation. Memory helps one to reason about the world, but it also changes one's tastes.

Thus, there are two fundamental questions to which memory is key: first, "What do I want to do now?" and second, "What do I think of this act?" In answering the first question, memory plays a dual role: as a source of information and as a factor affecting preferences. In answering the second, memory only serves as a source of information. Correspondingly, we would like to suggest that U offers an answer to the first, whereas V – to the second.

Consider the following example: every day Mary has to choose a restaurant. She chooses a U -maximizer. Assume that her aspiration level is high and that she therefore exhibits change seeking behavior. Assume next that Mary has a guest, and he asks her which is the best restaurant in town, namely, which restaurant should one go to if one has only one day to spend there (with no memory). Then, according to this interpretation, Mary will recommend a V -maximizing, rather than a U -maximizing act. Asked why she is not choosing this restaurant herself, Mary might say, Oh, I was there just yesterday. Having visited it recently, Mary attaches to the restaurant a relatively low U value. But the very fact that the restaurant was recently chosen need not change its V value.

The optimality rule discussed in this paper is therefore not as inconsistent

as it may appear at first glance: our decision makers are U -maximizers in their choices. This means that memory enters their decision considerations not only as a source of information. With a high aspiration level, this also allows them to keep switching among the alternative acts, and to continue trying acts whose past average performance happened to be poor. On the other hand, asking themselves, “What can I reasonably hope for?” or “What choice would I recommend to someone who hasn’t tried any of the options?” they base their answers on V -maximizing acts. As we have shown, adjusting their aspiration level based on the maximal V value also colors past experiences differently. In the long run, the dissatisfaction with V -maximizing acts decreases, and thus their relative frequency tends to 1.

The above notwithstanding, it certainly makes sense to consider two simpler alternatives. The first, “ U -rule”, prescribes that decisions will be made so as to maximize U , and that the aspiration level will be adjusted according to U as well. The second is the corresponding “ V -rule”. Both rules seem sensible, but none of them guarantees optimal choice in the long run. Using the U -rule, the aspiration level need not converge. (As a matter of fact, it is not obvious what is the right way to define the aspiration level adjustment rule in this case.) Using the V -rule, the decision maker may never retry certain alternatives that happened to have particularly low realizations in the first few periods. We omit the simple examples.

Infinitely Many Acts The discussion in this section focused on finitely many alternatives. Considering generalizations to large sets of acts, one would like to take into account the fact that infinitely many acts are typically endowed with additional structure. For instance, prices and quantities may be modeled as continua, which also have a metrizable topology. It is natural to reflect this topology in an act similarity function. For instance, having set a price at \$20, a seller may have some idea about the outcomes that are likely to result from a price of \$20.01. Since these two acts are similar, past experience with one of them enters the evaluation of another. Given a compact metric topological

space of acts and a similarity function that is monotone with respect to the metric, and assuming continuity of u , one would expect a similar optimality result to hold.

Other Adjustment Rules Our results do not hinge on the specific aspiration level adjustment rule. First, the aspiration level need not be adjusted in every period, and the adjustment periods do not have to be deterministically set. All that is required is that there be infinitely many of them with a high enough probability. Similarly, the realistic adjustment need not be done by a weighted average with fixed weights. The conclusion of Theorem 7.1 will hold whenever the following two conditions are satisfied. (i) The adjustment process guarantees convergence, that is, that for all $a, b \in \mathbb{R}$ and $\varepsilon > 0$, if $\bar{X}(\omega, t) \in (a, b)$ for all $t \geq T_1$ for some T_1 , then there exists T_2 such that for almost all $t \geq T_2$, $H_t(\omega) \in (a - \varepsilon, b + \varepsilon)$. (ii) The adjustment is not too fast: for all $R \in \mathbb{R}$ and all $T_0 \geq 1$ there is a number H_0 such that for all $H_1 > H_0$ and all $t \leq T_0$, $H_t > R$. Similarly, the conclusion of Theorem 7.2 holds under the following conditions: (i) the adjustment process guarantees convergence as above; and (ii) the aspiration level increases over an infinite but sparse set of periods. Neither the increase h nor the set of increase periods N_A need be deterministic or exogenously given. Both may depend of the state ω , on past acts, and on their results. It is essential, however, that, for each ω , h be bounded away from zero (and not too large), and that N_A be infinite but sparse.

Finally, one may assume that in the periods of ambitiousness, the aspiration is set so as to exceed (by h) its own previous value, rather than the maximal average performance level.

Retrospective Evaluation Note that when the aspiration level is updated in our model, the u value of past experiences is also updated. That is, outcomes that have been obtained in the past are re-evaluated according to the newly defined aspiration level. Thus we implicitly assume that the decision maker can reflect upon the outcomes themselves, sometimes realizing that they were

not as unsatisfactory as they seemed at the time.

Alternatively, one may assume that only the utility value of past experiences is retained in memory, and that the original evaluation of an outcome will be forever used to judge the act that led to it. Our first result does not hold in this case, since a very high initial aspiration level may make an expected utility maximizing act have a very low U value, to the extent that it may never be chosen again. While one may argue for the psychological plausibility of the alternative assumption, it seems that it is “more rational” to re-evaluate outcomes based on an adjusted aspiration level, rather than compare each outcome to a possibly different aspiration level. At any rate, the second result holds under the alternative assumption as well: having infinitely many periods in which the expected utility of any act is a negative number bounded away from zero guarantees that all acts will be chosen infinitely often with probability 1.

Games Pazgal (1997) has shown that if realistic but ambitious case-based decision makers repeatedly play a game of common interest, they will converge to the optimal play. His definition of realism is slightly different from the one we use here, in that the players are assumed to update their aspiration levels toward the best payoff encountered. Thus it implicitly assumes that the players know that the game is of common interest.

18.7 Proofs

Proof of Theorem 7.1: A few words on the strategy of the proof are probably in order. The general idea is very similar to the deterministic case described above: let the initial aspiration level be high enough so that each act is chosen a large enough number of times, and then notice that the aspiration level tends to the maximal expected utility. In the deterministic case, each act should be chosen at least once in order to get its average performance \bar{X} equal to its expectation. In the stochastic case, more choices are needed, and a law of large numbers will be invoked for a similar conclusion. Thus

the initial aspiration level should be high enough to guarantee that each of the acts is chosen enough times to get the average close to the expectation.

If the supports of the given distributions F_i were bounded, one could find high enough aspiration levels such that all possible realizations of all possible choices seem similarly unsatisfactory. This would guarantee that, as long as the aspiration level is beyond a certain bound, all acts are chosen with similar frequencies, and therefore all of them will be chosen enough times for the law of large numbers to apply. However, these distributions need not have a bounded support. They are only known to have a finite variance. Thus the proof is slightly more involved, as we explain below.

Let us first assume, without loss of generality, that for some $r \leq n$,

$$\mu_1 = \mu_2 = \cdots = \mu_r > \mu_{r+1} \geq \cdots \geq \mu_n .$$

Furthermore, we assume that $r < n$. (The theorem is trivially true otherwise.) Next denote

$$I = \arg \max_{i \in A} \mu_i = \{1, 2, \dots, r\}$$

$$\delta = \frac{\mu_1 - \mu_{r+1}}{3} .$$

The number δ is so chosen that, if the average values are δ -close to the corresponding expectations, then the maximal average value is obtained by a maximizer of the expectation.

We now turn to find the number of times that is needed to guarantee, with high enough probability, that the averages are, indeed, δ -close to the expectations. Given $\varepsilon > 0$ as in the theorem and $i \in A$, let $K_i \geq 1$ be such that: for every $k \geq K_i$ and every sequence of i.i.d random variables $X_i^1, X_i^2, \dots, X_i^k$, each with distribution F_i ,

$$Pr\left(\left|\frac{1}{k} \sum_{j=1}^k (X_i^j - \mu_i)\right| \leq \delta\right) \geq (1 - \varepsilon)^{1/2n}$$

where Pr is the measure induced by the distribution F_i . Notice that such K_i exists by the strong law of large numbers. (See, for instance, Halmos (1950).) Let $K = \max_{i \in A} K_i$.

We now turn to the construction of the initial aspiration level. As explained above, we would like to be able to assume that the F_i 's have bounded supports, in order to guarantee that each act is chosen at least K times. We will therefore find an event with a high enough probability, on which the random variables x_t are, indeed, bounded.

We start by finding, for each $i \in A$, bounds $\underline{b}_i, \bar{b}_i \in \mathbb{R}$ such that, for any random variable X_i distributed according to F_i ,

$$Pr(\underline{b}_i \leq X_i \leq \bar{b}_i) \geq (1 - \varepsilon)^{1/4nK}$$

where Pr is some probability measure that agrees with F_i . Notice that such bounds exist since F_i has a finite variance. Without loss of generality assume also that $\bar{b}_i > \mu_i + 2\delta$ for all $i \in A$. Next define

$$\underline{b} = \min_{i \in A} \underline{b}_i \quad \text{and} \quad \bar{b} = \max_{i \in A} \bar{b}_i .$$

The critical lower bound on the aspiration level (for the “experimentation” period, during which every act is chosen at least K times) is chosen to be

$$R = 2\bar{b} - \underline{b} .$$

Let us define, for every $T \geq 1$, the event

$$B_T = \{\omega \in \Omega \mid \forall t \leq T, \underline{b} \leq x_t(\omega) \leq \bar{b}\} .$$

Notice that, since the given measure P is consistent with $(F_i)_{i \in A}$, $P(B_T) \geq (1 - \varepsilon)^{T/4nK}$. Hence, provided that T is not too large, B_T will have a high enough probability. In order to show that T need not be too large to get enough ($\geq K$) observations of each act, we first show that, on B_T and with sufficiently high aspiration level, the first T choices are more or less evenly distributed among the acts:

Claim 7.1 *Let there be given $T \geq n$, and $\omega \in B_T$. Assume that for all $t \leq T$, $H_t(\omega) > R$. Then for all $i, j \in A$ and all t such that $n \leq t \leq T$,*

$$K(\omega, i, t) \leq 2K(\omega, j, t) .$$

Proof: Assume the contrary, and let t_0 be the minimal time t such that $n \leq t \leq T$ and

$$K(\omega, i, t_0) > 2K(\omega, j, t_0)$$

for some $i, j \in A$. Notice that $K(\omega, a, n) = 1$ for all $a \in A$, hence $t_0 > n$. It follows from minimality of t_0 that $a_{t_0-1}(\omega) = i$, i.e., that i was the last act chosen.

Consider the following bounds on the U values of the two acts:

$$\begin{aligned} U(\omega, i, t_0 - 1) &\leq K(\omega, i, t_0 - 1)(\bar{b} - H_{t_0-1}(\omega)) \\ U(\omega, j, t_0 - 1) &\geq K(\omega, j, t_0 - 1)(\underline{b} - H_{t_0-1}(\omega)) . \end{aligned}$$

The optimality of i at stage $t_0 - 1$ implies

$$\begin{aligned} U(\omega, i, t_0 - 1) &\geq U(\omega, j, t_0 - 1) ; \\ K(\omega, i, t_0 - 1)(\bar{b} - H_{t_0-1}(\omega)) &\geq K(\omega, j, t_0 - 1)(\underline{b} - H_{t_0-1}(\omega)) . \end{aligned}$$

Recalling that $H_{t_0-1}(\omega) > R \geq \bar{b} \geq \underline{b}$, this is equivalent to

$$\frac{K(\omega, j, t_0 - 1)}{K(\omega, i, t_0 - 1)} \geq \frac{\bar{b} - H_{t_0-1}(\omega)}{\underline{b} - H_{t_0-1}(\omega)} .$$

By minimality of t_0 we know that

$$\frac{K(\omega, j, t_0 - 1)}{K(\omega, i, t_0 - 1)} = \frac{1}{2} .$$

We therefore obtain

$$\underline{b} - H_{t_0-1}(\omega) \leq 2(\bar{b} - H_{t_0-1}(\omega))$$

which implies

$$H_{t_0-1}(\omega) \leq 2\bar{b} - \underline{b} = R ,$$

a contradiction. \square

We now set $T_0 = 2nK$, and will prove that, as long as the aspiration level is kept above R , after T_0 stages, each act will be chosen at least K times on the event B_{T_0} . Formally,

Claim 7.2 *Let there be given $\omega \in B_{T_0}$ and assume that $H_t(\omega) > R$ for all $t \leq T_0$. Then for $i \in A$,*

$$K(\omega, i, T_0) \geq K .$$

Proof: If $K(\omega, i, T_0) < K$ for some $i \in A$, then by Claim 7.1, $K(\omega, j, T_0) < 2K$ for all $j \in A$. Then we get

$$T_0 = \sum_{j \in A} K(\omega, j, T_0) < 2nK = T_0 ,$$

which is impossible. \square

We finally turn to choose the required level for the initial aspiration level. Choose a value

$$H_0 = H_0(\varepsilon) > \underline{b} + 2\left(\frac{1}{\alpha}\right)^{T_0}(\bar{b} - \underline{b})$$

and let us assume for the rest of the proof that $H_1 \geq H_0$. We verify that this bound is sufficiently high in the following:

Claim 7.3 *Let there be given $\omega \in B_{T_0}$, and assume that $H_1 \geq H_0$. Then for all $t \leq T_0$, $H_t(\omega) > R$.*

Proof: For all $1 < t \leq T_0$,

$$H_t(\omega) \geq \alpha H_{t-1}(\omega) + (1 - \alpha)\underline{b}$$

or

$$H_t(\omega) - \underline{b} \geq \alpha(H_{t-1}(\omega) - \underline{b}) .$$

Hence

$$\begin{aligned} H_t(\omega) - \underline{b} &\geq \alpha^t(H_1 - \underline{b}) > 2\alpha^t\left(\frac{1}{\alpha}\right)^{T_0}(\bar{b} - \underline{b}) \geq 2(\bar{b} - \underline{b}) \\ H_t(\omega) &> 2\bar{b} - \underline{b} = R . \quad \square \end{aligned}$$

Combining the above, we conclude that, for $H_1 \geq H_0$, $K(\omega, i, T_0) \geq K$ for all $\omega \in B_{T_0}$ and all $i \in A$. Furthermore, for a measure P , consistent with $(F_i)_{i \in A}$,

$$P(B_{T_0}) \geq (1 - \varepsilon)^{T_0/4nK} = (1 - \varepsilon)^{1/2} .$$

We now define the event on which the limit frequency of the expected-utility maximizing acts is 1: let $B \subset B_{T_0}$ be defined by

$$B = \left\{ \omega \in B_{T_0} \left| \begin{array}{l} \forall t \geq T_0, \quad \forall i \in A, \\ |\bar{X}(\omega, i, t) - \mu_i| < \delta \end{array} \right. \right\} .$$

By the choice of K and the independence assumption, we conclude that $P(B|B_{T_0}) \geq (1 - \varepsilon)^{1/2}$. Hence $P(B) \geq (1 - \varepsilon)$.

The proof of the theorem will therefore be complete if we prove the following:

Claim 7.4 *Assume that $H_1 \geq H_0$ and let P be a measure on $(\Omega(\alpha, H_1), \Sigma(\alpha, H_1))$ that is consistent with $(F_i)_{i \in A}$. Then, for P -almost all ω in B ,*

$$\exists f(\omega, I) = 1 .$$

(Recall that $I = \arg \max_{i \in A} \mu_i$.)

Proof: Given $\omega \in B$ and $\xi > 0$, we wish to show that, unless ω is in a certain P -null event (to be specified later), there exists a $T = T(\omega, \xi)$ such that for all $t \geq T$,

$$f(\omega, I, t) \geq 1 - \xi .$$

It is sufficient to find a $T = T(\omega, \xi)$ such that for some $i \in I$, for all $t \geq T$, and for all $j \notin I$,

$$\frac{K(\omega, j, t)}{K(\omega, i, t)} = \frac{f(\omega, j, t)}{f(\omega, i, t)} \leq \frac{\xi}{n(1 - \xi)} .$$

We remind the reader that for all $t \geq T_0$ and all $a \in A$ we have

$$|\overline{X}(\omega, a, t) - \mu_a| \leq \delta .$$

Also, since $H_{T_0}(\omega) > R > \mu_1$, for all $t \geq T_0$ we have

$$H_t(\omega) > \mu_1 - \delta = \mu_{r+1} + 2\delta .$$

That is, the aspiration level will be adjusted towards the average performance of one of the expected-utility maximizing acts, and will be bounded away from the expected utility and from the average performance value of sub-optimal acts.

We will need a uniform bound on $H_t(\omega)$. To this end, note that for all $a \in A$ and $t \leq T_0$, $\overline{X}(\omega, a, t) < R$, by definition of the set B_{T_0} . For $t \geq T_0$, the same inequality holds since $\overline{X}(\omega, a, t) < \mu_a + \delta < \overline{b}_a \leq R$. Since $H_{t+1}(\omega)$ is a convex combination of $H_t(\omega)$ and $\overline{X}(\omega, t) = \max_{a \in A} \overline{X}(\omega, a, t) < R$, we conclude that for all $t \geq 1$, $H_{t+1}(\omega) \leq \max\{H_t(\omega), R\}$. By induction, it follows that for all $t \geq 1$, $H_t(\omega) \leq H_1$.

Let $O(\omega) \subset A$ be the set of acts that are chosen infinitely often at ω . That is,

$$O(\omega) = \{a \in A \mid K(\omega, a, t) \longrightarrow \infty \text{ as } t \rightarrow \infty\} .$$

We would first like to establish the fact that some expected utility maximizing acts are indeed chosen infinitely often. Formally,

Claim 7.5 $O(\omega) \cap I \neq \emptyset$.

Proof: Let $\tilde{T} \geq T_0$ be such that for all $t \geq \tilde{T}$, $a_t(\omega) \in O(\omega)$. Assume the contrary, i.e., that $O(\omega) \cap I = \emptyset$. (In particular, $a_t(\omega) \notin I$ for all $t \geq \tilde{T}$.) For all $t \geq \tilde{T} \geq T_0$ we also know that

$$\bar{X}(\omega, j, t) < H_t(\omega) - \delta$$

for all $j \notin I$. Hence, for $j \notin I$,

$$U(\omega, j, t) = K(\omega, j, t)V(\omega, j, t) = K(\omega, j, t)[\bar{X}(\omega, j, t) - H_t(\omega)] < -\delta K(\omega, j, t) .$$

This implies that $U(\omega, j, t) \rightarrow -\infty$ as $K(\omega, j, t) \rightarrow \infty$. Thus, for all $j \in O(\omega) \setminus I$,

$$U(\omega, j, t) \xrightarrow[t \rightarrow \infty]{} -\infty .$$

On the other hand, consider some $i \in I \subset (O(\omega))^c$. Let L satisfy $L > K(\omega, i, t)$ for all $t \geq 1$. Then

$$U(\omega, i, t) = K(\omega, i, t)V(\omega, i, t) = K(\omega, i, t)[\bar{X}(\omega, i, t) - H_t(\omega)] > L(\underline{b} - H_1) .$$

It is therefore impossible that only members of I^c would be U -maximizers from some \tilde{T} on. \square

We now assume that for all $a \in O(\omega)$, $\bar{X}(\omega, a, t) \xrightarrow[t \rightarrow \infty]{} \mu_a$. By the strong law of large numbers, this is the case for all $\omega \in B$ apart from a P -null set. Choose $\zeta > 0$ such that

$$\zeta < \frac{\xi\delta}{6n(1-\xi)}$$

and let $T_1 \geq T_0$ be such that for all $t \geq T_1$ and all $i \in O(\omega) \cap I$,

$$|\bar{X}(\omega, i, t) - \mu_1| < \zeta .$$

For all $t \geq T_1$ we also conclude that

$$|\bar{X}(\omega, t) - \mu_1| < \zeta$$

(where, as above, $\bar{X}(\omega, t) = \max_{a \in A} \bar{X}(\omega, a, t)$). It follows that the aspiration level, $H_{t+1}(\omega)$, which is adjusted to be some average of its previous value $H_t(\omega)$ and $\bar{X}(\omega, t)$, will also converge to μ_1 . To be precise, there is $T_2 \geq T_1$ such that for all $t \geq T_2$,

$$|H_t(\omega) - \mu_1| < 2\zeta .$$

We wish to show that there exists $T(\omega, \xi)$ such that for all $t \geq T(\omega, \xi)$, all $i \in O(\omega) \cap I$ and all $j \notin I$ the following holds:

$$\frac{K(\omega, j, t)}{K(\omega, i, t)} \leq \frac{\xi}{n(1-\xi)} .$$

It will be helpful to start with:

Claim 7.6 *For all $t \geq T_2$, all $i \in O(\omega) \cap I$ and all $j \notin I$, if $a_i(\omega) = j$, then*

$$K(\omega, j, t) < \frac{\xi}{2n(1-\xi)} K(\omega, i, t) .$$

Proof: Let there be given t, i and j as above. Observe that

$$U(\omega, i, t) = K(\omega, i, t)V(\omega, i, t) = K(\omega, i, t)[\bar{X}(\omega, i, t) - H_t(\omega)] \geq -3K(\omega, i, t)\zeta$$

while

$$U(\omega, j, t) = K(\omega, j, t)V(\omega, j, t) = K(\omega, j, t)[\bar{X}(\omega, j, t) - H_t(\omega)] \leq -K(\omega, j, t)\delta .$$

The fact that $a_t(\omega) = j$ implies that $U(\omega, j, t) \geq U(\omega, i, t)$. Hence

$$\begin{aligned} -K(\omega, j, t)\delta &\geq -3K(\omega, i, t)\zeta \\ K(\omega, j, t) &\leq \frac{3\zeta}{\delta} K(\omega, i, t) . \end{aligned}$$

However, the choice of ζ (as smaller than $\frac{\xi\delta}{6n(1-\xi)}$) implies that

$$\frac{3\zeta}{\delta} < \frac{\xi}{2n(1-\xi)} .$$

We have thus established that

$$K(\omega, j, t) < \frac{\xi}{2n(1-\xi)} K(\omega, i, t)$$

for any t at which j is chosen (i.e., $a_t(\omega) = j$). \square

We proceed as follows: let $T_3 \geq T_2$ be such that for all $t \geq T_3$, $a_t(\omega) \in O(\omega)$. Let $T_4 \geq T_3$ be large enough so that for all $t \geq T_4$, $a \in O(\omega)$ and $c \notin O(\omega)$,

$$K(\omega, c, t) \leq \frac{\xi}{n(1-\xi)} K(\omega, a, t).$$

Finally, let $T_5 > T_4$ be such that for all $a \in O(\omega)$, $K(\omega, a, T_5) > K(\omega, a, T_2)$. We now have

Claim 7.7 *For all $t \geq T_5$, all $i \in O(\omega) \cap I$ and all $j \notin I$,*

$$K(\omega, j, t) \leq \frac{\xi}{n(1-\xi)} K(\omega, i, t).$$

Proof: Let there be given t, i and j as above. If $j \notin O(\omega)$, the choice of T_4 concludes the proof. Assume, then, that $j \in O(\omega)$. Then, by choice of T_5 , j has been chosen since T_2 . That is,

$$T_j^t \equiv \{s | T_2 \leq s < t, a_s(\omega) = j\} \neq \emptyset.$$

Let s be the last time at which j was chosen before time t , i.e., $s = \max T_j^t$. Note that

$$K(\omega, j, t) = K(\omega, j, s+1)$$

and

$$K(\omega, i, t) \geq K(\omega, i, s+1).$$

Hence it suffices to show that

$$K(\omega, j, s+1) \leq \frac{\xi}{n(1-\xi)} K(\omega, i, s+1).$$

By Claim 7.6 we know that

$$K(\omega, j, s) \leq \frac{\xi}{2n(1-\xi)} K(\omega, i, s) .$$

Since $s \geq T_2 \geq T_0$, $K(\omega, j, s) \geq K \geq 1$. This implies $\frac{\xi}{2n(1-\xi)} K(\omega, i, s) \geq 1$. Next, observe that

$$\begin{aligned} K(\omega, j, s+1) &= K(\omega, j, s) + 1 \leq \frac{\xi}{2n(1-\xi)} K(\omega, i, s) + 1 \\ &\leq \frac{\xi}{2n(1-\xi)} K(\omega, i, s) + \frac{\xi}{2n(1-\xi)} K(\omega, i, s) \\ &= \frac{\xi}{n(1-\xi)} K(\omega, i, s) = \frac{\xi}{n(1-\xi)} K(\omega, i, s+1) . \end{aligned}$$

This concludes the proof of Claim 7.7. \square

Thus T_5 may serve as the required $T(\omega, \xi)$. As a matter of fact, our claim regarding T_5 is slightly stronger than that we need to prove regarding $T(\omega, \xi)$. The latter should have the inequality of Claim 7.7 satisfied for *some* $i \in I$, while the former satisfies it for *all* $i \in O(\omega) \cap I$, while Claim 7.5 guarantees that this set indeed contains some $i \in I$.

At any rate, Claim 7.7 completes the proof of Claim 7.4, which, in turn, completes the proof of the theorem. \square

Proof of Theorem 7.2: The general idea of the proof, as well as the proof itself, is quite simple: as long as the aspiration level is close to the average performance of an expected utility maximizing act, the proof mimics that of Theorem 7.1. The problem is that the decision maker may “lock in” on sub-optimal acts, which may be almost-satisficing or even satisficing, and not try the optimal acts frequently enough. However, the fact that the decision maker is ambitious infinitely often (in the sense of setting the aspiration level beyond the maximal average performance) guarantees that this will not be the case. Thus, the fact that N_A is infinite ensures that every act will be chosen infinitely often. On the other hand, the fact that it is sparse implies

that these periods of ambitiousness will not change the limit frequencies obtained in the proof of Theorem 7.1.

In the formal proof it will prove convenient to take the following steps: we will restrict our attention to the event at which all acts that are chosen infinitely often have a limit average performance close to their expectation. On this event we will show that the expected utility maximizers among those acts have a limit choice frequency of 1. Finally, we will show that all acts are chosen infinitely often, whence the result follows.

We adopt some notation from the proof of Theorem 7.1. In particular, assume that for some $r < n$,

$$\mu_1 = \mu_2 = \cdots = \mu_r > \mu_{r+1} \geq \mu_{r+2} \geq \cdots \geq \mu_n$$

and denote

$$I = \arg \max_{i \in A} \mu_i = \{1, 2, \dots, r\} .$$

We will also use

$$O(\omega) = \{a \in A \mid K(\omega, a, t) \xrightarrow[t \rightarrow \infty]{} \infty\}$$

and the new notation

$$I(\omega) = \arg \max\{\mu_i \mid i \in O(\omega)\} .$$

We would like to focus on the event

$$B = \{\omega \in \Omega \mid \forall i \in O(\omega), \bar{X}(\omega, i, t) \xrightarrow[t \rightarrow \infty]{} \mu_i\} .$$

Since A is finite, the strong law of large numbers guarantees that $P(B) = 1$ for any consistent P . Thus it suffices to show that for every $\omega \in B$, $f(\omega, I) = 1$. We do this in two steps: we first show that $f(\omega, I(\omega)) = 1$, and then – that $I(\omega) = I$.

Claim 7.8 *For all $\omega \in B$, $\exists f(\omega, I(\omega)) = 1$.*

Proof: Let there be given $\omega \in B$, and denote $\mu = \mu_i$ for some $i \in I(\omega)$. Given the proof of Claim 7.4 in Theorem 7.1, it suffices to show that for every $\zeta > 0$, $|H_t(\omega) - \mu| < \zeta$ holds for all $t \notin N_0$ where $N_0 \subset \mathbb{N}$ is sparse.

Let $\zeta > 0$ be given, and assume without loss of generality that $\zeta < \delta = \frac{\mu - \mu_i}{3}$ for all $i \notin I(\omega)$ and that $\zeta < h$. Let T_1 be such that for all $t \geq T_1$ and all $i \in O(\omega)$,

$$|\bar{X}(\omega, i, t) - \mu_i| < \zeta/2 .$$

Let $T_2 \geq T_1$ be such that for all $t \geq T_2$, $i \in O(\omega)$ and $j \notin O(\omega)$, $\bar{X}(\omega, i, t) > \bar{X}(\omega, j, t)$. Thus, for $t \geq T_2$, if $t \notin N_A$, $H_t(\omega)$ is adjusted towards $\bar{X}(\omega, t)$, which equals $\bar{X}(\omega, i, t)$ for some $i \in O(\omega)$, where the latter is close to μ . Since for $t \in N_A$ $H_t(\omega)$ is set to $\bar{X}(\omega, t) + h$, there exists $T_3 \geq T_2$ such that for all $t \geq T_3$,

$$|H_t(\omega) - \mu| < 2h .$$

We now wish to choose a number k , such that any sequence of k periods following T_3 , at which $H_t(\omega)$ is adjusted “realistically”, that is, as an average of $H_{t-1}(\omega)$ and $\bar{X}(\omega, t)$, will guarantee that it ends up ζ -close to μ .

Let $k > \log_\alpha(\zeta/4h)$. Define

$$N_A \oplus k = \left\{ t \in \mathbb{N} \left| \begin{array}{l} t = t_1 + t_2 \text{ where} \\ t_1 \in N_A \text{ and } 0 \leq t_2 \leq k \end{array} \right. \right\} .$$

Note that for $t \geq T_3$, if $t \notin N_A \oplus k$, i.e., if t is at least k periods after the most recent “ambitious” update, we have

$$|H_t(\omega) - \mu| < \zeta .$$

Setting $N_0 = (N_A \oplus k) \cup \{1, \dots, T_2\}$ (and noting that it is sparse) completes the proof. \square

Claim 7.9 For all $\omega \in B$, $I(\omega) = I$.

Proof: It suffices to show that $O(\omega) = A$ for all $\omega \in B$. Assume, to the contrary, that for some $\omega \in B$, $j \in A$ and $L \geq 1$, $K(\omega, j, t) \leq L$ for all $t \geq 1$. Let $i \in O(\omega)$. For any $t \in N_A$,

$$U(\omega, i, t) = K(\omega, i, t)V(\omega, i, t) = K(\omega, i, t)[\bar{X}(\omega, i, t) - H_t(\omega)] < -hK(\omega, i, t).$$

Let $T_4 \geq T_3$ be such that for all $t \geq T_4$, $a_t(\omega) \neq j$. Recall that for all $t \geq T_4$, $H_t(\omega) < \mu + 2h$. Consider $t \in N_A$ such that $t \geq T_4$. Then

$$U(\omega, j, t) = K(\omega, j, t)V(\omega, j, t) = K(\omega, j, t)[\bar{X}(\omega, j, t) - H_t(\omega)] > -LC$$

where $C = \mu + 2h - \bar{X}(\omega, j, T_3)$. That is, $U(\omega, j, t)$ is bounded from below. Since for a large enough $t \in N_A$, $U(\omega, i, t)$ is arbitrarily small for all $i \in O(\omega)$, we obtain a contradiction to U -maximization. Thus we conclude that $O(\omega) = A$. This concludes the proof of the claim and of the theorem. \square

19 Learning the Similarity Function

19.1 Examples

Sarah wonders which of two new movies, a and b , she should see tonight. Let us consider her choice given two memories. In the first, M_1 , there is only one case: Jim saw movie a and liked it. In the second, M_2 , there are several cases involving other movies (but not a or b). It turns out that there are many movies that, in M_2 , were seen by both Sarah and Jim, but none that they both liked.

Given memory M_1 , Sarah is likely to decide to see movie a . After all, all that she knows is that Jim liked it, and this is more than can be said of movie b . Given memory M_2 , in which neither a nor b appear, Sarah is likely to be indifferent between the two movies. But given the union of M_1 and M_2 , Sarah may conclude that she has no chance of liking movie a , precisely because Jim did like it, and opt for movie b .

This is a direct violation of the combination axiom of Chapter 3.³ It is therefore a phenomenon that cannot be captured by W -maximization. The reason is that the combination axiom implicitly assumes that the algorithm by which one learns from cases, and the support that a case c provides to choice a do not change from one memory to another. But in the example above one learns how to learn from experience. Specifically, memory M_2 suggests that Sarah and Jim have very different tastes, and the fact that he liked a movie (M_1) should detract from its desirability in Sarah's eyes, rather than enhance it.

To consider another example, assume that Mary has to choose a car. Cars are either green or blue, and they are of make 1 or 2. Mary had a good experience with a green car of make 1 (G1), and a bad experience with a blue car of make 2 (B2). Now she has to choose between a green car of make 2 (G2) and a blue one of make 1 (B1). Both are similar to some degree to both

³This example is based on an example suggested to us by Sujoy Mukerji.

past cases. Specifically, each new choice shares the color attribute with one case and the make attribute with the other. Hence the choice between G2 and B1 entails an implicit ranking of the attributes in the similarity judgment. Indeed, if Mary is three years old, she is likely to base her similarity judgment mostly on color. But if Mary is thirty years old, she is likely to put more weight on the make of the car than on its color in judging similarity. Namely, she would consider G2 similar to B2, and B1 – similar to G1, and prefer B1 to G2.

The reason that the adult Mary puts more weight on the make attribute than on the color attribute is that past experience shows that the make of the car is a better predictor of quality than the car's color. In other words, past cases are used to define the similarity function, which, in turn, is used to learn from these cases about the future. For concreteness, assume that Mary has six cases in her memory. In the first four she has used red (R) and white (W) cars both of make 3 and of make 4. The last two are cases G1 and B2 discussed above. We would like to consider her preference between G2 and B1 given three different memories. In the table below, each column corresponds to a problem-act pair, and each row describes a possible history, specifying the utility of the outcome that was experienced in each case. Observe that Mary cannot choose among rows in this table. Rather, given a possible memory, represented by a row, she is asked to choose between G2 and B1.

possible memories	Problem-act pairs					
	R3	R4	W3	W4	G1	B2
x	1	-1	1	-1	1	-1
y	1	1	-1	-1	-1	1
$x + y$	2	0	0	-2	0	0

In memory x , past cases with other makes (3 and 4) and other colors (R and W) can be summarized by the observations that cars of make 3 are good, and that cars of make 4 are bad, regardless of color. This indicates

that make is a more important attribute than is color. Further, this example corresponds to the history we started with: G1 was a good experience, while B2 was a bad one. Thus it makes sense that history x would give rise to preferring B1 to G2. Memory y , by contrast, differs in two ways. First, the first four cases are consistent with the simple theory that all red cars are good, and all white cars are not, regardless of the make. Second, the experiences in G1 and in B2 are the opposite of their counterparts in the story above: now G1 ended up in a bad experience while B2 resulted in success. Thus memory y supports the child's similarity function, and this, in turn, favors B1 to G2, since B1 is of the same color as the car involved in the good experience B2.

Consider now a hypothetical memory generated by the sum of the utility vectors, $x + y$. This summation corresponds to a memory from which one cannot infer much about the relative importance of the color and make attributes. Moreover, the history of G1 and B2 makes them completely equivalent. Hence, even if one knew what similarity function to use, by symmetry one would have to be indifferent between G2 and B1.

Thus, under both x and y , a reasonable decision maker such as Mary would prefer B1 to G2. But given the memory $x + y$, she is indifferent between them. Assuming that the payoff numbers are measured on a utility scale, this is a contradiction to U' -maximization. Specifically, the formula

$$U'(a) = U'_{p,M}(a) = \sum_{(q,b,r) \in M} s((p,a), (q,b))u(r) . \quad (\bullet)$$

implies that if a is preferred to b given a set of cases $\{(q_i, b_i, r_i)\}_i$, and if a is preferred to b given the cases $\{(q_i, b_i, r'_i)\}_i$, then the same preference should be observed given the cases $\{(q_i, b_i, r''_i)\}_i$ when $u(r''_i) = u(r_i) + u(r'_i)$. This additivity condition trivially holds when the similarity function is fixed, but it need not hold when the similarity function is learned from experience.

Generally, we distinguish between two levels of inductive reasoning in the context of CBDT. “First order” induction is the process by which similar past cases are implicitly generalized to bear upon future cases, and, in particular, to affect the decision in a new problem. CBDT, as presented above, attempts to model this process, if only in a rudimentary way. “Second order” induction refers to the process by which the decision maker learns *how to conduct first-order induction*, namely, how to judge similarity of problems. The current version of CBDT does not model this process. Moreover, it implicitly assumes that it does not take place at all.

19.2 Counter-example to U -maximization

It may be insightful to explicitly see how second-order induction may also violate U -maximization. Consider the following example: this afternoon, John wants to order tickets to a concert, and he can do it either by phone (P) or by accessing a Web site (W). He would like to minimize the time and hassle of the process. John’s experience is as follows. He has once ordered tickets to a ballet by phone (at a different hall) in the afternoon, and the procedure was short and efficient. He has also ordered tickets to a concert (at the same hall) by the Web, and this experience was also painless, but then John made the transaction at night. John is aware of the fact that different halls vary in terms of the phone and internet services they offer. He is also aware of the fact that the time of day may greatly affect his waiting time. Thus, the decision problem boils down to the following question: should John assume that ordering tickets for the concert hall in the afternoon is more similar to doing so at night, or to ordering tickets for a ballet in the afternoon? That is, which feature of the problem is more important: the hall where the show is performed or the time of day?

In the past, John has also ordered tickets, both by phone and through the Web, for movie theaters. He has plenty of experience with movie theaters 1 and 2, where he ordered tickets in the morning and in the evening, with

varying degrees of success. John does not view his experience with movie theaters as relevant to the concert hall. But this experience might change the way he evaluates similarities. For concreteness, let us compare two scenarios, both represented by the matrix below. In this matrix, columns stand for decision problems encountered in the past, whereas rows – for acts chosen. The problems are assumed distinct. “M1” stands for ordering tickets to movie theater 1 in the morning; “E2” – for movie theater 2 in the evening, and so forth. “AB” stands for ordering tickets to the ballet in the afternoon, and “NC” – to the concert hall at night. The numerical entries in the matrix represent the payoffs obtained when an act was chosen in a given problem, that is, the values of the function u , and blank entries mean that the act corresponding to the row was not chosen in the problem corresponding to the column. For the purposes of U -maximization, we may think of the blank entries as zero payoffs.

In the first scenario, John’s experience with phone orders is represented by the vector x , where 1 stands for short wait, -1 – for long wait, and his experience with Web orders is represented by the vector y . In the second scenario, the history of phone orders is given by the vector z , and that of Web orders – by w . The vector d does not represent an act in our story and will only be used to facilitate the algebraic description. Observe that it would be incoherent to compare, say, the vectors x and z , since both have non-blank values in certain columns, whereas only one act can be chosen in a given past problem. By contrast, vectors x and y can be the histories of two different acts in a given memory, and so can vectors z and w .

Consider the first scenario, where John has to choose between using the phone, with history x , and using the Web, with history y . John’s experience with movie theaters clearly indicates that the morning is a better time to place an order than is the evening. This particular fact is not of great help to John, who has to order tickets in the afternoon. Moreover, John is now ordering tickets to a concert, rather than to a movie, and it may well be the

case that movie fans are on a different schedule than are concert audience. Still, John might infer that the time of day is a more important attribute than is the particular establishment. John's experience with halls include a success with a phone call made in the afternoon for ballet tickets, and a success with a Web order made at night for concert tickets. The former might be taken to suggest that the afternoon is a good time to call, favoring P to W. The latter tentatively offers the generalization that the concert hall Web system works fine, favoring W to P. Which of the two recommendations is more convincing? Given the experience with movie theaters, which supports generalizations based on the time of day rather than on the establishment, John is likely to give more weight to the afternoon order at the ballet and prefer a phone call, namely, choose an act with history x over one with history y .

By contrast, in the second scenario John faces a choice between an act with history z and an act with history w . Both histories suggest that movie theater 1 has an efficient ordering service, whereas movie theater 2 does not. Again, this information per se is not very useful when ordering tickets to a concert, which is in a hall having nothing to do with either movie theater. But this piece of information also indicates that establishments do differ from each other, while the time of day is of little or no importance. This conclusion would imply that John should rely on his night experience at the concert hall more than on his afternoon experience at the ballet, and opt for a Web order, namely, choose an act with history w over one with history z .

act profiles	Problems									
	M1	M2	E1	E2	M1	M2	E1	E2	AB	NC
x (P)	1	1					-1	-1	1	
y (W)			-1	-1	1	1				1
z (P)	1		1			-1		-1	1	
w (W)		-1		-1	1		1			1
d		-1	1			-1	1			

However, this preference pattern is inconsistent with U -maximization for any fixed similarity function s . Indeed, for every s , since $z - x = w - y = d$, we have

$$U(z) - U(x) = U(w) - U(y) = U(d)$$

implying

$$U(x) - U(y) = U(z) - U(w) .$$

That is, x is preferred to y if and only if z is preferred to w , in contradiction to the preference pattern we motivated above.

Similar examples may be constructed, in which the violation of U -maximization stems from learning that certain *values* of an attribute are similar, rather than that the attribute itself is of importance. For example, instead of learning that the establishment is an important factor, John may simply learn that the concert hall is similar to the ballet theater.

In the previous subsection we generated a counter-example to W -maximization by a direct violation of the combination axiom. For concreteness we have also constructed counter-examples to U' - and to U -maximization. Indeed, all these functions are additively separable across cases, whereas second-order induction renders the weight of a set of cases a non-additive set function. Several cases, taken in conjunction, may implicitly suggest a rule. Hence the effect of all of them together may differ from the sum of effects of each one, considered separately. Differently put, the “marginal contribution” of a case to preferences depends not only on the case itself, but also on the other cases it is lumped together with.⁴

⁴A possible generalization of CBDT functions that may account for this phenomenon involves the use of non-additive measures, where aggregation of utility is done by the Choquet integral. (See Schmeidler (1989).)

19.3 Learning and expertise

The distinction between the two levels of induction may also help classify types of learning and of expertise. A case-based decision maker learns in two ways: first, by introducing more cases into memory; second, by refining the similarity function based on past experiences. Knowing more is but one aspect of learning. Making better use of known cases is another. Correspondingly, the notion of expertise combines knowledge of cases and the ability to focus on the relevant analogies.

Knowledge of cases is relatively objective. By contrast, knowledge of the similarity function is inherently subjective.⁵ Correspondingly, it is easier to compare people's knowledge of cases than it is to compare their knowledge of similarity. It is natural to define a relation "knows more than" when cases are considered: a person (or a system) knows more cases than another if the memory of the former is a superset of that of the latter. The concept of "knowing more" is more evasive when applied to the similarity function. In hindsight, a similarity function that leads to better decision making can be viewed as better or more precise. But it is difficult to provide an a-priori formal definition of "having a better similarity function". It is relatively easy to find that an expert has a vast database to draw upon. It is harder to know whether she makes the "right" analogies when using it.

In Section 13 we discussed the two roles that rules may play in a case-based knowledge representation system. These two roles correspond to the two types of knowledge and to the two levels of induction. The first role, of summarizing many cases, conveys knowledge of cases. First-order induction suffices to formulate such rules: given a similarity function, one generates a rule by lumping similar cases together.⁶ The second role, of drawing attention

⁵True, even "objective" cases and alleged "facts" may be described, construed, and interpreted in different ways. But they are still closer to objectivity than are similarity judgments.

⁶Notice, however, that this is first-order *explicit* induction, that is, a process that generates a general rule, as opposed to the *implicit* induction performed by CBDT.

to similarity among cases, conveys knowledge of similarity. One needs to engage in second-order induction to formulate such rules: the similarity has to be learned in order to observe the regularity that the rule should express.

These distinctions may have implications for the implementation of computerized systems. A case-based expert system has to represent both types of knowledge. As a result of programming necessity, cases will typically be represented by a database, whereas similarity judgments might be implicit in the software. The discussion above suggests that such a distinction is also conceptually desirable. For instance, one may wish to use one expert's knowledge of cases with another expert's similarity judgments. It might therefore be useful to represent cases in a way that is independent of the similarity function.

Finally, we remind the reader that in Section 13 we argued that case-based systems enjoy a theoretical advantage when compared to rule-based systems: case-based knowledge representation incorporates modifications in a smooth way. This advantage exists also in the presence of second-order induction. Second-order induction may induce updates of similarity values, and this may lead to different decisions based on the same set of cases. But this process does not pose any theoretical difficulties such as those entailed by explicit induction.

20 Two Views of Induction: CBDT and Simplicism

20.1 Wittgenstein and Hume

Case-based decision making involves implicit first- and second-order induction. To what extent can it serve as a model of inductive reasoning or inductive inference? How does it compare with explicit induction, namely, with the formation of rules or general theories? This section is devoted to a few preliminary thoughts on these questions.

How do people use past cases to extrapolate the outcomes of present circumstances? Wittgenstein (1922, 6.363) suggested that people tend to engage in explicit induction:

“The procedure of induction consists in accepting as true the *simplest* law that can be reconciled with our experiences.”⁷

The notion of simplicity is rather vague and subjective. (See, for instance, Sober (1975) and Gärdenfors (1990).) Gilboa (1994) suggests employing Kolmogorov’s complexity measure for the definition of simplicity. In this model, a theory is a program that can generate a prediction for every instance it accepts as input. It is then argued, as a descriptive theory of inductive inference, that people tend to choose the shortest program that conforms to their observations. This theory is referred to as “simplicism”. Its prediction is well defined *given* a programming language. Yet the choice of language is no less subjective than the notion of simplicity. Indeed, simplicism merely translates the choice of a complexity measure to the choice of the appropriate programming language.

By contrast, our starting point is Hume’s claim that “from similar causes we expect similar effects”. That is, Hume offers an alternative descriptive

⁷Observe that this statement has a descriptive flavor, as opposed to the normative flavor of Occam’s razor.

theory of inductive reasoning, namely, that the process of implicit induction is based on the notion of similarity. Thus, Hume may be described as referring to implicit induction, based on the (vague and subjective) notion of similarity, whereas Wittgenstein may be viewed as defining explicit induction, based on the (vague and subjective) notion of simplicity.

It is tempting to contrast the two views of induction. One way to do so is by choosing a formal representation of each theory. As an exercise, we here take Wittgenstein's view of explicit induction as modeled by simplicism, and Hume's view of implicit induction as modeled by CBDT, and study their implications in a few simple examples. Two caveats are in order: first, any formal model of an informal claim is bound to commit to a certain interpretation thereof; thus the particular models we discuss may not do justice to the original views. Second, both similarity and language are inherently subjective notions. Hence, the formal models we discuss here by no means resolve all ambiguity. Even with the use of these models, much freedom is left in the way the two views are brought to bear on a particular example. Yet, we hope that the analysis of a few simple examples might indicate some of the advantages of both views as theories of human thinking.

20.2 Examples

Consider a simple learning problem. Every item has two observable attributes, color and shape. Each attribute might take one of two values. Color might be red (R) or blue (B) whereas shape might be a square (S) or a circle (C). The task is to learn a concept, say, of "nice" items, Σ , that is fully determined by the attributes. Formally, Σ is a subset of $\{RS, RC, BS, BC\}$. We are given positive and/or negative examples to learn from. An example states that one of RS , RC , BS , or BC is in Σ (+) or that it is not (-). We are asked to extrapolate whether a new item is in Σ , based on its observable attributes (again, one of the pairs RS , RC , BS , or BC). The set of examples we have observed, namely, our memory, may be summarized by a matrix, in

which the symbol “+” stands for “such items are in Σ ”, the symbol “-” stands for “such items are not in Σ ”, and a blank space is read as “such items have not yet been observed”. Finally, a “?” would indicate the next item we are asked about. For instance, the matrix

1	<i>S</i>	<i>C</i>
<i>R</i>	+	
<i>B</i>		?

describes a database in which a red square (*RS*) is known to be nice, and we are asked about a blue circle (*BC*).

Let us proceed with this example. Knowing only that a red square is nice, would a blue circle be nice as well? Not having observed any negative examples, the simplest theory in any reasonable language is likely to be “All items are nice”, predicting “+” for *BC*. Correspondingly, if we assume that all items bear some resemblance to each other, a case-based extrapolator will also come up with this prediction.

Next consider the matrices

2	<i>S</i>	<i>C</i>
<i>R</i>	+	-
<i>B</i>		?

3	<i>S</i>	<i>C</i>
<i>R</i>	-	+
<i>B</i>		?

In matrix 2, the simplest theory would probably be “an item is nice if and only if it is a square”, predicting that *BC* is not in Σ . The same prediction would be generated by CBDT: since *BC* is more similar to *RC* than it is to *RS*, the former (non-nice) example would outweigh the latter (nice) one. Similarly, simplicism and CBDT will concur in their prediction for matrix 3, predicting that *BC* is nice. We also get similar predictions for the following two matrices:

4	<i>S</i>	<i>C</i>
<i>R</i>	+	
<i>B</i>	-	?

5	<i>S</i>	<i>C</i>
<i>R</i>	-	
<i>B</i>	+	?

However, the two methods of extrapolation might also be differentiated. Consider the matrices

6	<i>S</i>	<i>C</i>
<i>R</i>	+	?
<i>B</i>		-

7	<i>S</i>	<i>C</i>
<i>R</i>	+	
<i>B</i>	?	-

The observations in both matrices are identical. The simplest rule that accounts for them is not uniquely defined: the theory “an item is nice if and only if it is red”, as well as the theory “an item is nice if and only if it is a square” are both consistent with evidence, and both would be minimizers of Kolmogorov’s complexity measure in a language that has *R* and *S* as primitives. (As opposed to, say, their conjunction.) Each of these simplest theories would yield a positive prediction in one matrix and a negative one in the other. By contrast, assuming that color and shape are symmetric, the CBDT classification would leave us undecided between a positive and a negative answer in both matrices.

In matrices 6 and 7 the prediction provided by CBDT appears more satisfactory than that of simplicism. In both matrices the evidence for and against a positive prediction is precisely balanced. Simplicism comes up with two equally simple but very different answers, whereas CBDT reflects the balance between them. Since CBDT uses quantitative similarity judgments, and produces quantitative evaluation functions, it deals with indifference more graciously than do simplest theories or rules. In a way that parallels our discussion in Section 13, we find that CBDT behaves more smoothly at the transition between different rules.

It is natural to suggest that simplicism be extended to a random choice among all simplest theories, or an expected prediction of a theory chosen randomly. Indeed, in matrices 6 and 7 above, if we were to quantify the predictions and take an average of the predictions of the two simplest theories, we will also be indifferent between a positive and a negative prediction. But once we allow weighted aggregation of theories, we would probably not want to restrict it to cases of indifference. For instance, suppose that a hundred theories, each of which has Kolmogorov complexity of 1,001 (say, bits), agree on a prediction, but disagree with the unique simplest theory, whose complexity is 1,000. It would be natural to extend the aggregated prediction method to this case as well, allowing many almost-simplest theories to outweigh a single simplest one. But then we are led down the slippery path leading to a Bayesian prior over all theories, which is a far cry from simplicism.

An example of more pronounced disagreement between CBDT and simplicism is provided by the following matrix.

8	<i>S</i>	<i>C</i>
<i>R</i>	+	+
<i>B</i>	−	?

The simplest theory that accounts for the data is “an item is nice if and only if it is red”, predicting that a blue circle is not nice. What will be the CBDT prediction? Had *RS* not been observed, the situation would have been symmetric to matrices 6 and 7, leaving a CBDT predictor indifferent between a positive and a negative answer. However, the similarity between *RS* and *BC* is positive, as argued above. (If it were not, CBDT would not yield a positive prediction in matrix 1.) Hence the additional observation tilts the balance in favor of a positive answer.

Observe, however, that the derivation of the CBDT prediction in matrix 8 relies on additive separability. That is, we only allow CBDT to employ first-order induction. But if we extend CBDT to incorporate second-order

induction, the three observations in matrix 8 do indicate that color is a more important attribute than is shape. If this is reflected in the weight that the similarity function puts on the two attributes, BC will be more similar to BS than to RC . If this difference is large enough, the fact that a blue square is not nice may outweigh the two nice red examples.

Second-order induction can also be defined in the context of simplicism. A simplicistic predictor chooses simplest theories given a language, but she can also learn the language, in an analogous way to learning the similarity function in CBDT. Consider Mary of Section 19 again. We have argued that, as compared with children, most adults learn to put less emphasis on color as a defining feature of a car's quality. In both models, this type of learning falls under the category of second-order induction: in CBDT, as we have seen, it would be modeled by changing the similarity function so that the weight attached to color in similarity judgments is reduced. In simplicism, it would be captured by including in the language other predicates, and perhaps dispensing with "color" altogether. Observe that in this respect, too, the quantitative nature of CBDT may provide more flexibility than the qualitative choice of language in simplicism.

To sum, CBDT appears to be more flexible than does simplicism. Implicit induction seems to avoid some of the difficulties posed by explicit induction. Yet, modeling only first order induction is unsatisfactory. It remains a challenge to find a formal model that will capture Hume's intuition and allow quantitative aggregation of cases, without excluding second-order induction and refinements of the similarity function.

Bibliography

- [A] Akaike, H. (1954), “An Approximation to the Density Function”, *Annals of the Institute of Statistical Mathematics*, **6**: 127-132.
- [All] Allais, M. (1953), “Le Comportement de L’Homme Rationel devant le Risque: critique des Postulates et Axioms de l’Ecole Americaine”, *Econometrica*, **21**: 503-546.
- [Al] Alt, F., (1936), “On the Measurability of Utility”, in J. S. Chipman, L. Hurwicz, M. K. Richter, H. F. Sonnenschein (eds.): *Preferences, Utility, and Demand, A. Minnesota Symposium*, pp. 424-431, ch. 20. New York: Harcourt Brace Jovanovich, Inc. [Translation of “Uber die Messbarkeit des Nutzens”, *Zeitschrift fuer Nationaloekonomie* **7**: 161-169.]
- [An] Anscombe, F. J., and R. J. Aumann (1963), “A Definition of Subjective Probability”, *The Annals of Mathematics and Statistics*, **34**: 199-205.
- [Ar] Aragonés, E. (1997), “Negativity Effect and the Emergence of Ideologies”, *Journal of Theoretical Politics*, **9**: 189-210.
- [Arr1] Arrow, K. (1986), “Rationality of Self and Others in an Economic System”, *Journal of Business*, **59**: S385-S399.
- [Arr2] Arrow, K. and G. Debreu (1954), “Existence of an Equilibrium for a Competitive Economy”, *Econometrica*, **22**: 265-290.

- [AL] Ashkenazi, G. and E. Lehrer (2000), "Well-Being Indices", mimeo.
- [Au] Aumann, R. J. (1962), "Utility Theory without the Completeness Axiom", *Econometrica* **30**: 445-462.
- [BG] Beja, A. and I. Gilboa (1992), "Numerical Representations of Imperfectly Ordered Preferences (A Unified Geometric Exposition)", *Journal of Mathematical Psychology*, **36**: 426-449.
- [B] Bewley, T. (1986), "Knightian Decision Theory: Part I", Discussion Paper 807, Cowles Foundation.
- [BM] Bush, R. R., and F. Mosteller (1955), *Stochastic Models for Learning*. New York: Wiley.
- [CW] Camerer, C. and M. Weber (1992), "Recent Developments in Modeling Preferences: Uncertainty and Ambiguity", *Journal of Risk and Uncertainty*, **5**: 325-370.
- [Ca] Carnap, R. (1923), "Uber die Aufgabe der Physik und die Anwendung des Grundsatzes der Einfachheit", *Kant-Studien*, **28**: 90-107.
- [CH] Cover, T. and P. Hart (1967), "Nearest Neighbor Pattern Classification", *IEEE Transactions on Information Theory* **13**: 21-27.
- [C] Cross, J. G. (1983), *A Theory of Adaptive Economic Behavior*. New York: Cambridge University Press.
- [deF] de Finetti, B. (1937), "La Prevision: Ses Lois Logiques, Ses Sources Subjectives", *Annales de l'Institute Henri Poincare*, **7**: 1-68.
- [DeG] DeGroot, M. H. (1975), *Probability and Statistics*. Reading, MA: Addison-Wesley.

- [DGL] Devroye, L., L. Györfi, and G. Lugosi (1996), *A Probabilistic Theory of Pattern Recognition*, New York: Springer-Verlag.
- [E] Ellsberg, D. (1961), "Risk, Ambiguity and the Savage Axioms", *Quarterly Journal of Economics*, **75**: 643-669.
- [Et] Etzioni, A. (1986), "Rationality is Anti-Entropic", *Journal of Economic Psychology*, **7**: 17-36.
- [Fi1] Fishburn, P. C. (1970), "Intransitive Indifference in Preference Theory: A Survey", *Operations Research*, **18**: 207-228.
- [Fi2] Fishburn, P. C. (1985), *Interval Orders and Interval Graphs*. New York: Wiley and Sons.
- [FH1] Fix, E. and J. Hodges (1951), "Discriminatory Analysis. Nonparametric Discrimination: Consistency Properties". Technical Report 4, Project Number 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX.
- [FH2] Fix, E. and J. Hodges (1952), "Discriminatory Analysis: Small Sample Performance". Technical Report 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX.
- [FFG] Flakenhainer, B., K. D. Forbus, and D. Gentner (1989), "The Structure-Mapping Engine: Algorithmic Example", *Artificial Intelligence*, **41**: 1-63.
- [FY] Foster, D. and H. P. Young (1990), "Stochastic Evolutionary Game Dynamics", *Theoretical Population Biology*, **38**: 219-232.
- [G] Gärdenfors, P. (1990), "Induction, Conceptual Spaces and AI", *Philosophy of Science*, **57**: 78-95.
- [GV] Gibbard, A. and H. Varian (1978), "Economic Models", *Journal of Philosophy*, **75**: 664-677.

- [GH1] Gick, M. L. and K. J. Holyoak (1980), “Analogical Problem Solving”, *Cognitive Psychology*, **12**: 306-355.
- [GH2] Gick, M. L. and K. J. Holyoak (1983), “Schema Induction and Analogical Transfer”, *Cognitive Psychology*, **15**: 1-38.
- [G87] Gilboa, I. (1987), “Expected Utility with Purely Subjective Non-Additive Probabilities,” *Journal of Mathematical Economics* **16**: 65-88.
- [Gi] Gilboa, I. (1994), “Philosophical Applications of Kolmogorov’s Complexity Measure”, in *Logic and Philosophy of Science in Uppsala*, D. Prawitz and D. Westerstahl (eds.), Synthese Library, Vol. 236, Kluwer Academic Press, pp. 205-230.
- [GL] Gilboa, I. and R. Lapsen (1995), “Aggregation of Semi-Orders: Intransitive Indifference Makes a Difference”, *Economic Theory*, **5**: 109-126.
- [GP] Gilboa, I. and A. Pazgal (1996), “History-Dependent Brand Switching: Theory and Evidence”, Northwestern University Discussion Paper.
- [GP2] Gilboa, I. and A. Pazgal (2000), “Cumulative Discrete Choice”, *Marketing Letters*, forthcoming.
- [GS1] Gilboa, I. and D. Schmeidler (1989), “Maxmin Expected Utility with a Non-Unique Prior”, *Journal of Mathematical Economics*, **18**: 141-153.
- [GS2] Gilboa, I. and D. Schmeidler (1995), “Case-Based Decision Theory”, *The Quarterly Journal of Economics*, **110**: 605-639.
- [GS4] Gilboa, I. and D. Schmeidler (1996a), “Case-Based Optimization”, *Games and Economic Behavior*, **15**: 1-26.

- [GS5] Gilboa, I. and D. Schmeidler (1996b), “A Cognitive Model of Well-being”, *Social Choice and Welfare*, forthcoming.
- [GS6] Gilboa, I. and D. Schmeidler (1997a), “Act Similarity in Case-Based Decision Theory”, *Economic Theory*, **9**: 47-61.
- [GS7] Gilboa, I., and D. Schmeidler (1997b), “Cumulative Utility Consumer Theory”, *International Economic Review*, **38**: 737-761.
- [GS3] Gilboa, I. and D. Schmeidler (2000a), “Case-Based Knowledge and Induction”, *IEEE Transactions on Systems, Man, and Cybernetics*.
- [GS8] Gilboa, I. and D. Schmeidler (2000b), “Cognitive Foundations of Inductive Inference and Probability: An Axiomatic Approach”, mimeo.
- [GS9] Gilboa, I. and D. Schmeidler (2000c), “They’re All Nearest Neighbors”, mimeo.
- [GSW] Gilboa, I., D. Schmeidler, and P. Wakker (1999), “Utility in Case-Based Decision Theory”, mimeo.
- [Git] Gittins, J. C. (1979), “Bandit Processes and Dynamic Allocation Indices”, *Journal of the Royal Statistical Society, B*, **41**: 148-164.
- [H] Halmos, P. R. (1950), *Measure Theory*. Princeton, NJ, Van Nostrand.
- [Ha] Hanson, N. R. (1958), *Patterns of Discovery*. Cambridge, England: Cambridge University Press.
- [HC] Harless, D. and C. Camerer (1994), “The Utility of Generalized Expected Utility Theories”, *Econometrica*, **62**: 1251-1289.
- [Har1] Harsanyi J. C. (1953), “Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking”, *Journal of Political Economy*, **61**: 434-435.

- [Har2] Harsanyi, J. C. (1955), "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility", *Journal of Political Economy*, **63**: 309-321.
- [HP] Herrnstein, R. J., and D. Prelec (1991), "Melioration: A Theory of Distributed Choice", *Journal of Economic Perspectives*, **5**: 137-156.
- [Ho] Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- [Hu] Hume, D. (1748), *An Enquiry Concerning Human Understanding*. Oxford: Clarendon Press.
- [KT] Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, **47**: 263-291.
- [KMR] Kandori, M., G. J. Mailath, and R. Rob (1993), "Learning, Mutation and Long-Run Equilibria in Games", *Econometrica*, **61**: 29-56.
- [KVS] Karni, E., D. Schmeidler, and K. Vind (1983), "On state dependent preferences and subjective probabilities," *Econometrica*, **51**: 1021-1031.
- [KM] Karni, E. and P. Mongin, (2000), "On the Determination of Subjective Probability by Choices", *Management Science*, **46**: 233-248.
- [KS] Karni E. and D. Schmeidler (1991), "Utility theory with uncertainty", in *Handbook of Mathematical Economics 4*, W. Hildenbrand and H. Sonnenschein (eds.), North Holland, pp. 1763-1831.
- [K] Keynes, J. M. (1921), *A Treatise on Probability*. London: MacMillan and Co.
- [KG] Kirkpatrick, S., C. D. Gellatt, et al. (1982), "Optimization by Simulated Annealing", *IBM Thomas J. Watson Research Center*. Yorktown Heights, NY.

- [Kn] Knight, F. H. (1921), *Risk, Uncertainty, and Profit*. Boston, New York: Houghton Mifflin.
- [Ko] Kolodner, J. L., Ed. (1988), *Proceedings of the First Case-Based Reasoning Workshop*. Los Altos, CA: Morgan Kaufmann Publishers.
- [KR] Kolodner, J. L. and C. K. Riesbeck (1986), *Experience, Memory, and Reasoning*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- [L] Levi, I. (1980), *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
- [Le] Lewin, S. (1986), "Economics and Psychology: Lessons For Our Own Day, From the Early Twentieth Century", *Journal of Economic Literature*, **34**: 1293-1323.
- [Lu] Luce, R. D. (1956), "Semiorders and a Theory of Utility Discrimination", *Econometrica*, **24**: 178-191.
- [M] Machina, M. (1987), "Choice Under Uncertainty: Problems Solved and Unsolved", *Economic Perspectives*, **1**: 121-154.
- [MS] March, J. G. and H. A. Simon (1958), *Organizations*. New York: Wiley.
- [Ma] Matsui, A. (2000), "Expected Utility and Case-Based Reasoning", *Mathematical Social Sciences*, **39**: 1-12.
- [MD] McDermott, D. and J. Doyle (1980), "Non-Monotonic Logic I", *Artificial Intelligence* **25**: 41-72.
- [Mo] Moser, P. K., Ed. (1986), *Empirical Knowledge*. Rowman and Littlefield Publishers.
- [My] Myerson, R. B. (1995), "Axiomatic Derivation of Scoring Rules Without the Ordering Assumption", *Social Choice and Welfare*, **12**: 59-74.

- [Na] Nash, J. F. (1951), “Non-Cooperative Games”, *Annals of Mathematics*, **54**: 286-295.
- [Pa] Parzen, E. (1962), “On the Estimation of a Probability Density Function and the Mode”, *Annals of Mathematical Statistics*, **33**: 1065-1076.
- [P] Pazgal, A. (1997), “Satisficing Leads to Cooperation in Mutual Interests Games”, *International Journal of Game Theory*, **26**: 439-453.
- [Po] Popper, K.R. (1934), *Logik der Forschung*; English edition (1958), *The Logic of Scientific Discovery*. London: Hutchinson and Co. Reprinted (1961), New York: Science Editions.
- [Q1] Quine, W. V. (1953), “Two Dogmas of Empiricism”, in *From a Logical Point of View*. Cambridge, MA: Harvard University Press.
- [Q2] Quine, W. V. (1969a), “Epistemology Naturalized”, in *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- [Q3] Quine, W. V. (1969b), “Natural Kinds”, in *Essays in Honor of Carl G. Hempel*, Rescher N. (ed.), Dordrecht, Holland: D. Reidel Publishing Company.
- [R] Ramsey, F. P. (1931), “Truth and Probability”, in *The Foundation of Mathematics and Other Logical Essays*. New York, Harcourt, Brace and Co.
- [Ra] Rawls, J. (1971), *A Theory of Justice*. Cambridge, MA: Belknap.
- [Re] Reiter, R. (1980), “A Logic for Default Reasoning”, *Artificial Intelligence*, **13**: 81-132.
- [RS] Riesbeck, C. K. and R. C. Schank (1989), *Inside Case-Based Reasoning*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

- [Ros] Rosenblatt, M. (1956), "Remarks on Some Nonparametric Estimates of a Density Function", *Annal of Mathematical Statistics*, **27**: 832-837.
- [Ro] Royall, R. (1966), *A Class of Nonparametric Estimators of a Smooth Regression Function*. Ph.D. Thesis, Stanford University, Stanford, CA.
- [S] Savage, L. J. (1954), *The Foundations of Statistics*. New York: John Wiley and Sons.
- [Sc] Schank, R. C. (1986), *Explanation Patterns: Understanding Mechanically and Creatively*. Hillsdale, NJ, Lawrence Erlbaum Associates.
- [Schm] Schmeidler, D. (1989), "Subjective Probability and Expected Utility without Additivity", *Econometrica*, **57**: 571-587.
- [Se] Selten, R. (1978), "The Chain-Store Paradox", *Theory and Decision*, **9**: 127-158.
- [Sh] Shapley, L. S. (1953), "A Value for n -Person Games", in *Contributions to the Theory of Games II*, H. W. Kuhn and A. W. Tucker (eds.), Princeton: Princeton University Press, pp. 307-317.
- [SOF] Shepperd, J. A., A. J. Ouellette, and J. K. Fernandez (1996), "Abandoning Unrealistic Optimism: Performance, Estimates and the Temporal Proximity of Self-Relevant Feedback", *Journal of Personality and Social Psychology*, **70**: 844-855.
- [Si1] Simon, H. A. (1957), *Models of Man*. New York: John Wiley and Sons.
- [Si2] Simon, H. A. (1986), "Rationality in Psychology and Economics", *Journal of Business*, **59**: S209-S224.

- [Sk] Skinner, B. F. (1938), *The Behavior of Organisms*. New York: Appleton-Century-Crofts.
- [So] Sober, E. (1975), *Simplicity*. Oxford: Clarendon Press.
- [Su] Suppe, F. (1974), *The Structure of Scientific Theories* (edited with a critical introduction by F. Suppe), Urbana, Chicago, London: University of Illinois Press.
- [T1] Tversky, A. (1977), "Features of Similarity", *Psychological Review*, **84**: 327-352.
- [T2] Tversky, A., and D. Kahneman (1981), "The Framing of Decisions and the Psychology of Choice", *Science*, **211**: 453-458.
- [vNM] von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.
- [Wk] Wakker, P. P. (1989), *Additive Representations of Preferences*. Dordrecht, Boston, London: Kluwer Academic Publishers.
- [W1] Watson, J.B. (1913), "Psychology as the Behaviorist Views It", *Psychological Review*, **20**: 158-177.
- [W2] Watson, J.B. (1930), *Behaviorism* (Revised Edition). New York: Norton.
- [Wi] Wittgenstein, L. (1922), *Tractatus Logico Philosophicus*. London, Routledge and Kegan Paul.
- [Y1] Young, P. H. (1975), "Social Choice Scoring Function", *SIAM Journal of Applied Mathematics*, **28**: 824-838.
- [Y2] Young, P. H. (1993), "The Evolution of Conventions", *Econometrica*, **61**: 57-84.

Index

- Akaike, 69
- Allais, 18
- Alt, 81
- ambitiousness, 159
- annealing, 163
- Anscombe, 28
- Aragones, 139, 143
- Arrow, 27, 32
- Ashkenazi, 77
- aspiration level, 47, 54, 137, 143, 157
- Aumann, 28, 53
- average utility, 55
- axiomatization, 20, 71, 128
- Bayes, 51, 124, 126
 - rule, 100
- Bayesian statistics, 29
- behaviorism, 24, 33
- Bewley, 53
- Bush, 138
- Camerer, 37
- Carnap, 21
- case-based reasoning, 41, 116
- change seeking, 54
- Choquet, 192
- classification, 111
- cognitive, 24
 - data, 128
 - plausibility, 33, 52, 101
 - specification, 25, 34, 63, 64
- combination axiom, 73, 75, 82, 186
- complementarity, 151
- conceptual framework, 13, 31, 101
- Cover, 111
- Cross, 41
- de Finetti, 37, 54
 - exchangeability, 73
- Debreu, 32
- decision tree, 126
- DeGroot, 82
- descriptive, 13, 16, 22, 27
- Devroye, 111
- Doyle, 107
- dynamic programming, 127
- empirical knowledge, 106, 107
- Eshel, 60
- Etzioni, 27

- expected utility theory, 23, 24, 31,
 37, 39, 50, 63, 66, 99, 163
- Falkenhainer, 43
 falsifiable, 22, 31
 Fernandez, 160
 first welfare theorem, 32
 Fishburn, 82
 Fix, 111
 Forbus, 43
 Foster, 163
 framework, 13
 framing effect, 25
- game theory, 172
 Gardenfors, 195
 Gellatt, 163
 genetic algorithms, 163
 Gentner, 43
 Gibbard, 30
 Gick, 43
 Gilboa-Schechtman, 144
 Gittins, 161
 Gyorfi, 111
- habit formation, 54
 Halmos, 174
 Hanson, 30, 108
 Harless, 37
 Harsanyi, 17, 20
 Hart, 111
 Herrnstein, 137
- Hodges, 111
 Holland, 163
 Holyoak, 43
 Hume, 40, 43, 107, 110, 195
 hypothesis testing, 29, 83
 hypothetical
 - cases, 54, 100
 - memory, 74
 - reasoning, 101
- induction, 105, 157, 195
 - explicit, 107
 - first-order, 193
 - implicit, 110, 113
 - second-order, 84, 194
- insufficient reason, 104
- Kahneman, 15, 25
 Kandori, 164
 Karni, 37, 81
 kernel-based estimate, 69
 Keynes, 41
 Kirkpatrick, 163
 Knight, 52
 Kolmogorov complexity, 195, 199
 Kolodner, 41
- language, 195
 Laplace, 104
 learning, 77, 138, 157, 163, 186,
 193
 Lehman, 84

- Lehrer, 77
Levi, 107
Lewin, 30
logical positivism, 21, 30, 31
Luce, 82
Lugosi, 111
Machina, 37
Mailath, 164
March, 49
Markov chain, 127
Matsui, 100
McDermott, 107
melioration theory, 137
memory, 45, 135, 143, 172
Mongin, 81
Morgenstern, 15, 18, 37
Moser, 108
Mosteller, 138
Mukerji, 186
multi-armed bandit, 161
Myerson, 74
Nash, 15
nearest neighbor, 84, 111
normative, 13, 16, 22, 27
objective, 27
observable, 20, 103
Occam's razor, 195
Ouellette, 160
Parzen, 69
Pazgal, 138, 143, 172
philosophy of science, 29
planning, 117
Popper, 22, 31
potential, 147
prediction, 67
Prelec, 137
Quine, 43, 106
Ramsey, 37
rationality, 25, 27, 35, 127
 bounded, 35, 105, 157
Rawls, 17, 20
realism, 159
Received View, 30
Reiter, 107
Riesbeck, 41, 115, 116
risk, 37
Rob, 164
Rosenblatt, 69
Royall, 111
rule, 113, 115, 193
rule based, 53, 106
satisficing, 47, 99, 136
Savage, 15, 18, 23–25, 33, 37, 38,
 54, 102
Schank, 41, 115, 116
Selten, 41
Shapley, 18
Shepperd, 160

- similarity, 43, 54, 104, 111, 115, 186
 - averaged, 55
 - of acts, 57, 152, 170
 - of cases, 59
- Simon, 27, 49, 136
- simplicism, 195
- Simpson's paradox, 82
- Skinner, 24
- Sober, 195
- states of the world, 40, 100
- subjective, 27
- substitution, 151
- Suppe, 21
- sure-thing principle, 33
- theory, 13, 16, 20, 24, 31
 - laden observations, 30, 108
- Tversky, 16, 25, 43
- U''-maximization, 59, 64
- U'-maximization, 57, 64, 188
- U-maximization, 169, 189
- utility, 13, 22, 24, 44, 113, 124, 158
 - cardinal, 141
 - cumulative, 54, 135
 - neo-classical, 149
- utility theory, 31
- V-maximization, 55, 64, 82, 169
- Varian, 30
- Vecchi, 163
- verificationism, 30
- Vind, 81
- von Neumann, 18, 37
- von-Neumann, 15
- W-maximization, 61, 71, 186
- Wakker, 53, 72, 82
- Walliser, 82
- Watson, 24
- Weber, 37
- Wittgenstein, 195
- Young, 74, 163, 164