REGRET MINIMIZATION AND RELATED DECISION RULES

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

YinYee Leung

August 2015

REGRET MINIMIZATION AND RELATED DECISION RULES

YinYee Leung, Ph.D.
Cornell University 2015

Our starting point is a setting where a decision maker's uncertainty is represented by a set of probability measures, rather than a single measure. Measure-by-measure (a.k.a. prior-by-prior) updating of such a set of measures upon acquiring new information is well-known to suffer from problems. To deal with these problems, we propose using *weighted sets of probabilities*: a representation where each measure is associated with a *weight*, which denotes its significance. We describe a natural approach to updating in such a situation. We then show how this representation can be used in decision-making, by modifying a standard approach to decision making—minimizing expected regret—to obtain *minimax weighted expected regret* (MWER). We provide an axiomatization that characterizes preferences induced by MWER both in the static and dynamic case.

This same concept of weighted probability distributions can also be applied to the widely-studied maxmin expected utility decision rule. Chateauneuf and Faro [2009] axiomatize a weighted version of maxmin expected utility over acts with nonnegative utilities, where weights are represented by a confidence function. We argue that their representation is only one of many possible, and we axiomatize a more natural form of maxmin weighted expected utility. We also provide stronger uniqueness results.

Next, we apply regret-minimization to dynamic decision problems. The menu-dependent nature of regret-minimization creates subtleties when it is ap-

plied to dynamic decision problems. If forgone opportunities are included, we can characterize when a form of dynamic consistency is guaranteed.

Finally, we look at a “dual” of minimax regret, called maximin safety. Much as regret is a form of distance to the best possible outcome, safety is a form of distance to be worst possible outcome. The idea behind maximin safety is that one would want to maximize the distance between one’s choice and the worst possible outcome. This decision rule might be appropriate in cases where it is important not to be “the last person”, for instance when a group of hikers is being chased by a bear. We examine its behavioral motivations and provide an axiomatization for the decision rule.

BIOGRAPHICAL SKETCH

Samantha Leung received her B.Sc. in Computer Science (with Co-op work experience) from the University of British Columbia, Canada, in 2009. In the same year, she entered the Ph.D. program at the Department of Computer Science at Cornell University.

ACKNOWLEDGEMENTS

I thank my advisor, Joseph Y. Halpern, for his guidance and patience, both with this body of work and with other aspects in my graduate school experience. In addition, I thank my thesis committee members Lawrence Edward Blume and Hod Lipson for their time and useful comments. I also thank Edward Lui for useful discussions and continued encouragement.

I thank Leandro Chaves Rêgo and the Theory and Decision reviewers for their insights and useful comments pertaining to Chapter 3, and Brad Gulko for his initiation of and contribution to Chapter 5.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

1.1 Introduction

Decision-making is an important aspect of intelligent behavior. Decision theorists are interested in two ways of studying decisions: *normative* and *descriptive*. Descriptive decision theory attempts to explain behavior that are observed in real humans. This thesis lies within the scope of normative decision theory – we are interested in how intelligent agents "ought to" make decisions. In decision theory, this is usually accomplished through the axiomatization of decision rules – giving a set of axioms and show that they characterize the behavior of the decision rule. Intuitively, if a decision rule is characterized by a set of reasonable or desirable axioms, then the decision rule itself should be considered reasonable or desirable.

Savage [1951] and Anscombe and Aumann [1963] showed that a decision maker maximizing expected utility with respect to a probability measure over the possible states of the world is characterized by a set of arguably desirable principles. However, as Allais [1953b] and Ellsberg [1961] point out using compelling examples, sometimes intuitive choices are incompatible with maximizing expected utility.

In the Ellsberg paradox [1961], there is an urn with 90 marbles, 30 of which are red, and 60 of which are either yellow or blue. The decision maker (henceforth DM) does not know how many blue or yellow balls there are. The DM is given a choice between two gambles: gamble A, under which the DM wins

$100 if the next marble drawn is red; and gamble B, under which the DM wins $100 if the next marble drawn is blue. Most humans choose gamble A. The DM is further asked to compare two different gambles: gamble C, under which the DM wins $100 if the next marble drawn is red or yellow; and gamble D, under which the DM wins $100 if the next marble drawn is blue or yellow. Most humans choose gamble D. Assuming that the DM is an expected-utility maximizer, there is no subjective belief on the proportions of yellow and blue balls that result in the observed preferences (that is, gamble A is preferred over gamble B and gamble D is preferred over gamble C).

One reason for this incompatibility is that there is often *ambiguity* in the problems we face; we often lack sufficient information to capture all uncertainty using a single probability measure over the possible states.

To this end, there is a rich literature offering alternative means of making decisions (see, e.g., [Al-Najjar and Weinstein 2009] for a survey). For example, we might choose to represent uncertainty using a set of possible states of the world, but using no probabilistic information at all to represent how likely each state is. With this type of representation, two well-studied rules for decision-making are *maximin utility* and *minimax regret*. Maximin says that you should choose the option that maximizes the worst-case payoff, while minimax regret says that you should choose the option that minimizes the *regret* you'll feel at the end, where, roughly speaking, regret is the difference between the payoff you achieved, and the payoff that you could have achieved had you known what the true state of the world was. Both maximin and minimax regret can be extended naturally to deal with other representations of uncertainty. For example, with a set of probability measures over the possible states, minimax regret becomes minimax expected regret (MER) [Hayashi 2011; Stoye 2011a]. Other works that use a set of probability measures include, for example, [Campos and Moral 1995; Cousa, Moral, and Walley 1999; Gilboa and Schmeidler 1993; Levi 1985; Walley 1991].

We now briefly explain our minimax weighted expected regret decision rule (MWER), the topic of Chapter 2, by first discussing MER. MER is a probabilistic variant of the minimax regret decision rule proposed by Niehans [1948] and Savage [1951].[1] Most likely, at some point, we've second-guessed ourselves and thought "had I known this, I would have done that instead". That is, in hindsight, we regret not choosing the act that turned out to be optimal for the realized state, called the *ex post* optimal act. The *regret* of an act a in a state s is the difference (in utility) between the ex post optimal act in s and a. Of course, typically one does not know the true state at the time of decision. Therefore the regret of an act is the worst-case regret, taken over all states. The *minimax regret* rule orders acts by their regret.

The definition of regret can be used if there is no probability on states. If a DM's uncertainty is represented by a single probability measure, then we can compute the *expected regret* of an act a: just multiply the regret of an act a at a state s by the probability of s, and then sum. It is well known that the order on acts induced by minimizing expected regret is identical to that induced by maximizing expected utility (see [Hayashi 2008a] for a proof). If an DM's uncertainty is represented by a set $\mathcal{P}$ of probabilities, then we can compute the expected regret of an act a with respect to each probability measure $\Pr \in \mathcal{P}$, and then take the worst-case expected regret. The MER (Minimax Expected Regret)

[1] Note that our definition of regret minimization, while standard, differs from that used by Loomes and Sugden [**?**], where probabilities are given, and where the DM not only feels regret but also "rejoice" if the chosen alternative is better than the unchosen ones.

rule orders acts according to their worst-case expected regret, preferring the act that minimizes the worst-case regret. If the set of measures is the set of *all* probability measures on states, then it is not hard to show that MER induces the same order on acts as (probability-free) minimax regret. Thus, MER generalizes both minimax regret (if $\mathcal{P}$ consists of all measures) and expected utility maximization (if $\mathcal{P}$ consists of a single measure).

MWER further generalizes MER. If we start with a *weighted* set of measures, then we can compute the weighted expected regret for each one (just multiply the expected regret with respect to $\Pr$ by the weight of $\Pr$) and compare acts by their worst-case weighted expected regret.

We also consider the updating of beliefs. We propose a method of updating a set of weighted probability measures, called *likelihood updating*. In Chapter 2, we provide an axiomatization for MWER, as well as for MWER with likelihood updating.

We have thus far neglected the other popular decision rule, maximin expected utility. Just as minimax regret generalizes to minimax expected regret (MER) when uncertainty is represented by a set of probability distributions, the maximin decision rule also has a counterpart when uncertainty is represented by a set of probability distributions. This "multiple priors" version of maximin is called maxmin expected utility (MEU or MMEU). Gilboa and Schmeider [**?**] axiomatize the MMEU decision rule in their seminal paper.

In Chapter 3, we change the focus to MMEU. In particular, we apply the concept of weighted probability measures, which we used for regret-minimization in Chapter 2 to obtain MWER, to MMEU, resulting in a decision rule which we call maxmin weighted expected utility (MWEU). It turns out that Chateauneuf and Faro [2009] had also axiomatized a weighted version of maxmin expected utility over acts with nonnegative utilities. We argue that their representation is only one of many possible, and we axiomatize a more natural form of maxmin weighted expected utility, dropping one of Chateauneuf and Faro's axioms. We also provide stronger uniqueness results for the set of weights in the representation theorem.

In Chapter 4, we return to the study of regret-minimization, this time with special attention on the dynamic aspects of regret-minimization. We look at both measure-by-measure updating, as well as likelihood updating, and characterize when a form of dynamic consistency is guaranteed when each of these updating rules are used. Additional examples of minimizing regret in dynamic decision problems, such as the secretary problem, can be found in the appendix.

Finally, in Chapter 5 we look at a notion very much related to regret, which we call *safety*. The *maximin safety* decision rule seeks to maintain a large margin from the worst outcome, in much the same way minimax regret seeks to minimize distance from the best. In a sense, safety is the "dual" of regret because it measures the distance to the worst outcome, while regret measures the distance to the best outcome. The maximization of safety therefore have similar qualities to minimizing regret. It turns out that Stoye [2011a] had already considered an equivalence notion to maximizing safety. However, he only briefly mentioned its possibility, which he calls the maximization of "joy", within a proof in the paper.

1.2 BIBLIOGRAPHIC NOTES

Chapter 2 is joint-work with Joseph Y. Halpern, who contributed significantly to the exposition, and provided the proof to the main theorem. A short version of Chapter 2 appeared in the proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI 2012) as "Weighted Sets of Probabilities and Minimax Weighted Expected Regret: New Approaches for Representing Uncertainty and Making Decisions".

Chapter 3 is joint-work with Joseph Y. Halpern. The material in Chapter 3 has been accepted to the journal Theory and Decisions under the title "Maxmin Weighted Expected Utility: A Simpler Characterization".

Chapter 4 is also joint-work with Joseph Y. Halpern, who provided several of the proofs. A 10-page proceedings version of Chapter 4 appeared in the proceedings of European Conferences on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2015) as "Minimizing Regret in Dynamic Decision Problems".

Chapter 5 is joint work initiated by coauthor Brad Gulko, who contributed much of the literature research, motivation, and examples. I was responsible for the proofs of the results. Chapter 5 appeared in the proceedings of the European Conferences on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2013) as "Maximin Safety: When Failing to Lose Is Preferable to Trying to Win."

The work in this thesis was partly supported by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by AFOSR grants FA9550-08-1-0438 and FA9550-09-

1-0266, by the Multidisciplinary University Research Initiative (MURI) program administered by the AFOSR under grant FA9550-12-1-0040, and by ARO grants W911NF-09-1-0281 and W911INF-14-1-0017.

The entire body of work was completed under the guidance of Joseph Y. Halpern, who contributed many ideas, proofs, and exposition.

CHAPTER 2

MINIMAX WEIGHTED EXPECTED REGRET

2.1 Overview

We consider a setting where a DM's uncertainty is represented by a set of probability measures, rather than a single measure. Measure-by-measure updating of such a set of measures upon acquiring new information is well-known to suffer from problems. To deal with these problems, we propose using *weighted sets of probabilities*: a representation where each measure is associated with a *weight*, which denotes its significance. We describe a natural approach to updating in such a situation and a natural approach to determining the weights. We then show how this representation can be used in decision-making, by modifying a standard approach to decision making—minimizing expected regret—to obtain *minimax weighted expected regret* (MWER). We provide an axiomatization that characterizes preferences induced by MWER both in the static and dynamic case.

2.2 Introduction

From deciding between crispy fries and a bland salad, to forming an investment portfolio, to military planning, our decisions can significantly impact our lives and those of others. These problems can often be abstracted as decision problems with uncertainty. For decisions based on the outcome of the toss of a fair coin, the uncertainty can be well characterized by probability. However, what

	1 defect	10 defect
deliver	10,000	-10,000
cancel	0	0
check	5,001	-4,999

Table 2.1: Payoffs for the quality-control problem. Acts are in the leftmost column. The remaining two columns describe the outcome for the two sets of states that matter.

is the probability of you gaining weight if you eat fries at every meal? What if you have salads instead? Even experts would not agree on a single probability.

Representing uncertainty by a single probability measure and making decisions by maximizing expected utility leads to further problems. Consider the following stylized quality control-problem, which serves as a running example in this chapter. A business owner (the DM) is contracted to produce $1,000$ items, and will be rewarded when she delivers the items, but punished if she delivers a batch with too many defective items. She has recently switched a raw material supplier, and does not know whether the supplier is reliable, and provides good quality materials, or unreliable, and provides bad quality materials. For simplicity, assume that using good quality raw materials results in one defective item in the batch of $1,000$, while using bad quality raw materials results in ten defective items.[1]

The owner's choices and their consequences are summarized in Table 2.1. Decision theorists typically model decision problems with states, acts, and outcomes: the world is in one of many possible states, and the DM chooses an *act*, a function mapping states to outcomes. For now, we use the simplest possible state space for this problem: $\{$one defect, ten defects$\}$. These two possible states

[1]It is more natural to assume that the quality of the raw materials affects the distribution of the number of defects, rather than directly affecting the final number; we make the latter assumption here for simplicity.

are sufficient to capture the owner's uncertainty in the quality-control problem. (However, we later consider a more refined state space.)

The owner can choose among three acts: $deliver$: deliver the products; $cancel$: cancel the contract; or $check$: inspect enough items to determine the number of defects, then decide to deliver or cancel the contract. The client will tolerate at most one defect in the lot of $1,000$. Therefore, if the owner chooses $deliver$, and if there is only one defect, then the client is happy, and the owner obtains a utility of $10,000$; if there are ten defects, then the outcome then the client will penalize the owner, resulting in a utility of $-10,000$. If the owner chooses to $cancel$, then the contract is canceled, and the owner gets a utility of 0. Finally, checking the items costs $4,999$ units of utility but is reliable, so if the owner chooses $check$, and if there is one defect, then the products will be delivered after the check, and the client will be happy; the owner nets a utility of $5,001$. On the other hand, if there are ten defects, then after the check, the contract will be canceled, and the owner nets a utility of $-4,999$.

To maximize expected utility, we must assume some probability over states. What measure should be used? There are two hypotheses that the owner entertains: (1) the raw material is of high quality and (2) the raw material is of low quality. Each of these hypotheses places a different probability on states.

If the raw material is of high quality, then with probability 1 there will only be one defect; if the raw material is of low quality, then with probability 1 there will be ten defects. One way to model the owner's uncertainty about the quality of the material is to take each hypothesis to be equally likely. However, not having any idea about which hypothesis holds is very different from believing that all hypotheses are equally likely. It is easy to check that taking each hypothesis to be equally likely makes *check* the act that maximizes utility, but taking the probability that the raw material has low quality $.51$ makes *cancel* the act that maximizes expected utility, and taking the probability that the raw material has high quality to be $.51$ makes *deliver* the act that maximizes expected utility. What makes any of these choices the "right" choice?

It is easy to construct many other examples where a single probability measure does not capture uncertainty, and does not result in what seem to be reasonable decisions, when combined with expected utility maximization. A natural alternative, which has often been considered in the literature, is to represent the DM's uncertainty by a *set* of probability measures. For example, in the quality-control problem, the owner's beliefs could be represented by two probability measures, $\Pr_1$ and $\Pr_{10}$, one for each hypothesis. Thus, $\Pr_1$ assigns uniform probability to all states with exactly one defective items, and $\Pr_{10}$ assigns uniform probability to all states with exactly ten defective items.

But this representation also has problems. Consider the quality-control problem again. Why should the owner be sure that there is exactly either one or ten defective items? Of course, we can replace these two hypotheses by hypotheses that say that the probability of an item being defective is either $.001$ or $.01$, but this doesn't solve the problem. Why should the DM be sure that the probability is either exactly $.001$ or exactly $.01$? Couldn't it also be $.0999$? Representing uncertainty by a set of measures still places a sharp boundary on what measures are considered possible and impossible.

A second problem involves updating beliefs. How should beliefs be updated if they are represented by a set of probability measures? The standard approach for updating a single measure is by conditioning. The natural extension of conditioning to sets of measure is measure-by-measure updating: conditioning each measure on the information (and also removing measures that give the information probability 0).

However, measure-by-measure (or prior-by-prior) updating can produce some rather counter-intuitive outcomes. In the quality-control problem, suppose that the owner knows that the first 100 items that came off the assembly line are good. We denote this piece of information by E. Intuition tells us that there is now more reason to believe that there is only one defective item. The simple two-state state space we used is not sufficient to capture this new information, so we now expand the state space from $\{\text{one defect}, \text{ten defects}\}$ to $\{\text{good}, \text{defective}\}^{1000}$. That is, each item that gets produced is numbered from 1 to $1,000$, and each one can be either defective, or good. We can adapt the two hypotheses in the obvious way to this new state space, and the hypotheses can be conditioned on the new information.

However, $\Pr_1 \mid E$ places uniform probability on all states where the first 100 items are good, and there is exactly one defective item among the last 900 items. Similarly, $\Pr_{10} \mid E$ places uniform probability on all states where the first 100 items are good, and there are exactly ten defective items among the last 900. $\Pr_1 \mid E$ still places probability 1 on there being one defective product, just like $\Pr_1$, and $\Pr_{10} \mid E$ still places probability 1 on there being ten defective products. There is no way to capture the fact that the owner now views the hypothesis $\Pr_{10}$ as less likely, even if the owner was told that the first 990 items are all good!

Of course, both of these problems would be alleviated if we placed a probability on hypotheses, but, as we have already observed, simply maximizing expected utility with respect to this second-order probability distribution has its problems. In this thesis, we propose an intermediate approach: representing uncertainty using *weighted sets of probabilities*. That is, each probability measure is associated with a weight. These weights can be viewed as probabilities; indeed, if the set of probabilities is finite, we can normalize them so that they are effectively probabilities. Moreover, in one important setting, we update them in the same way that we would update probabilities, using likelihood (see below). On the other hand, these weights do not act like probabilities if the set of probabilities is infinite. For example, if we had a countable set of hypotheses, we could assign them all weight 1 (so that, intuitively, they are all viewed as equally likely), but there is no uniform measure on a countable set. To avoid complications about measurability, we think of our representation as a weighted set of probabilities. However, one can equally well think of the representation as a probability on probabilities.

More importantly, when it comes to decision making, we use the weights quite differently from how we would use second-order probabilities on probabilities. Second-order probabilities would let us define a probability on events (by taking expectation) and maximize expected utility, in the usual way. Using the weights, we instead define a novel decision rule, *minimax weighted expected regret (MWER)*, that has some rather nice properties. If all the weights are 1, then MWER is just the standard *minimax expected regret (MER)* rule (described below). If the set of probabilities is a singleton, then MWER agrees with (subjective) expected utility maximization (SEU). More interestingly perhaps, if, through updating, the weights converge to a single hypothesis/probability measure (which happens in one important special case, discussed below), MWER converges to SEU. Thus, the weights give us a smooth, natural way of interpolating between MER and SEU.

In summary, weighted sets of probabilities allow us to represent ambiguity (uncertainty about the correct probability distribution). Real individuals are sensitive to this ambiguity when making decisions, and the MWER decision rule takes this into account. Updating the weighted sets of probabilities using likelihood allows the initial ambiguity to be resolved as more information about the true distribution is obtained.

MWER further generalizes MER. If we start with a *weighted* set of measures, then we can compute the weighted expected regret for each one (just multiply the expected regret with respect to $\Pr$ by the weight of $\Pr$) and compare acts by their worst-case weighted expected regret.

Sarver [2008] also proves a representation theorem that involves putting a multiplicative weight on a regret quantity. However, his representation is fundamentally different from MWER. In his representation, regret is a factor only when comparing two *sets* of acts; the ranking of individual acts is given by expected utility maximization. By way of contrast, we do not compare sets of acts.

It is standard in decision theory to axiomatize a decision rule by means of a representation theorem. For example, Savage [1954] showed that if an DM's preferences $\succeq$ satisfied several axioms, such as completeness and transitivity, then the DM is behaving as if she is maximizing expected utility with respect to some utility function and probabilistic belief.

If uncertainty is represented by a set of probability measures, then we can generalize expected utility maximization to *maxmin expected utility (MMEU)*. MMEU compares acts by their worst-case expected utility, taken over all measures. MMEU has been axiomatized by Gilboa and Schmeidler [1989b]. MER was axiomatized by Hayashi [2008a] and Stoye [2011a]. We provide an axiomatization of MWER. We make use of ideas introduced by Stoye [2011a] in his axiomatization of MER, but the extension seems quite nontrivial.

We also consider a dynamic setting, where beliefs are updated by new information. If observations are generated according to a probability measure that is stable over time, then, as we suggested above, there is a natural way of updating the weights given observations, using ideas of likelihood. The idea is straightforward. After receiving some information E, we update each probability $\Pr \in \mathcal{P}$ to $\Pr \mid E$, and take its weight to be $\alpha_{\Pr} = \Pr(E)/\sup_{\Pr' \in \mathcal{P}} \Pr'(E)$. If more than one $\Pr \in \mathcal{P}$ gets updated to the same $\Pr \mid E$, the $\sup$ of all such weights is used. Thus, the weight of $\Pr$ after observing E is modified by taking into account the likelihood of observing E assuming that $\Pr$ is the true probability. We refer to this method of updating weights as *likelihood updating*.

If observations are generated by a stable measure (e.g., we observe the outcomes of repeated flips of a biased coin) then, as the DM makes more and more observations, the weighted set of probabilities of the DM will, almost surely, look more and more like a single measure. The weight of the measures in $\mathcal{P}$ closest to the measure generating the observations converges to 1, and the weight of all other measures converges to 0. This would not be the case if uncertainty were represented by a set of probability measures and we did measure-by-measure updating, as is standard. As we mentioned above, this means that MWER converges to SEU.

We provide an axiomatization for dynamic MWER with likelihood updating. We remark that a dynamic version of MMEU with measure-by-measure updating has been axiomatized by Jaffray [1994], Pires [2002], and Siniscalchi [2011].

Likelihood updating is somewhat similar in spirit to an updating method implicitly proposed by Epstein and Schneider [2007]. They also represented uncertainty by using (unweighted) sets of probability measures. They choose a threshold α with $0 < \alpha < 1$, update by conditioning, and eliminate all measures whose relative likelihood does not exceed the threshold. This approach also has the property that, over time, all that is left in $\mathcal{P}$ are the measures closest to the measure generating the observations; all other measures are eliminated. However, it has the drawback that it introduces a new, somewhat arbitrary, parameter α.

Chateauneuf and Faro [2009] also consider weighted sets of probabilities (they model the weights using what they call *confidence functions*). They then define and provide a representation of a generalization of MMEU using weighted sets of probabilities that parallels our generalization of MER. Chateauneuf and Faro do not discuss the dynamic situation; specifically, they do not consider how weights should be updated in light of new information. We discuss the relationship of our work that of Chateauneuf and Faro in more detail in Section A.2.

Ideas related to putting weights on a set of probability measures can be dated back to at least Gärdenfors and Sahlin [1982, 1983]; see also [Good 1980] for discussion and further references. Walley [1997] suggested putting a possibility measure [Dubois and Prade 1998; Zadeh 1978] on probability measures; this was also essentially done by Cattaneo [2007], Chateauneuf and Faro [2009], and de Cooman [2005]. All of these authors and others (e.g., Klibanoff et al. [2005]; Maccheroni et al. [2006]; Nau [1992]) proposed approaches to decision making using their representations of uncertainty.

Klibanoff et al. [2005] propose a model of decision making that associates weights with probability measures, but makes decisions based on a "weighted" expected utility function. Maccheroni et al. [2006] study a model of decision making where additive, instead of multiplicative, weights are associated with probability measures. Hayashi [2008a] considers a model of expected-regret-minimization where regrets computed with respect to each state are taken to a positive power before expectations are taken. Others have also proposed and studied approaches of representing uncertainty that are similar to weighted probabilities (see, e.g. [Abdellaoui, Baillon, Placido, and Wakker 2011; Cooman 2005; Moral 1992; Walley 1997]).

The rest of this chapter is organized as follows. Section 2.3 introduces the weighted sets of probabilities representation, and Section 2.4 introduces the MWER decision rule. Axiomatic characterizations of static and dynamic MWER are provided in Sections 2.5 and 2.6, respectively.

2.3 Weighted Sets of Probabilities

A set $\mathcal{P}^+$ of *weighted probability measures* on a set S consists of pairs $(\Pr, \alpha_{\Pr})$, where $\alpha_{\Pr} \in [0,1]$ and $\Pr$ is a probability measure on S.[2] Let $\mathcal{P} = \{\Pr : \exists\alpha((\Pr, \alpha) \in \mathcal{P}^+)\}$. We assume that, for each $\Pr \in \mathcal{P}$, there is exactly one α such that $(\Pr, \alpha) \in \mathcal{P}^+$. We denote this number by $\alpha_{\Pr}$, and view it as the *weight of* $\Pr$. We further assume for convenience that weights have been normalized so that there is at least one measure $\Pr \in \mathcal{P}$ such that $\alpha_{\Pr} = 1$.[3]

[2]In this chapter, for ease of exposition, we take the state space S to be finite, and assume that all sets are measurable. We can easily generalize to arbitrary measure spaces.

[3]While we could take weights to be probabilities, and normalize them so that they sum to 1, if $\mathcal{P}$ is finite, this runs into difficulties if we have an infinite number of measures in $\mathcal{P}$. For example, if we are tossing a coin, and $\mathcal{P}$ includes all probabilities on heads from $1/3$ to $2/3$, using a uniform probability, we would be forced to assign each individual probability measure

As we observed in the introduction, one way of updating weighted sets of probabilities is by using likelihood updating. We use $\mathcal{P}^+ \mid E$ to denote the result of applying likelihood updating to $\mathcal{P}^+$. Define $\overline{\mathcal{P}}^+(E) = \sup\{\alpha_{\Pr} \Pr(E) : \Pr \in \mathcal{P}\}$; if $\overline{\mathcal{P}}^+(E) > 0$, set $\alpha_{\Pr,E} = \sup_{\{\Pr' \in \mathcal{P} : \Pr' \mid E = \Pr \mid E\}} \frac{\alpha_{\Pr'} \Pr'(E)}{\overline{\mathcal{P}}^+(E)}$. Note that given a measure $\Pr \in \mathcal{P}$, there may be several distinct measures $\Pr'$ in $\mathcal{P}$ such that $\Pr' \mid E = \Pr \mid E$. Thus, we take the weight of $\Pr \mid E$ to be the $\sup$ of the possible candidate values of $\alpha_{\Pr,E}$. By dividing by $\overline{\mathcal{P}}^+(E)$, we guarantee that $\alpha_{\Pr,E} \in [0,1]$, and that there is some measure $\Pr$ such that $\alpha_{\Pr,E} = 1$, as long as there is some pair $(\alpha_{\Pr}, \Pr) \in \mathcal{P}$ such that $\alpha_{\Pr} \Pr(E) = \overline{\mathcal{P}}^+(E)$. If $\overline{\mathcal{P}}^+(E) > 0$, we take $\mathcal{P}^+ \mid E$ to be

$$\{(\Pr \mid E, \alpha_{\Pr,E}) : \Pr \in \mathcal{P}\}.$$

If $\overline{\mathcal{P}}^+(E) = 0$, then $\mathcal{P}^+ \mid E$ is undefined.

In computing $\mathcal{P}^+ \mid E$, we update not just the probability measures in $\mathcal{P}$, but also their weights. The new weight combines the old weight with the likelihood. Clearly, if all measures in $\mathcal{P}$ assign the same probability to the event E, then likelihood updating and measure-by-measure updating coincide. This is not surprising, since such an observation E does not give us information about the relative likelihood of measures. We stress that using likelihood updating is appropriate only if the measure generating the observations is assumed to be stable. For example, if observations of heads and tails are generated by coin tosses, and a coin of possibly different bias is tossed in each round, then likelihood updating would not be appropriate.

It is well known that, when conditioning on a single probability measure, the order that information is acquired is irrelevant; the same observation easily

a weight of 0, which would not work well in the definition of MWER.

extends to sets of probability measures. As we now show, it can be further extended to weighted sets of probability measures.

Proposition 2.3.1. *Likelihood updating is consistent in the sense that for all* $E_1, E_2 \subseteq S$, $(\mathcal{P}^+ \mid E_1) \mid E_2 = (\mathcal{P}^+ \mid E_2) \mid E_1 = \mathcal{P}^+ \mid (E_1 \cap E_2)$, *provided that* $\mathcal{P}^+ \mid (E_1 \cap E_2)$ *is defined.*

Proof. By standard results, $(\Pr \mid E_1) \mid E_2 = (\Pr \mid E_2) \mid E_1 = \Pr \mid (E_1 \cap E_2)$. Since the weight of the measure $\Pr \mid E_1$ is proportional to $\alpha_{\Pr} \Pr(E_1)$, the weight of $(\Pr \mid E_1) \mid E_2$ is proportional to $\alpha_{\Pr} \Pr(E_1) \Pr(E_2 \mid E_1) = \alpha_{\Pr} \Pr(E_1 \cap E_2)$. Likewise, the weight of $(\Pr \mid E_2) \mid E_1$ is proportional to $\alpha_{\Pr} \Pr(E_2) \Pr(E_1 \mid E_2) = \alpha_{\Pr} \Pr(E_1 \cap E_2)$. Since, in all these cases, the $\sup$ of the weights is normalized to 1, the weights of corresonding measures in $\mathcal{P}^+ \mid (E_1 \cap E_2)$, $(\mathcal{P}^+ \mid E_1) \mid E_2$ and $(\mathcal{P}^+ \mid E_2) \mid E_1$ must be equal. $\square$

2.4 Formal Definitions and MWER

In this section, we introduce the decision setting that will be used in all the chapters in this thesis.

Given a set S of states and a set X of outcomes, an *act* f (over S and X) is a function mapping S to X. We use $\mathcal{F}$ to denote the set of all acts. For simplicity in this thesis, we take S to be finite. Associated with each outcome $x \in X$ is a utility: $u(x)$ is the utility of outcome x. We call a tuple (S, X, u) a *(non-probabilistic) decision problem*. To define regret, we need to assume that we are also given a set $M \subseteq \mathcal{F}$ of acts, called the *menu*. The reason for the menu is that, as is well known (and we will demonstrate by example shortly), regret can depend on the menu. Moreover, we assume that every menu M has utilities bounded from above. That is, we assume that for all menus M, $\sup_{g\in M} u(g(s))$ is finite. This ensures that the regret of each act is well defined.[4]

For a menu M and act $f \in M$, the regret of f with respect to M and decision problem (S, X, u) in state s is

$$reg_M(f, s) = \left(\sup_{g\in M} u(g(s))\right) - u(f(s)).$$

That is, the regret of f in state s (relative to menu M) is the difference between $u(f(s))$ and the highest utility possible in state s among all the acts in M. The regret of f with respect to M and decision problem (S, X, u), denoted $reg_M^{(S,X,u)}(f)$, is the worst-case regret over all states:

$$reg_M^{(S,X,u)}(f) = \max_{s\in S} reg_M(f, s).$$

We typically omit superscript (S, X, u) in $reg_M^{(S,X,u)}(f)$ if it is clear from context.

The minimax regret decision rule chooses an act that minimizes $\max_{s\in S} reg_M(f, s)$. In other words, the minimax regret choice function is

$$C_M^{reg,u}(M') = \operatorname*{argmin}_{f\in M'} \max_{s\in S} reg_M(f, s).$$

The choice function returns the set of all acts in M' that minimize regret with respect to M. Note that we allow the menu M', the set of acts over which we are minimizing regret, to be different from the menu M of acts with respect to which regret is computed. For example, if the DM considers forgone opportunities, they would be included in M, although not in M'.

[4]Stoye [2011b] assumes that, for each menu M, there is a finite set A_M of acts such that M consists of all the convex combinations of the acts in A_M. We clearly allow a larger set of menus than Stoye. We return to the issue of what menus to consider after we discuss the representation theorem in Section A.2, and again when we discuss choice functions in Section 2.5.

If there is a probability measure $\Pr$ over the σ-algebra Σ on the set S of states, then we can consider the *probabilistic decision problem* $(S, \Sigma, X, u, \Pr)$. In that context, we omit mentioning Σ, and instead consider the *probabilistic decision problem* $(S, X, u, \Pr)$.

The *expected regret* of f with respect to M is

$$reg_M^{\Pr}(f) = \sum_{s \in S} \Pr(s) reg_M(f, s).$$

If there is a set $\mathcal{P}$ of probability measures over the σ-algebra Σ on the set S of states, states, then we consider the $\mathcal{P}$-decision problem $\mathcal{D} = (S, \Sigma, X, u, \mathcal{P})$. The maximum expected regret of $f \in M$ with respect to M and $\mathcal{D}$ is

$$reg_M^{\mathcal{P}}(f) = \sup_{\Pr \in \mathcal{P}} \left(\sum_{s \in S} \Pr(s) reg_M(f, s) \right).$$

The minimax expected regret (MER) decision rule minimizes $reg_M^{\mathcal{P}}(f)$.

Of course, we can define the choice functions $C_M^{reg,\Pr}$ and $C_M^{reg,\mathcal{P}}$ using $reg_M^{\Pr}$ and $reg_M^{\mathcal{P}}$, by analogy with C_M^{reg}.

For example, the MER choice function is defined as

$$C_{\mathcal{P}}^{reg,u}(M) = \operatorname*{argmin}_{f \in M} \sup_{\Pr \in \mathcal{P}^+} reg_{M,\mathcal{P}}(f).$$

In addition to choice functions, another standard way to represent preferences is to use *preference relations*. For example, the maximin expected utility decision rule can be expressed as

$$f \succeq^{maximin} g \Leftrightarrow E[u(f)] \geq E[u(g)].$$

Similarly, MMEU can be expressed as the following preference relation:

$$f \succeq^{MMEU,\mathcal{P}} g \Leftrightarrow \inf_{\Pr \in \mathcal{P}} E_{\Pr}[u(f)] \geq \inf_{\Pr \in \mathcal{P}} E_{\Pr}[u(g)].$$

To represent MER, whose ranking of acts depend on the menu with respect to which the regrets are computed, we need more than just a single preference relation. We need a family of preference relations, one for each menu. For example, given a menu M, the comparison of acts $f, g \in M$ can be expressed as

$$f \succeq_M^{reg,\mathcal{P}} g \Leftrightarrow reg_M^{\mathcal{P}}(f) \leq reg_M^{\mathcal{P}}(g).$$

We now define MWER formally. For simplicity in this chapter, we take S to be finite, and take the σ-algebra Σ to be the powerset of S. Moreover, we assume that every menu M has utilities bounded from above. That is, we assume that for all menus M, $\sup_{g \in M} u(g(s))$ is finite. This ensures that the regret of each act is well defined.[5]

If beliefs are modeled by weighted probabilities $\mathcal{P}^+$, then we consider the $\mathcal{P}^+$-decision problem $(S, X, u, \mathcal{P}^+)$. The maximum weighted expected regret of $f \in M$ with respect to M and $(S, X, u, \mathcal{P}^+)$ is

$$reg_{M,\mathcal{P}^+}(f) = \sup_{\Pr \in \mathcal{P}} \left(\alpha_{\Pr} \sum_{s \in S} \Pr(s) reg_M(f, s) \right).$$

If $\mathcal{P}^+$ is empty, then $reg_M^{\mathcal{P}^+}$ is identically zero.

The MWER choice function is thus defined as

$$C_{\mathcal{P}^+}^{reg,u}(M) = \operatorname*{argmin}_{f \in M} \sup_{(\Pr,\alpha) \in \mathcal{P}^+} reg_{M,\mathcal{P}^+}(f).$$

As a family of preference relations, the MWER decision rule is defined for all $f, g \in X^S$ as

$$f \succeq_{M,\mathcal{P}^+}^{S,X,u} g \text{ iff } reg_{M,\mathcal{P}^+}^{(S,X,u)}(f) \leq reg_{M,\mathcal{P}^+}^{(S,X,u)}(g).$$

[5]Stoye [2011b] assumes that, for each menu M, there is a finite set A_M of acts such that M consists of all the convex combinations of the acts in A_M. We clearly allow a larger set of menus than Stoye. We return to the issue of what menus to consider after we discuss the representation theorem in Section A.2, and again when we discuss choice functions in Section 2.5.

	1 defect		10 defects	
	Payoff	Regret	Payoff	Regret
deliver	10,000	0	-10,000	10,000
cancel	0	10,000	0	0
check	5,001	4,999	-4,999	4,999

Table 2.2: Payoffs and regrets for quality-control problem.

That is, f is preferred to g if the maximum expected regret of f is less than that of g. We can similarly define $\succeq_{M,reg}$, $\succeq_{M,\Pr}^{S,X,u}$, and $\succeq_{M,\mathcal{P}}^{S,X,u}$ by replacing $reg_{M,\mathcal{P}^+}^{(S,X,u)}$ by $reg_{M}^{(S,X,u)}$, $reg_{M,\Pr}^{(S,X,u)}$, and $reg_{M,\mathcal{P}}^{(S,X,u)}$, respectively. Again, we usually omit the superscript (S, X, u) and subscript $\Pr$ or $\mathcal{P}^+$, and just write $\succeq_M$, if it is clear from context.

To see how these definitions work, consider the quality-control problem from the introduction, there are $1,000$ states with one defective item, and $C(1000, 10)$ states with ten defective items, where $C(m, n)$ is the number of combinations of n items from a collection of 1000 unique items. The regret of each action in a state depends only on the number of defective items, and is given in Table 2.2. It is easy to see that the action that minimizes regret is $check$, with $deliver$ and $cancel$ having equal regret. If we represent uncertainty using the two probability measures $\Pr_1$ and $\Pr_{10}$, the expected regret of each of the acts with respect to $\Pr_1$ (resp., $\Pr_{10}$) is just its regret with respect to states with one (resp. ten) defective items. Thus, the action that minimizes maximum expected regret is again $check$.

As we said above, the ranking of acts based on MER or MWER can change if the menu of possible choices changes. For example, suppose that we introduce a new choice in the quality-control problem, whose gains and losses are twice those of $deliver$, resulting in the payoffs and regrets described in Table 2.3. In

	1 defect		10 defects	
	Payoff	Regret	Payoff	Regret
deliver	10,000	10,000	-10,000	10,000
cancel	0	20,000	0	0
check	5,001	14,999	-4,999	4,999
new	20,000	0	-20,000	20,000

Table 2.3: Payoffs and regrets for the quality-control problem with a new choice added.

this new setting, *deliver* has a lower maximum expected regret ($10,000$) than *check* ($14,999$), so MER prefers *deliver* over *check*. Thus, the introduction of a new choice can affect the relative order of acts according to MER (and MWER), even though other acts are preferred to the new choice. By way of contrast, the decision rules MMEU and SEU are *menu-independent*; the relative order of acts according to MMEU and SEU is not affected by the addition of new acts.

We next consider a dynamic situation, where the DM acquires information about a stable randomizing mechanism (i.e., a stationary probability distribution). Specifically, in the context of the quality-control problem, suppose that the owner learns E—the first 100 items are good. Initially, suppose that The owner has no reason to believe that one hypothesis is more likely than the other, so assigns both hypotheses weight 1. Note that $P_1(E) = 0.9$ and $\Pr_{10}(E) = C(900,10)/C(1000,10) \approx 0.35$. Thus,

$$\mathcal{P}^+ \mid E = \{(\Pr_1 \mid E, 1), (\Pr_{10} \mid E, C(900,10)/(.9C(1000,10))\}.$$

We can also see from this example that MWER interpolates between MER and expected utility maximization. Suppose that a passer-by tells The owner that the first N cupcakes are good. If $N = 0$, MWER with initial weights 1 is the same as MER. On the other hand, if $N \geq 991$, then the likelihood of $\Pr_{10}$ is 0, and the only measure that has effect is $\Pr_1$, which means minimizing maximum weighted expected regret is just maximizing expected utility with respect to $\Pr_1$. If $0 < N < 991$, then the likelihoods (hence weights) of $\Pr_1$ and $\Pr_{10}$ are 1 and $\frac{C(1000-N,10)}{C(1000,10)} \times \frac{1000}{1000-N} < ((999-N)/999)^9$. Thus, as N increases, the weight of $\Pr_{10}$ goes to 0, while the weight of $\Pr_1$ stays at 1.

2.5 A Characterization of MWER

We now provide a representation theorem for MWER. That is, we provide a collection of properties (i.e., axioms) that hold of MWER such that a preference order on acts that satisfies these properties can be viewed as arising from MWER. To get such an axiomatic characterization, we restrict to what is known in the literature as the *Anscombe-Aumann* (AA) framework [1963], where outcomes are lotteries on prizes. This framework is standard in the decision theory literature; axiomatic characterizations of SEU [1963], MMEU [1989b], and MER [2008a, 2011a] have already been obtained in the AA framework. We draw on these results to obtain our axiomatization.

In this section, we provide a characterization of MWER using choice functions as the primitives. In Appendix A.2, we provide a characterization of MWER with menu-indexed preference orders as the primitives, which allows us to compare our axioms to axioms that have been used to characterize other decision rules.

A choice function maps every finite set M of acts to a subset M' of M. Intuitively, the set M' consists of the "best" acts in M. Thus, a choice function gives less information than a preference order; it gives only the top elements of the preference order. Stoye [2011b] provides a representation theorem for MER where the axioms are described in terms of choice functions.

As usual, a *choice function* C is a function from menus to menus, where $C(M) \subseteq M$.[6] We start by taking the domain of a choice to be the set $\mathcal{M}_B$ of all bounded menus. As we now show, we can get a representation theorem for MWER by using all the axioms given by Stoye [2011a], except for the "betweenness" axiom that restricts the representation to consist of probability distributions (of weight one). For completeness, we reproduce the axioms below.

Given a set Y (which we view as consisting of *prizes*), a *lottery* over Y is just a probability with finite support on Y. Let $\Delta(Y)$ consist of all finite probabilities over Y. In the AA framework, the set of outcomes has the form $\Delta(Y)$. So now acts are functions from S to $\Delta(Y)$. (Such acts are sometimes called *Anscombe-Aumann acts*.) We can think of a lottery as modeling objective uncertainty, while a probability on states models subjective uncertainty; thus, in the AA framework we have both objective and subjective uncertainty. The technical advantage of considering such a set of outcomes is that we can consider convex combinations of acts. If f and g are acts, define the act $\alpha f + (1-\alpha)g$ to be the act that maps a state s to the lottery $\alpha f(s) + (1-\alpha)g(s)$.

In this setting, we assume that there is a utility function U on prizes in Y. The utility of a lottery l is just the expected utility of the prizes obtained, that is,

$$u(l) = \sum_{\{y \in Y \,:\, l(y) > 0\}} l(y)U(y).$$

This makes sense since $l(y)$ is the probability of getting prize y if lottery l is played. The expected utility of an act f with respect to a probability $\Pr$ is then just $u(f) = \sum_{s \in S} \Pr(s)u(f(s))$, as usual. We also assume that there are at least two prizes y_1 and y_2 in Y, with different utilities $U(y_1)$ and $U(y_2)$.

[6]We use $\subseteq$ to denote subset, and $\subset$ to denote strict subset.

Given a set Y of prizes, a utility U on prizes, a state space S, and a set $\mathcal{P}^+$ of weighted probabilities on S, we can define a family $\succeq_{M,\mathcal{P}^+}^{S,\Delta(Y),u}$ of preference orders on Anscombe-Aumann acts determined by weighted regret, one per menu M, as discussed above, where u is the utility function on lotteries determined by U. For ease of exposition, we usually write $\succeq_{M,\mathcal{P}^+}^{S,Y,U}$ rather than $\succeq_{M,\mathcal{P}^+}^{S,\Delta(Y),u}$.

Axiom 2.1 (Nontriviality). *$C(M) \subset M$ for some menu M.*

Let $g(s)^*$ denote the constant act that maps all states to the outcome $g(s)$. Given a state s, define the choice function C_s by taking $f \in C_s(M)$ if and only if $f(s)^* \in C(\{g(s)^* : g \in M\})$. Thus, C_s is a choice function that is concerned only with state s.

Axiom 2.2 (Monotonicity). *If $f \in M$ and $f \in C_s(M)$ for all s, then $f \in C(M)$.*

Intuitively, this axioms says that if, for each state s, f is a best choice in M when restricted to s, then f is a best choice in M overall.

Given a menu M and an act f, let $\lambda M + (1-\lambda)f$ be the menu $\{\lambda h + (1-\lambda)f : h \in M\}$.

Axiom 2.3 (Independence). *$C(\lambda M + (1-\lambda)f) = \lambda C(M) + (1-\lambda)f$.*

Axiom 2.4 (Independence of Irrelevant Alternatives (IIA) for Constant Acts). *If M and N consist of constant acts, then*

$$C(M \cup N) \cap M \in \{C(M), \emptyset\}.$$

Axiom 2.5 (Independence of Never Strictly Optimal Alternatives (INA)). *If $C_s(M \cup N) \cap M \neq \emptyset$ for all s, then*

$$C(M \cup N) \cap M \in \{C(M), \emptyset\}.$$

Axiom 2.6 (Mixture Continuity)**.** *If $f \notin M$ and $C(M \cup \{f\}) = \{f\}, g \in M$, and h is an act, then there exists $\lambda \in (0,1)$ such that $C(M \cup \{\lambda f + (1-\lambda)h\}) = \{\lambda f + (1-\lambda)h\}$, and $\lambda g + (1-\lambda)h \notin C(M \cup \{f, \lambda g + (1-\lambda)h\})$.*

Axiom 2.7 (Ambiguity Aversion)**.** *If $\lambda \in [0,1]$, and $M \supseteq \{f, g, \lambda f + (1-\lambda)g\}$, and $\{f, g\} \subseteq C(M)$, then $\lambda f + (1-\lambda)g \in C(M)$.*

The MWER choice function is defined as

$$C^{S,Y,U}_{\mathcal{P}^+}(M) = \operatorname*{argmin}_{f \in M} reg^{(S,X,u)}_{M,\mathcal{P}^+}(f).$$

We now state and prove a representation theorem for MWER. Roughly, the representation theorem states that a choice function satisfies Axioms 2.1–2.5 if and only if it has a MWER representation with respect to some utility function and weighted probabilities. In the representation theorem for SEU [1963], not only is the utility function unique (up to affine transformations, so that we can replace U by $aU + b$, where $a > 0$ and b are constants), but the probability is unique as well. Similarly, in the MMEU representation theorem of Gilboa and Schmeidler [1989b], the utility function is unique, and the set of probabilities is also unique, as long as one assume that the set is convex and closed.

To get uniqueness in the representation theorem for MWER, we need to consider a different representation of weighted probabilities. Define a *sub-probability measure* $\mathbf{p}$ on S to be like a probability measure (i.e., a function mapping measurable subsets of S to $[0,1]$ such that $\mathbf{p}(T \cup T') = \mathbf{p}(T) + \mathbf{p}(T')$ for disjoint sets T and T'), without the requirement that $\mathbf{p} = 1$. We can identify a weighted probability distribution $(\Pr, \alpha)$ with the sub-probability measure $\alpha \Pr$. (Note that given a sub-probability measure $\mathbf{p}$, there is a unique pair $(\alpha, \Pr)$ such that $\mathbf{P} = \alpha \Pr$: we simply take $\alpha = \mathbf{p}(S)$ and $\Pr = \mathbf{p}/\alpha$.) A set B of sub-probability measures is *downward-closed* if, whenever $\mathbf{p} \in B$ and $\mathbf{q} \leq \mathbf{p}$, then $\mathbf{q} \in B$. We get a unique set of sub-probability measures in our representation theorem if we restrict to sets that are convex, downward-closed, closed, and contain at least one (proper) probability measure. (The latter requirement corresponds to having $\alpha_{\Pr} = 1$ for some $\Pr \in \mathcal{P}^+$.) For convenience, we will call a set *regular* if it is convex, downward-closed, and closed.

We identify each set of weighted probabilities $\mathcal{P}^+$ with the set of sub-probability measures

$$C(\mathcal{P}^+) = \{\alpha \Pr : (\Pr, \alpha_{\Pr}) \in \mathcal{P}^+, 0 \leq \alpha \leq \alpha_{\Pr}\}.$$

Note that if $(\alpha, \Pr) \in \mathcal{P}^+$, then $C(\mathcal{P}^+)$ includes all the sub-probability measures between the all-zero measure and $\alpha_{\Pr} \Pr$.

We need to restrict to closed and convex sets of sub-probability measures to get uniqueness in the representation of MWER for much the same reason that we need to restrict to closed and convex sets to get uniqueness in the representation of MMEU. Convexity is needed because a set B of sub-probability measures always induce the same MWER preferences as its convex hull. For example, consider the quality-control problem and the expected regrets in Table 2.2, and the distribution $a \Pr_1 + (1-a) \Pr_{10}$, for some $a \in (0, 1)$. The weighted expected regret of any act with respect to $a \Pr_1 + (1-a) \Pr_{10}$ is bounded above by the maximum weighted expected regret of that act with respect to $\Pr_1$ and $\Pr_{10}$. Therefore, adding $a \Pr_1 + (1-a) \Pr_{10}$ to $\mathcal{P}^+$ for some weight $a \in (0, 1)$ does not change the resulting family of preferences. Similarly, we need to restrict to closed sets for uniqueness, since if we start with a set B of sub-probability measures that is not closed, taking the closure of B would result in the same family of preferences.

While convexity is easy to define for a set of sub-probability measures, there seems to be no natural notion of convexity for a set $\mathcal{P}^+$ of weighted probabilities. Moreover, the requirement that $\mathcal{P}^+$ is closed is different from the requirement that $C(\mathcal{P}^+)$ is closed. The latter requirement seems more reasonable. For example, fix a probability measure $\Pr$, and let $\mathcal{P}^+ = \{(1, \Pr)\} \cup \{(0, \Pr') : \Pr' \neq \Pr\}$. Thus, $\mathcal{P}^+$ essentially consists of a single probability measure, namely $\Pr$, with weight 1; all the weighted probability measures $(0, \Pr')$ have no impact. This represents the uncertainty of an DM who is sure that that $\Pr$ is true probability. Clearly $\mathcal{P}^+$ is not closed, since we can find a sequence $\Pr_n$ such that $(0, \Pr_n) \to (0, \Pr)$, although $(0, \Pr) \notin \mathcal{P}^+$. But $C(\Pr^+)$ is closed.

Restricting to closed, convex sets of sub-probability measures does not suffice to get uniqueness; we also need to require downward-closedness. This is so because if $\mathbf{p}$ is in B, then adding any $\mathbf{q} \leq \mathbf{p}$ to the set leaves all regrets unchanged. Finally, the presence of a proper probability measure is also required, since for all $a \in (0, 1]$, scaling each element in the set B by a leaves the family of preferences unchanged.

In summary, if we consider arbitrary sets of sub-probability measures, then the set of sub-probability measures that represent a given family of MWER preferences is unique if we require it to be regular and contain a probability measure.

Although we have assumed that the set of menus is $\mathcal{M}_B$, other sets have been considered in the literature. In particular, Stoye considers the set $\mathcal{M}_C$ of menus that are convex hulls of a finite number of acts, and the set $\mathcal{M}_F$ of finite menus [2011a, 2011b]. As we now show, the representation theorem holds for both $\mathcal{M}_F$ and $\mathcal{M}_B$. In Appendix A.2, we show that if we consider preference orders as opposed to choice functions, then the corresponding representation theorem holds for $\mathcal{M}_C$ as well as $\mathcal{M}_F$ and $\mathcal{M}_B$.

Theorem 2.5.1. *For all Y, U, S, and $\mathcal{P}^+$, the choice function $C_{\mathcal{P}^+}^{S,Y,U}$ satisfies Axioms 2.1–5.9 for all $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_F$). Conversely, if a choice function C satisfies Axioms 2.1–5.9 for all $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_F$), then there exists a utility U on Y and a weighted set $\mathcal{P}^+$ of probabilities on S such that $C(\mathcal{P}^+)$ is regular and $C = C_{\mathcal{P}^+}^{S,Y,U}$. Moreover, U is unique up to affine transformations, and $C(\mathcal{P}^+)$ is unique in the sense that if $\mathcal{Q}^+$ represents C, and $C(\mathcal{Q}^+)$ is regular, then $C(\mathcal{Q}^+) = C(\mathcal{P}^+)$.*

Proof. Showing that $\succeq_{M,\mathcal{P}^+}^{S,Y,U}$ satisfies Axioms 2.1–5.9 is fairly straightforward; we leave details to the reader. Essentially the same proof works for $\mathcal{M}_B$ and $\mathcal{M}_F$. For the proof of the converse, we rely heavily on parts of the proof by Stoye [2011a]. Although Stoye considers $\mathcal{M}_F$, the arguments also work if we consider $\mathcal{M}_B$.

Stoye [2011a] shows that Axioms 2.1–5.9 imply that a menu-independent, revealed preference order $\succeq_C$ can be constructed based on the behavior of the choice function C on the set $\mathcal{M}_0$ of menus with nonpositive acts and utility frontier 0 (i.e., for every state, some act has a utility of 0). The preference order $\succeq_C$ can be thought of as the revealed preference order corresponding to the choice function C. Its definition is as follows:

$$f \succ_C g \Leftrightarrow \exists M \in \mathcal{M}_0 : f \in C(M), g \in M \backslash C(M),$$

$$f \sim_C g \Leftrightarrow \exists M \in \mathcal{M}_0 : f \in C(M), g \in C(M).$$

The arguments given by Stoye [2011a] to establish that $\succeq_C$ satisfies completeness, transitivity, nontriviality, monotonicity, mixture continuity, independence, INA, and ambiguity aversion apply verbatim to our setting here. We do not repeat Stoye's arguments here, but from here on we assume without further comment that $\succeq_C$ satisfies all these properties.

Theorem 2.5.1 can be completed by finding a MWER representation for $\succeq_C$. This follows from the following lemma. Let M_0 denote the set of all acts with nonpositive utilities.

Lemma 2.5.2. *If a preference order $\succeq$ on acts with nonpositive utilities satisfies completeness, transitivity, nontriviality, monotonicity, mixture continuity, independence, INA, and ambiguity aversion, then there exists a utility U on Y and a weighted set $\mathcal{P}^+$ of probabilities on S such that $C(\mathcal{P}^+)$ is regular and $\succeq = \succeq_{M_0,\mathcal{P}^+}^{S,Y,U}$. Moreover, U is unique up to affine transformations, and $C(\mathcal{P}^+)$ is unique in the sense that if $\mathscr{Q}^+$ represents $\succeq$, and $C(\mathscr{Q}^+)$ is regular, then $C(\mathscr{Q}^+) = C(\mathcal{P}^+)$.*

The proof of Lemma 2.5.2 is given in Appendix A.1. We use techniques in the spirit of those used by by Gilboa and Schmeidler [1989a] to represent (unweighted) MMEU. However, there are technical difficulties that arise from the fact that we do not have a key axiom that is satisfied by MMEU: *C-independence* (discussed below). The heart of the proof involves dealing with the lack of C-independence. With Lemma 2.5.2, the proof of Theorem 2.5.1 is complete. □

The axioms used in Theorem 2.5.1 can be adapted to describe choice functions over $\mathcal{M}_C$, the set of finitely generated convex menus. However, we do not know whether the equivalent of Theorem 2.5.1 holds if we restrict the domain of the choice function to be $\mathcal{M}_C$. However, in Appendix A.2 we show that if we consider preference orders instead of a choice function, then there is a collection of axioms that provide a representation theorem for $\mathcal{M}_F$, $\mathcal{M}_B$, and $\mathcal{M}_C$. The result extends to $\mathcal{M}_C$ because the set of preference orders $\succeq_M$ with respect to the set $\mathcal{M}_C$ of all convex menus also determines the preference orders with respect to the set $\mathcal{M}_F$ of all finite menus. This is not the case when we take choice functions as the primitives. As Stoye [2011a, 2013] points out, the main subtlety lies in the fact that the choice functions over convex menus can always return an interior point of the convex set, thus not providing observations of choice between the vertices of the set. Stoye believes that this is a technicality that can be overcome, so that Theorem 2.5.1 should hold even if we restrict the domain of the choice function to be $\mathcal{M}_C$. However, this conjecture has yet to be verified. In the case of preference orders, the preference orders with respect to the set $\mathcal{M}_B$ also determine the preference orders $\succeq_M$ with respect to the set $\mathcal{M}_C$, so once we have a proof for $\mathcal{M}_B$, it readily extends to $\mathcal{M}_F$ and $\mathcal{M}_C$.

As we observed, in general, we have ambiguity aversion (Axiom 5.9) for regret. *Betweenness* [1983] is a stronger notion than ambiguity aversion, which states that if an DM is indifferent between two acts, then he must also be indifferent among all convex combinations of these acts. While betweenness does not hold for regret, Stoye [2011a] gives a weaker version that does hold. A menu M has *state-independent outcome distributions* if the set $L(s) = \{y \in \Delta(Y) : \exists f \in M(f(s) = y)\}$ is the same for all states s.

Axiom 2.8 ([2011a]). *For all acts f, constant acts p, scalars $\lambda \in (0,1)$, and menus $M \supseteq \{p, f, \lambda f + (1-\lambda)p\}$ with state independent outcome distributions, if $p \notin C(M)$ and $f \notin C(M)$, then $\lambda f + (1-\lambda)p \notin C(M)$.*

The assumption that the menu has state-independent outcome distributions is critical in Axiom 2.8. Stoye [2011a] shows that Axioms 2.1–2.5 together with Axiom 2.8 characterize MER. Non-probabilistic regret (which we denote

	1 defect		10 defects	
	Payoff	Regret	Payoff	Regret
deliver	10,000	0	-10,000	10,000
$\frac{1}{2}deliver + \frac{1}{2}cancel$	5,000	5,000	-5,000	5,000
cancel	0	10,000	0	0
check	5,001	4,999	-4,999	4,999

Table 2.4: Payoffs and regrets for the quality-control problem, with *deliver* mixed with the constant act *cancel*.

	1 defect		10 defects	
	Payoff	Regret	Payoff	Regret
deliver	10,000	0	-10,000	20,000
$\frac{1}{2}deliver + \frac{1}{2}cancel$	5,000	5,000	-5,000	15,000
cancel	0	10,000	0	10,000
$check1$	-5,000	15,000	5,000	5,000
$check2$	-10,000	20,000	10,000	0

Table 2.5: Payoffs and regrets for the quality-control problem, with state-independent outcome distributions.

REG) can be viewed as a special case of MER, where $\mathcal{P}$ consists of all distributions. This means that it satisfies all the axioms that MER satisfies. As Stoye [2011b] shows, REG is characterized by Axioms 2.1–2.5 and one additional axiom, which he calls Symmetry. We omit the details here.

To see why Axiom 2.8 is needed, suppose that we change the payoffs in the quality-control problem so that *deliver* has the same maximum expected regret as *cancel* $(10,000)$. However, as seen in Table 2.4, $\frac{1}{2}deliver + \frac{1}{2}cancel$ has lower maximum expected regret $(5,000)$ than *deliver* $(10,000)$, showing that the variant of Axiom 2.8 without the state-independent outcome distribution requirement does not hold.

Although Axiom 2.8 is sound for unweighted minimax expected regret, it is no longer sound once we add weights. For example, suppose that we modified the quality-control problem so that all states we care about have the same outcome distributions, as required by Axiom 2.8. Then the payoffs and regrets will be those shown in Table 2.5. Suppose that the weights on $\Pr_1$ and $\Pr_{10}$ are 1 and 0.5, respectively. Then $deliver$ has the same maximum weighted expected regret as $cancel$ $(10,000)$. However, $\frac{1}{2}deliver + \frac{1}{2}cancel$ has lower maximum weighted expected regret $(7,500)$ than $deliver$, showing that Axiom 2.8 with weighted probabilities does not hold.

As mentioned in the introduction, Chateauneuf and Faro [2009] axiomatize a weighted version of maxmin expected utility, when utilities are restricted to be nonnegative. The expected utilities are multiplied by the reciprocal of the weights, instead of the weights themselves. Preferences are then defined by using the maxmin expected utility rule with the weighted expected utilities. To obtain uniqueness of the representation, while we require (1) a measure with weight 1, (2) downward-closedness, (3) closedness, and (4) convexity of the sub-probability measures to get uniqueness, Chateauneuf and Faro [2009] require slightly different properties. In particular, weights are represented by a function $\phi : \Delta(S) \to [0, 1]$, and Chateauneuf and Faro require that there be (1) a measure with weight 1, (2) upper semicontinuity of the function ϕ (i.e., the set $\{p \in \Delta : \phi(p) \geq \alpha\}$ is closed in the weak* topology, for all $\alpha \in [0, 1]$), and (3) quasi-concavity of ϕ (i.e., $\phi(\beta p_1 + (1-\beta)p_2) \geq \min\{\phi(p_1), \phi(p_2)\}$ for all $\beta \in [0, 1]$). It is not hard to show that convexity of the sub-probability measures imply quasi-concavity of the weights, and closedness of the sub-probability measures implies weak* upper semicontinuity of the weights. However, the converse does not hold: weak* upper semicontinuity and quasi-concavity of the weights do not imply convexity of the sub-probability measures. Therefore, our conditions to obtain uniqueness of the representation are more stringent than those of Chateauneuf and Faro.

2.6 Characterizing likelihood updating for MWER

We next consider a more dynamic setting, where DMs learn information. For simplicity, we assume that the information is always a subset E of the state space. If the DM is representing her uncertainty using a set $\mathcal{P}^+$ of weighted probability measures, then we would expect her to update $\mathcal{P}^+$ to some new set $\mathscr{Q}^+$ of weighted probability measures, and then apply MWER with uncertainty represented by$\mathscr{Q}^+$. In this section, we characterize what happens in the special case that the DM uses likelihood updating, so that $\mathscr{Q}^+ = (\mathcal{P}^+ \mid E)$.

For this characterization, we assume that the DM has a family of choice functions C_E indexed by the information E. Each choice function C_E satisfies Axioms 2.1–5.9, since the DM makes decisions after learning E using MWER. Somewhat surprisingly, all we need is one extra axiom for the characterization; we call this axiom MDC, for "menu-dependent dynamic consistency".

To explain the axiom, we need some notation. As usual, we take fEh to be the act that agrees with f on E and with h off of E; that is

$$fEh(s) = \begin{cases} f(s) & \text{if } s \in E \\ h(s) & \text{if } s \notin E. \end{cases}$$

In the quality-control problem, the act *check* can be thought of as $(\mathit{deliver})E(\mathit{cancel})$, where E is the set of states where there is only one defective item.

Roughly speaking, MDC says that you prefer f to g once you learn E if and only if, for all acts h, you also prefer fEh to gEh before you learn anything. This seems reasonable, since learning that the true state was in E is conceptually similar to knowing that none of your choices matter off of E.

To state MDC formally, we need to be careful about the menus involved. Let $MEh = \{fEh : f \in M\}$. We can identify unconditional preferences with preferences conditional on S; that is, we identify C with C_S. We also need to restrict the sets E to which MDC applies. Recall that conditioning using likelihood updating is undefined for an event such that $\overline{\mathcal{P}}^+(E) = 0$. That is, $\alpha_{\Pr} \Pr(E) = 0$ for all $\Pr \in \mathcal{P}$. As is commonly done, we capture the idea that conditioning on E is possible using the notion of a *non-null* event.

Definition 2.6.1 (Null event). *An event E is* null *if, for all $f, g \in \Delta(Y)^S$ and menus M with $fEg, g \in M$, we have $fEg \in C(M) \Leftrightarrow g \in C(M)$.*

Axiom 2.9 (MDC). *Let MEg denote the menu $\{hEg : h \in M\}$. For all $M \subseteq L, f \in M$,*

$$f \in C_E(M) \Leftrightarrow \exists g(fEg \in C(MEg)).$$

Theorem 2.6.2. *For all Y, U, S, and $\mathcal{P}^+$, the choice function $C^{S,Y,U}_{\mathcal{P}^+|E}$ for events E such that $\overline{\mathcal{P}}^+(E) > 0$ satisfies Axioms 2.1–5.9 and Axiom 2.9. Conversely, if a choice function C_E on the acts in $\Delta(Y)^S$ satisfies Axioms 2.1–5.9 and Axiom 2.9, then there exists a utility U on Y and a weighted set $\mathcal{P}^+$ of probabilities on S such that $C(\mathcal{P}^+)$ is regular, and for all non-null E, $C_E = C^{S,Y,U}_{\mathcal{P}^+|E}$. Moreover, U is unique up to affine transformations, and $C(\mathcal{P}^+)$ is unique in the sense that if $\mathcal{Q}^+$ represents C_E, and $C(\mathcal{Q}^+)$ is regular, then $C(\mathcal{Q}^+) = C(\mathcal{P}^+)$.*

Proof. Since $C = C_S$ satisfies Axioms 2.1–5.9, there must exist a weighted set $\mathcal{P}^+$ of probabilities on S and a utility function U such that $f \in C(M)$ iff $f \in$

$C^{S,Y,U}(M)$. We now show that if E is non-null, then $\overline{\mathcal{P}}^+(E) > 0$, and $f \in C_E(M)$ iff $f \in C^{S,Y,u}_{\mathcal{P}^+|E}(M)$.

For the first part, it clearly is equivalent to show that if $\overline{\mathcal{P}}^+(E) = 0$, then E is null. So suppose that $\overline{\mathcal{P}}^+(E) = 0$. Then $\alpha_{\Pr}\Pr(E) = 0$ for all $\Pr \in \mathcal{P}$. This means that $\alpha_{\Pr}\Pr(s) = 0$ for all $\Pr \in \mathcal{P}$ and $s \in E$. Thus, for all acts f and g,

$$\begin{array}{ll} & \mathit{reg}_{M,\mathcal{P}^+}(fEg) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\alpha_{\Pr}\sum_{s\in S}\Pr(s)\mathit{reg}_M(fEg,s)\right) \\ = & \sup_{\Pr\in\mathcal{P}}\ \left(\alpha_{\Pr}\left(\sum_{s\in E}\Pr(s)\mathit{reg}_M(f,s)\right) + \sum_{s\in E^c}\Pr(s)\mathit{reg}_M(g,s)\right) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\alpha_{\Pr}\sum_{s\in S}\Pr(s)\mathit{reg}_M(g,s)\right) \\ = & \mathit{reg}_{M,\mathcal{P}^+}(g). \end{array}$$

Thus, $fEg \in C(M) \Leftrightarrow g \in C(M)$ for all acts f, g and menus M containing fEg and g, which means that E is null.

For the second part, we first show that if $\overline{\mathcal{P}}^+(E) > 0$, then for all $f, h \in M$, we have that

$$\mathit{reg}_{MEh,\mathcal{P}^+}(fEh) = \overline{\mathcal{P}}^+(E)\mathit{reg}_{M,\mathcal{P}^+|E}(f).$$

We proceed as follows:

$$\begin{array}{ll} & \mathit{reg}_{MEh,\mathcal{P}^+}(fEh) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\alpha_{\Pr}\sum_{s\in S}\Pr(s)\mathit{reg}_{MEh}(fEH,s)\right) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\alpha_{\Pr}\Pr(E)\sum_{s\in E}\Pr(s\mid E)\mathit{reg}_M(f,s)\ + \alpha_{\Pr}\sum_{s\in E^c}\Pr(s)\mathit{reg}_{\{h\}}(h,s)\right) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\alpha_{\Pr}\Pr(E)\sum_{s\in E}\Pr(s|E)\mathit{reg}_M(s,f)\right) \\ = & \sup_{\Pr\in\mathcal{P}}\left(\overline{\mathcal{P}}^+(E)\alpha_{\Pr,E}\sum_{s\in E}\Pr(s|E)\mathit{reg}_M(f,s)\right) \\ & [\text{since } \alpha_{\Pr,E} = \sup_{\{\Pr'\in\mathcal{P}:\Pr'|E=\Pr|E\}}\frac{\alpha_{\Pr'}\Pr'(E)}{\overline{\mathcal{P}}^+(E)}] \\ = & \overline{\mathcal{P}}^+(E)\cdot\mathit{reg}_{M,\mathcal{P}^+|E}(f). \end{array}$$

Thus, for all $h \in M$,

$$\begin{aligned} & reg_{MEh,\mathcal{P}^+}(fEh) \leq reg_{MEh,\mathcal{P}^+}(gEh) \\ \text{iff } & \overline{\mathcal{P}}^+(E) \cdot reg_{M,\mathcal{P}^+|E}(f) \leq \overline{\mathcal{P}}^+(E) \cdot reg_{M,\mathcal{P}^+|E}(g) \\ \text{iff } & reg_{M,\mathcal{P}^+|E}(f) \leq reg_{M,\mathcal{P}^+|E}(g). \end{aligned}$$

It follows that the choice function induced by $\mathcal{P}^+$ satisfies MDC. Moreover, if Axioms 2.1–5.9 and MDC hold, then for a weighted set $\mathcal{P}^+$ that represents C, we have

$$\begin{aligned} & f \in C_E(M) \\ \text{iff} \quad & \text{for some } h, fEh \in C(MEh) \\ \text{iff} \quad & reg_{M,\mathcal{P}^+|E}(f) \leq reg_{M,\mathcal{P}^+|E}(g) \text{ for all } g, \end{aligned}$$

as desired.

Finally, the uniqueness of $C(\mathcal{P}^+)$ follows from Theorem 2.5.1, which says that C is already sufficient to guarantee the uniqueness of $C(\mathcal{P}^+)$. □

2.7 Chapter Conclusion

We proposed an alternative belief representation using *weighted sets of probabilities*, and described a natural approach to updating in such a situation and a natural approach to determining the weights. We also showed how weighted sets of probabilities can be combined with regret to obtain a decision rule, MWER, and provided an axiomatization that characterizes static and dynamic preferences induced by MWER.

One issue that must be dealt with when MWER is combined with likelihood updating is *dynamic inconsistency*. It is not hard to construct examples where a

DM decides on a plan that says that if he learns E, he should perform act f, but then when he actually learns E, he performs f' instead. Such dynamic inconsistency arises with MMEU when using measure-by-measure updating as well. Siniscalchi [2011] proposes an approach to dealing with dynamic consistency in the context of MMEU combined with measure-by-measure updating by using backward induction to decide which action to take. We believe that these ideas can be applied to MWER combined with likelihood updating as well. We hope to return to these issues in future work.

CHAPTER 3

MAXMIN WEIGHTED EXPECTED UTILITY

3.1 Overview

Chateauneuf and Faro [Chateauneuf and Faro 2009] axiomatize a weighted version of maxmin expected utility over acts with nonnegative utilities, where weights are represented by a confidence function. We argue that their representation is only one of many possible, and we axiomatize a more natural form of maxmin weighted expected utility. We also provide stronger uniqueness results.

3.2 Introduction

Maxmin expected utility (MMEU), axiomatized by Gilboa and Schmeidler [Gilboa and Schmeidler 1989b], is one of the best-studied alternatives to subjective expected utility (SEU) maximization [Savage 1954]. Its compatibility with *ambiguity-averse* preferences makes it an attractive descriptive decision model, in light of experimental evidence (e.g., the Allais Paradox [Allais 1953a] and the Ellsberg Paradox [Ellsberg 1961]) showing that intuitive decisions may violate the ambiguity neutrality, or "independence", property implied by the SEU model. In the (multiple priors) MMEU decision model, there is a set of possible probability distributions over the statespace, each giving rise to a (potentially different) expected utility value for each object of choice. A MMEU DM chooses an option that maximizes the minimum of such expected utility values.

However, even MMEU may be too restrictive a model for representing reasonable decision-making. For example, Chateauneuf and Faro [2009] (henceforth CF) point out that MMEU does not allow for "attraction for smoothing an uncertain act with the help of a positive constant act", a property that is intuitively reasonable and is demonstrated in Example 3.6.2.

To deal with this, CF consider a "weighted" version of maxmin expected utility [Gilboa and Schmeidler 1989b]. Recall that in the MMEU model, beliefs are represented by a set of probability measures over the state space. The distributions that are in the set are viewed as the possible distributions over the states. However, sometimes it makes sense to treat some distributions as "more likely" than other distributions, rather than just separating the distributions into two groups ("possible" and "impossible"). CF provide a method of treating distributions differently, by assigning a *confidence value* to each distribution. As mentioned in earlier chapters, others have independently studied similar models. In Chapter 2 we considered associating multiplicative weights with probability measures in expected-regret-minimization.

In the CF model, a high confidence value on a probability measure can be interpreted as the probability measure being "significant" or "likely to be the correct distribution," while a low confidence value on a probability measure is interpreted as the probability measure being insignificant or unlikely to be the correct distribution. These confidence values are used to scale the expected utilities of the acts in a way that reflects the relative significance of each probability measure. Since larger weights should always magnify the influence of a distribution, one must restrict to either nonnegative or nonpositive utilities. CF choose to restrict to nonnegative utilities, and they multiply the expected utilities by the multiplicative inverse of the associated confidence value. The maxmin expected utility criterion is then used to compare utility acts based on these "weighted" expected utilities. In this thesis, we use the term *weight* to refer to the final real number that we multiply the expected utilities by. In the CF model, the weight is obtained by taking the multiplicative inverse of the confidence value. Multiplying by the inverse ensures that probability measures with low confidence have a smaller effect, since they are less likely to give the minimum expected utility. This generalization of the maxmin expected utility decision rule allows for a "smoothing" effect. Instead of simply being in or out of the set of probability measures considered possible, probability measures now have finer weights associated with them.

However, CF also introduce a numerical confidence threshold $\alpha_0 > 0$; a probability measure is "discarded" (i.e., ignored) if its confidence value is below this threshold α_0. This threshold affects the resulting behavior of the decision model, as captured by the axioms characterizing the decision model. Having this threshold seems to us incompatible with the intuition behind weights. If a probability measure has low weight, we should perhaps take it less seriously than one with high weight, but there seems to be no good reason to ignore it altogether. Therefore, we define a simpler version of the decision rule where there is no threshold α_0. This simplified decision rule is characterized by removing one of the CF axioms.

Another problem with the CF approach is that of using the multiplicative inverse of the confidence value as the weight on the expected utilities. This choice seems rather arbitrary. Why not use the square of the inverse? We show that any monotonically decreasing transformation that maps $(0, 1]$ onto $\mathbb{R}^+$ (the nonnegative reals) satisfies the same axioms. Although all these transformations are characterized by the same axioms, different transformations may lead to quite different decisions.

It is not clear which transformation function is the "right" one. There is no compelling argument for using $\frac{1}{x}$ rather than, say, $\frac{1}{x^2}$. Our axiomatization leads to some important observations:

1. What is important is the composition $t \circ \phi$ of the transformation function t and the confidence function ϕ, not the confidence function itself nor the transformation function itself; it is the composition that determines the preferences.
2. Confidence values have no cardinal meaning: a confidence value of $\frac{1}{2}$ can have the same meaning as a confidence value of $\frac{1}{3}$ if the transformation t changes.

Moreover, as our results show, the confidence value and the transformation interact. In Chapter 2, we were able to get a strong uniqueness result in the context of regret by multiplying the probability measure by the weight. That is, instead of considering the set of probability measures and the associated weights separately, we consider what we called subprobability measures, which are probability measures "scaled" by a weight in $[0, 1]$. By looking at these subprobability measures, we were able to find natural properties to ensure uniqueness of the representation. Here, we show that by multiplying the probability measure by the weight, we can get a uniqueness result analogous to that for regret.

With weighted regret, there is no need to apply a transformation to the confidence values. The weights are simply the confidence values. Equivalently, the identify function is a valid transformation for weighted regret. We show that for maxmin weighted expected utility, if we restrict to *nonpositive* utilities instead of nonnegative utilities, we can also take the transformation to be the identity function. That is, we can just multiply the expected utilities by a confidence value without applying any transformations. We then replace the axiom saying that there is a worst outcome with one saying that there is a best outcome. This results in essentially the same representation theorem.

The rest of this chapter is organized as follows. Section 3.3 presents the decision setting, which differs slightly from that from the previous chapters. Section 3.4 presents the CF model and some of their results. Section 3.5 considers a generalization of the CF model. Section 3.6 presents a simpler model and provides a representation theorem.

3.3 Formal Definitions

In this section we provide definitions that will be used to present the CF results, as well as to develop our new results. We restrict to what is known in the literature as the *Anscombe-Aumann* (AA) framework [1963], where outcomes are restricted to lotteries. This framework is standard in the decision theory literature; axiomatic characterizations of SEU [Anscombe and Aumann 1963] and MMEU [Gilboa and Schmeidler 1989b] have been obtained in the AA framework.

We assume that the state space S is associated with a σ algebra, and we let $\Delta(S)$ denote the set of all probability distributions on S. Unlike in the previous chapters, we do not assume that the state space S is finite. Given a set X (which we view as consisting of *prizes* or *outcomes*), a *lottery* over X is just a probability distribution on X with finite support. Let $\underline{\Delta}(X)$ be the set of all lotteries. In the AA framework, the set of outcomes is $\underline{\Delta}(X)$. So now acts are functions from the state space S to $\underline{\Delta}(X)$. (Such acts are sometimes called *Anscombe-Aumann acts*.) We denote the set of all acts by $\mathcal{F}$. The technical advantage of considering such a set of outcomes is that we can consider convex combinations of acts. If f and g are acts, define the act $\alpha f + (1-\alpha)g$ to be the act that maps a state s to the lottery $\alpha f(s) + (1-\alpha)g(s)$.

Given a utility function U on prizes in X, the utility of a lottery $l \in \underline{\Delta}(X)$ is just the expected utility of the prizes obtained, that is,

$$u(l) = \sum_{\{x \in X \,:\, l(x) > 0\}} l(x)U(x).$$

This makes sense since $l(x)$ is the probability of getting prize x if lottery l is played. The expected utility of an act f with respect to a probability p on states is then just $u(f) = \int_S u(f(s))dp$, as usual.

3.4 CF Maxmin Expected Utility with Confidence Functions

The CF approach is formalized as follows. Let $\phi : \Delta(S) \to [0,1]$ be a confidence function on the probability measures, and let u be a utility function on lotteries over X with values in $\mathbb{R}^+$ (all instances of $\mathbb{R}^+$ in this thesis include 0). Let $L_{\alpha_0}\phi$ denote the set $\{p \in \Delta(S) : \phi(p) \geq \alpha_0\}$ for $\alpha_0 \in (0,1]$.

Definition 3.4.1. *Define $\succeq_\phi^{+,\alpha_0}$ so that*

$$f \succeq_\phi^{+,\alpha_0} g \Leftrightarrow \min_{p \in L_{\alpha_0}\phi} \frac{1}{\phi(p)} \int_S u(f)dp \geq \min_{p \in L_{\alpha_0}\phi} \frac{1}{\phi(p)} \int_S u(f)dp.$$

The superscript $+$ on $\succeq_{\phi}^{+,\alpha_0}$ indicates that the preference is defined for nonnegative utilities. Note that, according to Definition 3.4.1, a probability measure that has a confidence value (according to ϕ) lower than α_0 is simply discarded. The analogy to maxmin expected utility of Gilboa and Schmeidler [Gilboa and Schmeidler 1989b] is that the probability measure is not in the belief set. Indeed, if $\alpha_0 = 1$, then the CF approach essentially reduces to maxmin expected utility.

CF call confidence functions satisfying the following properties *regular* fuzzy sets*.

Definition 3.4.2. *The set of regular* fuzzy sets consists of all mappings* $\phi : \Delta(S) \to [0, 1]$ *satisfying the following properties:*

(a) ϕ *is* normal*:* $\{p \in \Delta(S) : \phi(p) = 1\} \neq \emptyset$.

(b) ϕ *is* weakly* upper semicontinuous*:* $\{p \in \Delta(S) : \phi(p) \geq \alpha\}$ *is weakly* closed for all* $\alpha \in [0, 1]$.

(c) ϕ *is* quasi-concave*:*

$$\forall \beta \in [0, 1](\phi(\beta p_1 + (1 - \beta)p_2) \geq \min\{\phi(p_1), \phi(p_2)\}).$$

One role of regular* fuzzy sets in the CF representation is that the condition provides a canonical representation. That is, every preference order satisfying appropriate axioms can be represented by some utility function, some $\alpha_0 > 0$, and some regular* fuzzy ϕ. Moreover, there is a ϕ^* within the set of regular* fuzzy sets generating these preferences such that ϕ^* is maximal in the sense that for every probability measure p, ϕ^* assigns weakly larger confidence to p than every other regular* fuzzy set generating these preferences.

CF consider the following axioms. In the axioms, the acts f and g are viewed as being universally quantified; given an outcome $x \in X$, we write x^* to denote the constant act that maps all states to the outcome x.

Axiom 3.1.

a. (Transitivity)*: $f \succeq g \succeq h \Rightarrow f \succeq h$.*

b. (Completeness)*: $f \succeq g$ or $g \succeq f$.*

c. (Nontriviality)*: $f \succ g$ for some acts f and g.*

Axiom 3.2 (Monotonicity)**.** *If $(f(s))^* \succeq (g(s))^*$ for all $s \in S$, then $f \succeq g$.*

Axiom 3.3 (Continuity)**.** *For all $f, g, h \in \mathcal{F}$, the sets $\{\alpha \in [0,1] : \alpha f + (1-\alpha)g \succeq h\}, \{\alpha \in [0,1] : h \succeq \alpha f + (1-\alpha)g\}$ are closed.*

Axiom 3.4 (Worst Independence)**.** *There exists a* worst outcome $\underline{x} \in X$ *such that $f \succeq \underline{x}^*$ for every $f \in \mathcal{F}$. Moreover,*

$$f \sim g \Rightarrow \alpha f + (1-\alpha)\underline{x}^* \sim \alpha g + (1-\alpha)\underline{x}^*.$$

Axiom 3.4 is reminiscent of Gilboa and Schmeidler's [Gilboa and Schmeidler 1989b] C-independence axiom of MMEU; C-independence is stronger in the sense that the independence property needs to hold not only for $\underline{x}^*$, but all other constant acts as well.

Axiom 3.5 (Independence on Constant Acts)**.**

$$\forall x, y, z \in X (x^* \sim y^* \Leftrightarrow \frac{1}{2}x^* + \frac{1}{2}z^* \sim \frac{1}{2}y^* + \frac{1}{2}z^*).$$

Axiom 3.5 is a weaker version of the more common independence axiom for constant acts, where instead of $\frac{1}{2}$ mixtures, all convex mixtures of the constant acts are allowed. CF chose to present this weaker axiom, since it was shown by

Herstein and Milnor [1953] that Axioms 3.1, 3.3 and 3.5 are sufficient to satisfy the premises of the von-Neumann-Morgenstern theorem, which says that there is an expected-utility representation for preferences over constant acts. While we could have used the more standard/stronger versions of the continuity and independence axioms, to make comparisons easier, we use the versions used by CF.

Axiom 3.6 (Ambiguity Aversion)**.**

$$f \sim g \Rightarrow pf + (1-p)g \succeq g.$$

Ambiguity aversion says that when there are two equally good alternatives, the DM prefers to hedge between these two alternatives. Ambiguity aversion is also sound for MMEU [Gilboa and Schmeidler 1989b].

Axiom 3.7 (Bounded Attraction for Certainty)**.** *There exists $\delta \geq 1$ such that for all $f \in \mathcal{F}$ and $x, y \in X$:*

$$x^* \sim f \Rightarrow \frac{1}{2}x^* + \frac{1}{2}y^* \succeq \frac{1}{2}f + \frac{1}{2}\left(\frac{1}{\delta}y^* + \left(1 - \frac{1}{\delta}\right)\underline{x}^*\right).$$

As CF point out, Axiom 5.9 implies that if an agent is indifferent between an act f and a constant act x^*, then she could strictly the convex combination of f with a constant act y^* to the combination of x^* and y^*. In particular, if we let $y^* = x^*$, then Axiom 5.9 implies that $pf + (1-p)y^* \succeq x^* = px^* + (1-p)y^*$ for all $p \in [0,1]$. CF explain that Axiom 3.7 imposes a bound on the affinity for smoothing out an uncertain act with a constant act. Continuing with our example and letting $\underline{x}^* = 0^*$ (assuming that outcomes are numbers), Axiom 3.7 implies that $\frac{1}{2}x^* + \frac{1}{2}y^* \succeq \frac{1}{2}f + \frac{1}{2\delta}y^*$ for some fixed δ specified by Axiom 3.7. The fact that there exists a $\delta > 1$ such that $\frac{1}{2}x^* + \frac{1}{2}y^* \succeq \frac{1}{2}f + \frac{1}{2\delta}y^*$ follows from monotonicity. The power of Axiom 3.7 comes from the fact that there is a single $\delta \geq 1$ such that this preference holds for all $x, y \in X$, and $f \in \mathcal{F}$.

The Bounded Attraction for Certainty axiom in the CF representation captures the lower bound α_0 in the model. Recall that if the confidence value of a probability measure is less than α_0, then that measure is considered "impossible", or ignored. CF show that the δ in the Bounded Attraction for Certainty axiom can be taken to be $\frac{1}{\alpha_0}$ in the representation. δ is roughly interpreted as an upper bound on how much the mixing of a constant act to an act can make the act more preferable. We essentially take $\alpha_0 = 0$; all probability measures into account, regardless of their weight, as long as the weight is positive. Since weighted regret already says that regret due to probability measures with low confidence is not taken seriously, there seems to be no reason to ignore probability measures of low confidence altgether. In any case, since we take $\alpha_0 = 0$, we would expect decision rule to satisfy an unbounded version of attraction for certainty. Our representation theorem shows that such an axiom is not needed to characterize maxmin weighted expected utility.

CF prove the following representation theorem:

Theorem 3.4.3 (CF representation theorem [2009])**.** *A binary relation $\succeq$ on $\mathcal{F}$ satisfies Axioms 3.1–3.7 if and only if there exists a unique non-constant function $u : X \rightarrow \mathbb{R}^+$ such that $u_{x_*} = 0$, unique up to positive linear transformations, a minimal confidence level $\alpha_0 \in (0, 1]$, and a regular* fuzzy set $\phi : \Delta(S) \rightarrow [0, 1]$ such that $\succeq = \succeq_\phi^{+,\alpha_0}$.*

Note that although CF guarantee the existence of a representation with a regular* fuzzy set, the confidence function does not necessarily need to be regular* fuzzy in order to satisfy Axioms 3.1–3.7. For example, if there are two states, s_1 and s_2, p_i is the point-mass on state s_i for $i \in \{1,2\}$, $\phi(p_1) = \phi(p_2) = 1$, and $\phi(p) = 0$ for all other probability measures p, then ϕ is not a regular* fuzzy set, since it is not quasi-concave. Nevertheless, $\succeq_{\phi}^{+,\frac{1}{2}}$ is determined by maxmin expected utility and thus must satisfy Axioms 3.1–3.7, because Axioms 3.1–3.7 are strictly weaker than the axioms for maxmin expected utility [Gilboa and Schmeidler 1989b].

3.5 t-Maxmin Weighted Expected Utility

In this section we consider a generalization of the CF approach, which we call the *t-maxmin weighted decision rule*. The t-maxmin weighted rule applies a monotonically decreasing transformation function t to the confidence values, and then uses the maxmin criterion on expected utilities multiplied by the transformed confidence values. The CF decision rule is the special case of the t-weighted maxmin decision rule, where $t(x) = \frac{1}{x}$.

Let $\phi : \Delta(S) \to [0,1]$ be a confidence function, let $t : (0,1] \to \mathbb{R}^+$ be a transformation function, and let u be a nonnegative utility function.

Definition 3.5.1 (t-maxmin weighted expected utility)**.** *Define* $\succeq_{t,\phi}^{+,\alpha_0}$ *so that*

$$f \succeq_{t,\phi}^{+,\alpha_0} g \Leftrightarrow \min_{p \in L_{\alpha_0}\phi} t(\phi(p)) \int_S u(g)dp \geq \min_{p \in L_{\alpha_0}\phi} t(\phi(p)) \int_S u(f)dp.$$

The threshold value α_0 affects the preferences $\succeq_{\phi}^{+,\alpha_0}$ only if it is larger than the smallest confidence value. That is, let $\alpha_0^*(\phi) = \max\{\alpha_0, \inf_{p\in\Delta(S)} \phi(p)\}$. It is easy to see that, for all $0 < \alpha \leq \alpha_0^*(\phi)$, we have $\succeq_{\phi}^{+,\alpha} = \succeq_{\phi}^{+,\alpha_0^*(\phi)}$.

Theorem 3.5.2 shows that it is not necessary to use the transformation $t(x) = \frac{1}{x}$ to map confidence values into weights with which expected utilities are multiplied. Other functions, such as $t(x) = \frac{1}{x^2}$, represent the same class of preference orders. However, there are some constraints on the allowed transformation functions t, since we need to "simulate" $\frac{1}{\phi(p)}$ with $t(\phi'(p))$. In addition to being strictly decreasing (a property of $t(x) = \frac{1}{x}$), the condition that there exists some $\beta > 0$ such that $[\beta, \beta/\alpha_0^*(\phi)] \subseteq range(t)$ guarantees that we can "simulate" $\frac{1}{\phi(p)}$ with $t(\phi'(p))$ for some ϕ' and α_0'. Continuity guarantees that we can find a preimage $\phi'(p)$ for every value in the range of t.

Theorem 3.5.2. *For all measurable spaces (S, Σ), consequences X, nonnegative utility functions u, confidence functions $\phi : \Delta(S) \to [0, 1]$, thresholds $\alpha_0 > 0$, and strictly decreasing, continuous transformation functions $t : (0, 1] \to \mathbb{R}^+$ such that there exists some $\beta > 0$ such that $[\beta, \beta/\alpha_0^*(\phi)] \subseteq range(t)$, there exists $\alpha_0' > 0$ and ϕ' such that*

$$\succeq_{\phi}^{+,\alpha_0} = \succeq_{t,\phi'}^{+,\alpha_0'};$$

moreover, if ϕ is regular and $t(1) = \beta$, then ϕ' is regular*.*

Theorem 3.5.2 highlights another perspective of the t-weighted maxmin expected utility representation. In addition to viewing $\phi : [0, 1]$ as a confidence function which is transformed and then applied to probability measures, we can also view $t(\phi(p))$ as a *weight* applied to the probability measure p. In this thesis, we use the term *weight* to refer to a value in $\mathbb{R}^+$ with which the expected probability is multiplied, while the term *confidence* refers to a value in $[0, 1]$ in the sense used by Chateaneuf and Faro. In the theorem statement (and later in the chapter), we take U^+ to denote a nonnegative utility function.

A corollary of Theorem 3.5.2 is a representation theorem for the CF axioms, that is, Axioms 3.1–3.7. Theorem 3.5.3 requires that $t(1) > 0$, since if $t(1) \leq 0$ and the confidence function is normal then the preferences will be trivial. Theorem 3.5.3 provides a stronger uniqueness result than Theorem 3.4.3.

Theorem 3.5.3. *Let $t : (0,1] \to \mathbb{R}^+$ be a continuous, strictly decreasing function with $t(1) > 0$ and $\lim_{x\to 0^+} t(x) > c$ for $c \in \mathbb{R}^+$. For all X, U^+, S, $\alpha_0 > 0$, and ϕ, if U^+ is nonconstant and $\alpha_0^*(\phi) \geq c$, then the preference order $\succeq_{t,\phi}^{+,\alpha_0}$ satisfies Axioms 3.1–3.7, with $\delta = \frac{c}{t(1)}$ in Axiom 3.7. Conversely, if the preference order $\succeq$ on the acts in $\mathcal{F}$ satisfies Axioms 3.1–3.7 with $t(1)\delta \leq c$ in Axiom 3.7, then there exists a nonnegative utility function U^+ on X, a threshold $\alpha_0 > 0$, and a confidence function $\phi : \Delta(S) \to [0,1]$ such that ϕ is regular* fuzzy, $t \circ \phi$ has convex upper support, and $\succeq = \succeq_{t,\phi}^{+,\alpha_0}$. Moreover, U^+ is unique up to positive linear transformations, and if S is finite, there is a sense in which ϕ is unique (see Theorem 4.5.3).*

Proof. That $\succeq_{t,\phi}^{+,\alpha_0}$ satisfies Axioms 3.1–3.7 follows from Theorem 3.4.3 and Theorem 3.5.2, since $\succeq_{t,\phi}^{+,\alpha_0} = \succeq_{\phi'}^{+,\alpha_0'}$ for some α_0' and ϕ', and $\succeq_{\phi'}^{+,\alpha_0'}$ satisfies Axioms 3.1–3.7.

Proving the converse also involves Theorems 3.4.3 and 3.5.2. If a preference order satisfies Axioms 3.1–3.7, then by Theorem 3.4.3 there exists a CF representation. Moreover, the α_0 in the construction of the representation in CF's proof of Theorem 3.4.3 is equal to $\frac{1}{\delta}$, where δ is the number in Axiom 3.7. Also recall that $\alpha_0 \leq \alpha_0^*$. Therefore, if $\lim_{x\to 0^+} t(x) > t(1)\delta$ and $t(1) > 0$, then for $\beta = t(1)$, we have $[\beta, \beta/\alpha_0^*(\phi)] \subseteq [\beta, \beta\delta] \in range(t)$ over the domain $(0,1]$. By Theorem 3.5.2, we can conclude that there exists a t-weighted maxmin expected utility representation.

The uniqueness claim follows from Theorem 4.5.3 below, which requires only Axioms 3.1–5.9. □

It is well known that for MMEU and regret, the preference order determined by a set P of probability measures is the same as that determined by the convex hull of P. Thus, to get uniqueness, Gilboa and Schmeidler [1989b] consider only convex sets of probability measures. In Chapter 2, we show that a set of sub-probability measures determine the same minimax weighted expected regret (MWER) preferences as its convex hull. Proposition 3.5.5 shows that the generalized probability measures behave in much the same way as the probability measures in MMEU and the sub-probability measures in MWER.

Given a set V of generalized probabilities, define the relation $\succeq_V$ by taking

$$f \succeq_V g \Leftrightarrow \inf_{p \in V} \int_S u(f)dp \geq \inf_{p \in V} \int_S u(g)dp.$$

It is not difficult to see that we can convert back and forth between the upper support of a weighting function and the weighting function itself. Therefore, we lose no information by looking at the upper support of a weighting function.

Proposition 3.5.4. $\succeq_{\overline{V}^{\alpha_0}_{t \circ \phi}} = \succeq^{+,\alpha_0}_{t,\phi}$.

Proof.

$$\begin{aligned}
f \succeq_{\overline{V}^{\alpha_0}_{t \circ \phi}} g \text{ iff } & \inf_{p' \in \overline{V}^{\alpha_0}_{t \circ \phi}} \int_S u(f)dp' \geq \inf_{p' \in \overline{V}^{\alpha_0}_{t \circ \phi}} \int_S u(g)dp' \\
\text{iff } & \inf_{\{q : q = t(\phi(p))p, \phi(p) > \alpha_0\}} \int_S u(f)dq \geq \inf_{\{q : q = t(\phi(p))p, \phi(p) > \alpha_0\}} \int_S u(g)dq \\
\text{iff } & \inf_{\{p : \phi(p) > \alpha_0\}} t(\phi(p)) \int_S u(f)dp \geq \inf_{\{p : \phi(p) > \alpha_0\}} t(\phi(p)) \int_S u(f)dp \\
\text{iff } & f \succeq^{+,\alpha_0}_{t,\phi} g,
\end{aligned}$$

if $\phi(p)$ is lower semi-continuous. □

Recall that, given a set V in a mixture space, $Conv(V) = \{\alpha x + (1 - \alpha)y : x, y \in V, \alpha \in [0, 1]\}$ is the convex hull of V.

Proposition 3.5.5. *If V, V' are sets of generalized probability measures and $Conv(V) = Conv(V')$, then $\succeq_V = \succeq_{V'}$.*

Proof. It suffices to show that V represents the same preferences as $Conv(V)$. Let V be a set of generalized probability measures. Given $\beta \in [0, 1]$, $p_1, p_2 \in V$, and an act $f \in \mathcal{F}$, we have

$$\beta \int u(f)dp_1 + (1 - \beta) \int u(f)dp_2 \geq \min\{\int u(f)dp_1, \int u(f)dp_2\}.$$

This means that $\beta p_1 + (1 - \beta)p_2$ can be added to V without changing the preferences, as required. □

3.5.1 Impact of the threshold

In the following example, we examine how Axiom 7 qualitatively affects the weighted maxmin expected utility preferences.

Example 3.5.6. Suppose there are two states: $S = \{s_0, s_1\}$. Consider the confidence function ϕ defined by $\phi(p) = \sqrt{p(s_1)}$. Like CF, we let $t(x) = \frac{1}{x}$, and let $\alpha_0 > 0$ be a fixed threshold value. Let $\succeq_\phi^{+,\alpha_0}$ be resulting preference relation. Let f be an act such that $u(f(s_0)) = 0$ and $u(f(s_1)) = 1$. Let c^* be a constant act with utility $c > 0$. Then we have that

$$f \succeq_\phi^{+,\alpha_0} c^* \Leftrightarrow \inf_{\{p: \sqrt{p(s_1)} \geq \alpha_0\}} \sqrt{p(s_1)} \geq c.$$

This means that f is strictly preferred to all constant acts c^* with $c < \alpha_0$, but is considered strictly worse than all constant acts c^* with $c > \alpha_0$.

Now compare this to the preference order obtained by considering the same confidence function c and weight function t, but with no threshold on the confidence. Then we have that

$$f \succeq_{\phi}^{+} c^* \Leftrightarrow \inf_{p\in\Delta(S)} \sqrt{p(s_1)} \geq c.$$

Since $\min_{p\in\Delta(S)} \sqrt{p(s_1)} = 0$, this means that f is strictly worse than all constant acts c with $c > 0$. Clearly, imposing a threshold has a nontrivial impact on the preference order.

We can also show how CF's Axiom 7 is violated by $\succeq_{\phi}^{+}$. Suppose that the worst outcome in this example (i.e., $\underline{x}$) is 0. If there is no threshold (or, equivalently, if $\alpha_0 = 0$), then $f \sim 0^*$. Thus, Axiom 7 implies that, for some fixed $\epsilon > 0$, for all outcomes y, we have that $\frac{1}{2}y^* \succeq \frac{1}{2}f + \epsilon y^*$. However,

$$\begin{array}{lll} & \frac{1}{2}y^* & \succeq_{\phi}^{+,0} \frac{1}{2}f + \epsilon y^* \\ \text{iff} & \frac{y}{2} & \geq \inf_{p\in\Delta(S)} \left(\frac{1}{\sqrt{p(s_1)}}(p(s_1)(\frac{1}{2} + \epsilon y) + (1-p(s_1))\epsilon y) \right) \\ & & = \inf_{p\in\Delta(S)} \left(\frac{\epsilon y}{\sqrt{p(s_1)}} + \frac{1}{2}\sqrt{p(s_1)} \right) \end{array}$$

It is easy to see that

$$\inf_{p\in\Delta(S)} \left(\frac{\epsilon y}{\sqrt{p(s_1)}} + \frac{1}{2}\sqrt{p(s_1)} \right) = \sqrt{2\epsilon y},$$

which means that for all $y < 8\epsilon$, we have that $\frac{1}{2}y \prec \frac{1}{2}f + \epsilon y$, contradicting Axiom 7.

3.6 Maxmin Weighted Expected Utility

3.6.1 Removing the threshold

As discussed in the previous section, it does not seem natural to discard probability measures if their confidence values do not meet some fixed threshold

$\alpha_0 > 0$. We can naturally extend the definition of t-weighted maxmin expected utility to remove the threshold α_0.

Definition 3.6.1 (t-maxmin weighted expected utility without α_0)**.** *Define* $\succeq^+_{t,\phi}$ *so that*

$$f \succeq^+_{t,\phi} g \Leftrightarrow \inf_{\{p:\phi(p)>0\}} t(\phi(p)) \int_S u(g)dp \geq \inf_{\{p:\phi(p)>0\}} t(\phi(p)) \int_S u(f)dp.$$

Clearly $\succeq^{+,\alpha_0}_{t,\phi} = \succeq^+_{t,\phi'}$ where $\phi'(p) = \phi(p)$ if $\phi(p) \geq \alpha_0$ and $\phi'(p) = 0$ if $\phi(p) < \alpha_0$. Thus, $\succeq^+_{t,\phi}$ is at least as expressive as $\succeq^{+,\alpha_0}_{t,\phi}$.

If we consider CF's preference order $\succeq^+_\phi$ without a threshold α_0, then as Example 3.6.2 below shows, Axiom 3.7 no longer holds.

Example 3.6.2. Let $S = \{s_1, s_2\}$. Let the constant act $\tilde{1}$ have constant utility 1, so that the minimum weighted expected utility of $\tilde{1}$ is 1 as long as ϕ is normal. Let $p_c \in \Delta(S)$ be the measure such that $p_c(s_1) = c$ for $c \in [0,1]$. Let ϕ be a confidence function on $\Delta(S)$ such that the confidence value for $p_c \in \Delta(S)$ is

$$\phi(p_c) = \begin{cases} 1\text{, if } c \geq \frac{1}{2} \\ \frac{1}{2^1}\text{, if } c \in [\frac{1}{8}, \frac{1}{2}) \\ \frac{1}{2^2}\text{, if } c \in [\frac{1}{32}, \frac{1}{8}) \\ \dots \\ \frac{1}{2^n}\text{, if } c \in [\frac{1}{2^{2n+1}}, \frac{1}{2^{2n-1}})\text{, for } n \in \mathbb{N}. \end{cases}$$

Clearly, ϕ is normal, since $\phi(p_{\frac{1}{2}}) = 1$. It is also easy to see from the definition that ϕ is weakly* upper semicontinuous. Lastly, to check quasi-concavity, note that a function which is nondecreasing up to a point and is nonincreasing from that point on is quasiconcave. Therefore ϕ is quasi-concave.

We describe the utility of an act f on a state space $S = \{s_1, \ldots, s_n\}$ using a *utility profile* with the format $(u(f(s_1)), \ldots, u(f(s_n)))$. Consider the sequence of acts $\{f_n\}_{n \geq 1}$ with utility profiles as follows

$$
\begin{aligned}
f_1 &= \left(2, \frac{2}{7}\right) \\
f_2 &= \left(4, \frac{4}{31}\right) \\
f_3 &= \left(8, \frac{8}{127}\right) \\
\ldots \\
f_n &= \left(2^n, \frac{2^n}{2^{2n+1} - 1}\right).
\end{aligned}
$$

Suppose, by way of contradiction, that there is a fixed $\delta \in \mathbb{R}$ such that $\succeq_\phi^+$ satisfies Axiom 3.7. In Appendix B.2, we show that for all $n \geq 1$, $f_n \sim_\phi^+ \tilde{1}$.

Now let $\tilde{m}$ be a constant act with constant utility m. The act $\frac{1}{2} f_n + \frac{1}{2\delta}\tilde{\delta}$ has utility $2^{n-1} + \frac{1}{2}$ in state s_1 and utility $\frac{2^{n-1}}{2^{2n+1}-1} + \frac{1}{2}$ in state s_2. If $c \in [\frac{1}{2^{2m+1}}, \frac{1}{2^{2m-1}})$ for $m \geq 1$, then the weighted expected utility of $\frac{1}{2} f_n + \frac{1}{2\delta}\tilde{\delta}$ with respect to p_c is at least $2^{n-m-2} + 2^{m-2}$. This means that if $n \geq 4 + 2\log_2 \delta$, then the minimum weighted expected utility of $\frac{1}{2} f_n + \frac{1}{2\delta}\tilde{\delta}$ is strictly greater than δ. The details are worked out in Appendix B.2.

On the other hand, the minimum weighted expected utility of $\frac{1}{2}\tilde{1} + \frac{1}{2}\tilde{\delta}$ is $\frac{1}{2}(1 + \delta) < \delta$ for $\delta \geq 1$. Thus, $\frac{1}{2} f_n + \frac{1}{2}\frac{1}{\delta}\tilde{\delta} \succ_{t,\phi}^+ \frac{1}{2}\tilde{1} + \frac{1}{2}\tilde{\delta}$ for sufficiently large n, violating Axiom 3.7 with $x_* = \tilde{0}$. Although Axiom 3.7 is violated, it is easy to see that Axioms 3.1–5.9 hold. Indeed, as we show, we can get a representation theorem for Axioms 3.1–5.9.

3.6.2 Maxmin weighted expected utility

It is useful to think of the CF model not as probability measures accompanied by confidence values, but rather as a set of "super-probability measures." By *super-probability measure* we mean that by multiplying a probability measure by a positive scalar in $[1, \infty)$, we get a scaled positive vector whose components may sum up to more than 1. A super-probability measure is therefore a nonnegative vector whose components sum to at least 1. This notion is analogous to the *sub-probability measures* used in Chapter 2, where a sub-probability measure is a nonnegative vector whose components sum to at most 1. Intuitively, a sub-probability measure is obtained by multiplying a probability measure by a scalar weight that is at most 1. We are also interested in sets containing both super and sub-probability measures. We will call these sets of *generalized probability measures*.

It is often helpful to consider the set of generalized probability measures *supporting* the weighting function. For generalized probability measures p and p', let $p' \geq p$ if for all $s \in S, p'(s) \geq p(s)$.

Definition 3.6.3 (Upper Support)**.** *The* upper support *of a nonnegative weighting function $t \circ \phi$ is the set $\overline{V}_{t \circ \phi} = \{p' : \exists p(\phi(p) > 0$ and $p' \geq t(\phi(p)))\}$.*

The upper support of $t \circ \phi$ contains the set of generalized probabilities $t(\phi(p))p$, as well as all generalized probabilities that are larger. Including these larger generalized probabilities does not change the underlying preferences of the upper support, since these larger generalized probabilities will never provide minimum expected utilities. While adding larger generalized probabilities does not affect the minimum expected utility, working with the upper support turns out to be technically convenient, as we shall see.

Define a relation $\succeq_{\overline{V}_{t\circ\phi}}$ by taking

$$f \succeq_{\overline{V}_{t\circ\phi}} g \Leftrightarrow \inf_{p\in\overline{V}_{t\circ\phi}} \int_S u(f)dp \geq \inf_{p\in\overline{V}_{t\circ\phi}} \int_S u(g)dp.$$

Just as before, we can convert back and forth between the upper support of a weighting function and the weighting function itself. The proof is analogous to that for Proposition 3.5.4 and is left to the reader.

Proposition 3.6.4. $\succeq_{\overline{V}_{t\circ\phi}} = \succeq^+_{t,\phi}$.

For the results beyond this point, we assume that the state space S is finite, since we make use of results due to Halpern and Leung [Halpern and Leung 2012], which are proved under the assumption of a finite state space.

Theorem 3.6.5. *Let $t : (0,1] \to \mathbb{R}^+$ be a strictly decreasing function with $t(1) > 0$. For all X, nonconstant U^+, S, and normal ϕ, the preference order $\succeq^+_{t,\phi}$ satisfies Axioms 3.1–5.9. Furthermore, if t is continuous, $\lim_{x\to 0^+} t(x) = \infty$, and the preference order $\succeq$ on the acts in $\mathcal{F}$ satisfies Axioms 3.1–5.9, then there exists a a nonnegative utility function U^+ on X and a regular* fuzzy confidence function $\phi : \Delta(S) \to [0,1]$ such that $t \circ \phi$ has convex upper support, and $\succeq = \succeq^+_{t,\phi}$. Moreover, U^+ is unique up to positive linear transformations, and ϕ is unique in the sense that if ϕ' is such that $\succeq^+_{t,\phi'} = \succeq$ and $\phi' \circ t$ has convex upper support, then $\phi = \phi'$.*

Theorem 4.5.3 characterizes t-maxmin weighted expected utility without the threshold α_0 of CF. By doing so, we show that the lower bound α_0 on the confidence or weight of probabilities is not a crucial part of the characterization of a weighted version of MMEU. Moreover, we provide a uniqueness result that is in some sense stronger than that by CF [Chateauneuf and Faro 2009], in that our uniqueness result directly identifies a "representative" set of beliefs, while the CF construction [2009] needs to be maximal in order to be unique. For example, consider a state space S with two states, and the regular* fuzzy set ϕ such that $\phi(p) = 1$ for all $p \in \Delta(S)$. Consider a second regular* fuzzy set ϕ' where $\phi'(p) = \frac{1}{1+\min_{s\in S} p(s)}$. It is not difficult to check that both sets induce the same maxmin preferences in the Chateaneuf and Faro representation, since the supports of the two regular* fuzzy sets have the same convex hull.

The requirement that $\lim_{x\to 0^+} t(x) = \infty$ is necessary to model probability measures that are arbitrarily close to being "ignored". This requirement was not necessary in the representation that made use of a lower bound α_0. However, there is another natural way to relax the constraints on t without introducing a lower bound α_0. As we show in the next section, if instead of restricting to nonnegative utilities, we restrict to nonpositive utilities, then we can drop the requirement that $\lim_{x\to 0^+} t(x) = \infty$, thus allowing a larger set of transformation functions.

3.6.3 Nonpositive utilities

Although the preceding results provide a relatively simple characterization of t-weighted maxmin expected utility, we have not yet presented the full picture. In the preceding results, just as in the CF model [Chateauneuf and Faro 2009], we have restricted utilities of acts to be nonnegative. It is easy to see why the restriction to nonnegative utilities was necessary. A larger weight makes positive utilities better but negative utilities worse. If we were to allow utilities to range over positive and negative values, the resulting decision rule would have very different, rather unintuitive behavior.

It turns out we can get a simpler decision rule, characterized almost exactly[1] by Axioms 3.1–5.9, if we look at *nonpositive utilities* instead of nonnegative utilities; in this section, we consider a representation that is restricted to *nonpositive* utilities, rather than nonnegative utilities. We use the notation U^- to indicate a nonpositive utility function.

Definition 3.6.6 (Weighted maxmin representation)**.** *Given a confidence function* $\phi : \Delta(S) \to [0, 1]$ *and strictly increasing transformation function* $t : [0, 1] \to \mathbb{R}^+$*, define* $\succeq^-_{t,\phi}$ *as follows:*

$$f \succeq^-_{t,\phi} g \Leftrightarrow \min_{p\in\Delta(S)} t(\phi(p)) \sum_{s\in S} p(s)u(f, s) \geq \min_{p\in\Delta(s)} t(\phi(p)) \sum_{s\in S} p(s)u(g, s).$$

The $^-$ superscript on $\succeq^-_{t,\phi}$ denotes that the relation is defined on acts with nonpositive utilities. One benefit of using nonpositive utilities instead of nonnegative utilities is that we no longer need to transform confidence values $\phi(p)$ in $(0, 1]$ into multiplicative weights $t(\phi(p)) \in [0, \infty)$. Instead, because a larger multiplicative confidence value results in utilities that are more negative, we can simply use the confidence function as the weights. Equivalently, we can take t to be the identity. Arguably this is the most natural choice for t, and minimizes concerns regarding which transformation function to use.

We show that preferences generated by the weighted maxmin representation is characterized by Axioms 3.1–5.9, with Axiom 3.4 replaced by the following axiom:

[1]Because we restrict to nonpositive utilities instead of nonnegative utilities, instead of a worst outcome/act we now have a best outcome/act instead. Thus Axiom 3.4 no longer holds and is replaced by Axiom 3.8.

Axiom 3.8 (Best Act Independence)**.** *There exists a best outcome* $\overline{x} \in X$ *such that* $\overline{x}^* \succeq f$ *for every* $f \in \mathcal{F}$*. Moreover,*

$$f \sim g \Rightarrow \alpha f + (1-\alpha)\overline{x}^* \sim \alpha g + (1-\alpha)\overline{x}^*.$$

In the case of nonpositive utilities, as is in the case of minimax weighted expected regret (MWER), it is useful to look at the *lower support* $\underline{V}_{t \circ \phi}$ formed by the set of sub-probabilities, defined by

$$\underline{V}_{t \circ \phi} = \{p' : \exists p (p' \leq t(\phi(p))p)\}.$$

Theorem 3.6.7. *Let* $t : [0,1] \to \mathbb{R}^+$ *be a strictly increasing, continuous transformation such that* $t(1) > 0 \geq t(0)$*. For all* X*, nonconstant* U^-*,* S*, and regular* fuzzy* ϕ*, the preference order* $\succeq^-_{t,\phi}$ *satisfies Axioms 3.1–3.3, 3.5–5.9, and 3.8. Conversely, if a preference order* $\succeq$ *on the acts in* $\mathcal{F}$ *satisfies Axioms 3.1–3.3, 3.5–5.9, and 3.8, then there exists a nonpositive utility function* U^- *on* X *and a confidence function* $\phi : \Delta(S) \to [0,1]$ *such that* ϕ *is regular* fuzzy, has convex lower support, and* $\succeq = \succeq^-_{t,\phi}$*. Moreover,* U^- *is unique up to positive linear transformations, and* ϕ *is unique in the sense that if* ϕ' *is such that* $\succeq^-_{t,\phi'} = \succeq$ *and* $\phi \circ t$ *has convex lower support, then* $\phi = \phi'$*.*

Note that the transformation t in Theorem 3.6.7 has domain $[0,1]$ instead of $(0,1)$. This is because in a setting with nonpositive utilities, a confidence value of 0 can be mapped to a weight of 0, contributing nothing to the definition of the preferences. This is analogous to a measure being ignored in the case of nonnegative utilities. Furthermore, t is required to be strictly increasing, instead of decreasing, since a larger multiplier amplifies the significance of a negative utility value. We need that $t(1) > 0$, since if $t(1) = 0$ then the preferences will be trivial. In the second part of the theorem, we need $t(0) \leq 0$ in order to find a representation for all possible preferences that satisfy the axioms. For example, suppose the preference $\succeq$ is such that $(c, 0) \sim (c', 0)$ for all $c, c' \in \mathbb{R}-$. Intuitively, this means that the first state is ignored. More precisely, any probability measure giving positive probability to the first state should be ignored. If $t(0) > 0$, then we do not have the representation power to ignore these probability measures. Therefore, we are unable to find a representation for $\succeq$.

3.6.4 The case of general acts

We have considered two different settings, one restricted to nonnegative utilities, and one restricted to nonpositive utilities. One might wonder whether a maxmin weighted expected utility representation could apply to a setting that include both positive and negative utilities. Recall that in the case of nonnegative utilities, a large positive multiplier on the utility *decreases* the impact of the constraint or weighted probability measure, while in the case of nonpositive utilities, a large positive multiplier on the utility *increases* the impact of the constraint or weighted probability measure. As a result, to have reasonable behavior when dealing with both positive and negative utilities, the multiplier on a utility value must depend not only on the probability measure, but also on the utility value itself (whether it is positive or negative).

CHAPTER 4

MINIMIZING REGRET IN DYNAMIC DECISION PROBLEMS

4.1 Outline

The menu-dependent nature of regret-minimization creates subtleties when it is applied to dynamic decision problems. It is not clear whether forgone opportunities should be included in the menu. We explain commonly observed behavioral patterns as minimizing regret when forgone opportunities are present. If forgone opportunities are included, we can characterize when a form of dynamic consistency is guaranteed.

4.2 Introduction

In this chapter, we reconsider the minimax *weighted* expected regret (MWER) decision rule introduced in Chapter 2. Recall that, for MWER, uncertainty is represented by a set of *weighted* probability measures. Intuitively, the weight represents how likely the probability measure is to be the true distribution over the states, according to the DM. The weights work much like a "second-order" probability on the set of probability measures.

Real-life problems are often dynamic, with many stages where actions can be taken; information can be learned over time. Before applying regret minimization to dynamic decision problems, there is a subtle issue that we must consider. In static decision problems, the regret for each act is computed with respect to a *menu*. That is, each act is judged against the other acts in the menu.

Typically, we think of the menu as consisting of the *feasible acts*, that is, the ones that the DM can perform. The analogue in a dynamic setting would be the feasible *plans*, where a plan is just a sequence of actions leading to a final outcome. In a dynamic decision problem, as more actions are taken, some plans become *forgone opportunities*. These are plans that were initially available to the DM, but are no longer available due to earlier actions of the DM. Since regret intuitively captures comparison of a choice against its alternatives, it seems reasonable for the menu to include all the feasible plans at the point of decision-making. But should the menu include forgone opportunities?

Consequentialists would argue that it is irrational to care about forgone opportunities [Hammond 1976; Machina 1989]; we should simply focus on the opportunities that are still available to us, and thus not include forgone opportunities in the menu. And, indeed, when regret has been considered in dynamic settings thus far (e.g., by Hayashi [2011]), the menu has not included forgone opportunities. However, introspection tells us that we sometimes do take forgone opportunities into account when we feel regret. For example, when considering a new job, one might compare the available options to what might have been available if one had chosen a different career path years ago. As we show, including forgone opportunities in the menu can make a big difference in behavior. Consider procrastination: we tell ourselves that we will start studying for an exam (or start exercising, or quit smoking) tomorrow; and then tomorrow comes, and we again tell ourselves that we will do it, starting tomorrow. This behavior is hard to explain with standard decision-theoretic approaches, especially when we assume that no new information about the world is gained over time. However, we give an example where, if forgone opportunities are not included in the menu, then we get procrastination; if they are, then we do not get procrastination.

This example can be generalized. Procrastination is an example of *preference reversal*: the DM's preference at time t for what he should do at time $t+1$ reverses when she actually gets to time $t + 1$. We prove in Section 4.4 that if the menu includes forgone opportunities and the DM acquires no new information over time (as is the case in the procrastination problem), then a DM who uses regret to make her decisions will not suffer preference reversals. Thus, we arguably get more rational behavior when we include forgone opportunities in the menu.

What happens if the DM does get information over time? It is well known that, in this setting, expected utility maximizers are guaranteed to have no preference reversals. Epstein and Le Breton [1993] have shown that, under minimal assumptions, to avoid preference reversals, the DM must be an expected utility maximizer. On the other hand, Epstein and Schneider [2003] show that a DM using MMEU never has preference reversals if her beliefs satisfy a condition they call *rectangularity*. Hayashi [2011] shows that rectangularity also prevents preference reversals for MER under certain assumptions. Unfortunately, the rectangularity condition is often not satisfied in practice. Other conditions have been provided that guarantee dynamic consistency for ambiguity-averse decision rules (see, e.g., [Al-Najjar and Weinstein 2009] for an overview).

We consider the question of preference reversal in the context of regret. Hayashi [2011] has observed that, in dynamic decision problems, both changes in menu over time and updates to the DM's beliefs can result in preference reversals. In Section 4.5, we show that keeping forgone opportunities in the menu is necessary in order to prevent preference reversals. But, as we show by example, it is not sufficient if the DM acquires new information over time. We then provide a condition on the beliefs that is necessary and sufficient to guarantee that a DM making decisions using MWER whose beliefs satisfy the condition will not have preference reversals. However, because this necessary and sufficient condition may not be easy to check, we also give simpler sufficient condition, similar in spirit to Epstein and Schneider's [2003] rectangularity condition. Since MER can be understood as a special case of MWER where all weights are either 1 or 0, our condition for dynamic consistency is also applicable to MER.

The remainder of the chapter is organized as follows. Section 4.3 discuss preliminaries. Section 4.4 introduces forgone opportunities. Section 4.5 gives conditions under which consistent planning is not required.

4.3 Preliminaries

4.3.1 Static decision setting and regret

In Chapter 2, we introduced another representation of uncertainty, *weighted set of probability measures*. A weighted set of probability measures generalizes a set of probability measures by associating each measure in the set with a weight, intuitively corresponding to the reliability or significance of the measure in capturing the true uncertainty of the world. Minimizing weighted expected regret with respect to a weighted set of probability measures gives a variant of minimax regret, called Minimax Weighted Expected Regret (MWER). A set $\mathcal{P}^+$ of *weighted probability measures* on (S, Σ) consists of pairs $(\Pr, \alpha_{\Pr})$, where $\alpha_{\Pr} \in [0, 1]$ and $\Pr$ is a probability measure on (S, Σ). Let $\mathcal{P} = \{\Pr : \exists \alpha (\Pr, \alpha) \in \mathcal{P}^+\}$. We assume that, for each $\Pr \in \mathcal{P}$, there is exactly one α such that $(\Pr, \alpha) \in \mathcal{P}^+$.

We denote this number by $\alpha_{\Pr}$, and view it as the *weight of* $\Pr$. We further assume for convenience that weights have been normalized so that there is at least one measure $\Pr \in \mathcal{P}$ such that $\alpha_{\Pr} = 1$.

4.3.2 Dynamic decision problems

A *dynamic decision problem* is a single-player extensive-form game where there is some set S of states, nature chooses $s \in S$ at the first step, and does not make any more moves. The DM then performs a finite sequence of actions until some outcome is reached. Utility is assigned to these outcomes. A *history* is a sequence recording the actions taken by nature and the DM. At every history h, the DM considers possible some other histories. The DM's *information set* at h, denoted $I(h)$, is the set of histories that the DM considers possible at h. Let $s(h)$ denote the initial state of h (i.e., nature's first move); let $R(h)$ denote all the moves the DM made in h after nature's first move; finally, let $E(h)$ denote the set of states that the DM considers possible at h; that is, $E(h) = \{s(h') : h' \in I(h)\}$. We assume that the DM has *perfect recall*: this means that $R(h') = R(h)$ for all $h' \in I(h)$, and that if h' is a prefix of h, then $E(h') \supseteq E(h)$.

A *plan* is a (pure) strategy: a mapping from histories to histories that result from taking the action specified by the plan. We require that a plan specify the same action for all histories in an information set; that is, if f is a plan, then for all histories h and $h' \in I(h)$, we must have the last action in $f(h)$ and $f(h')$ must be the same (so that $R(f(h)) = R(f(h'))$). Given an initial state s, a plan determines a complete path to an outcome. Hence, we can also view plans as acts: functions mapping states to outcomes. We take the acts in a dynamic decision problem to be the set of possible plans, and evaluate them using the decision rules discussed above.

A major difference between our model and that used Epstein and Schneider [2003] and Hayashi [2009] is that the latter assume a *filtration* information structure. With a filtration information structure, the DM's knowledge is represented by a fixed, finite sequence of partitions. More specifically, at time t, the DM uses a partition $F(t)$ of the state space, and if the true state is s, then all that the DM knows is that the true state is in the cell of $F(t)$ containing s. Since the sequence of partitions is fixed, the DM's knowledge is independent of the choices that she makes, and her options and preferences cannot depend on past choices. This assumption significantly restricts the types of problems that can be naturally modeled. For example, if the DM prefers to have one apple over two oranges at time t, then this must be her time t preference, regardless of whether she has already consumed five apples at time $t-1$. Moreover, consuming an apple at time t cannot preclude consuming an apple at time $t+1$. Since we effectively represent a decision problem as a single-player extensive-form game, we can capture all of these situations in a straightforward way. The models of Epstein, Schneider, and Hayashi can be viewed as a special case of our model.

In a dynamic decision problem, as we shall see, two different menus are relevant for making a decision using regret-minimization: the menu with respect to which regrets are computed, and the menu of feasible choices. We formalize this dependence by considering *choice functions* of the form $C_{M,E}$, where $E, M \neq \emptyset$. $C_{M,E}$ is a function mapping a nonempty menu M' to a nonempty subset of M'. Intuitively, $C_{M,E}(M')$ consists of the DM's most preferred choices from the menu M' when she considers the states in E possible and her decision are made relative to menu M. (So, for example, if the DM is making her choices choices using regret minimization, the regret is taken with respect to M.) Note that there may be more than one plan in $C_{M,E}(M')$; intuitively, this means that the DM does not view any of the plans in $C_{M,E}(M')$ as strictly worse than some other plan.

What should M and E be when the DM makes a decision at a history h? We always take $E = E(h)$. Intuitively, this says that all that matters about a history as far as making a decision is the set of states that the DM considers possible; the previous moves made to get to that history are irrelevant. As we shall see, this seems reasonable in many examples. Moreover, it is consistent with our choice of taking probability distributions only on the state space.

The choice of M is somewhat more subtle. The most obvious choice (and the one that has typically been made in the literature, without comment) is that M consists of the plans that are still feasible at h, where a plan f is *feasible* at a history h if, for all strict prefixes h' of h, $f(h')$ is also a prefix of h. So f is feasible at h if h is compatible with all of f's moves. Let M_h be the set of plans feasible at h. While taking $M = M_h$ is certainly a reasonable choice, as we shall see, there are other reasonable alternatives.

Before addressing the choice of menu in more detail, we consider how to apply regret in a dynamic setting. If we want to apply MER or MWER, we must update the probability distributions. Epstein and Schneider [2003] and Hayashi [2009] consider *prior-by-prior updating*, the most common way to update a set of probability measures, defined as follows:

$$\mathcal{P}|^p E = \{\Pr|E : \Pr \in \mathcal{P}, \Pr(E) > 0\}.$$

We can also apply prior-by-prior updating to a weighted set of probabilities:

$$\mathcal{P}^+|^p E = \{(\Pr|E, \alpha) : (\Pr, \alpha) \in \mathcal{P}^+, \Pr(E) > 0\}.$$

Prior-by-prior updating can produce some rather counter-intuitive outcomes. For example, suppose we have a coin of unknown bias in $[0.25, 0.75]$, and flip it 100 times. We can represent our prior beliefs using a set of probability measures. However, if we use prior-by-prior updating, then after each flip of the coin the set $\mathcal{P}^+$ representing the DM's beliefs does not change, because the beliefs are independent. Thus, in this example, prior-by-prior updating is not capturing the information provided by the flips.

We consider likelihood updating, introduced in Chapter 2:

$$\mathcal{P}^+|^l E = \{(\Pr|E, \alpha_E^l) : (\Pr, \alpha) \in \mathcal{P}^+, \Pr(E) > 0\}.$$

In computing $\mathcal{P}^+|^l E$, we update not just the probability measures in $\Pr \in \mathcal{P}$, but also their weights, which are updated to α_E^l. Although prior-by-prior updating does not change the weights, for purposes of exposition, given a weighted probability measure $(\Pr, \alpha)$, we use α_E^p to denote the "updated weight" of $\Pr|E \in \mathcal{P}^+|^p E$; of course, $\alpha_E^p = \alpha$.

Intuitively, probability measures that are supported by the new information will get larger weights using likelihood updating than those not supported by the new information. Clearly, if all measures in $\mathcal{P}$ start off with the same weight and assign the same probability to the event E, then likelihood updating will give the same weight to each probability measure, resulting in measure-by-measure updating. This is not surprising, since such an observation E does not give us information about the relative likelihood of measures.

Let $reg_M^{\mathcal{P}^+|^l E}(f)$ denote the regret of act f computed with respect to menu M and beliefs $\mathcal{P}^+|^l E$. If $\mathcal{P}^+|^l E$ is empty (which will be the case if $\overline{\mathcal{P}}^+(E) = 0$) then $reg_M^{\mathcal{P}^+|^l E}(f) = 0$ for all acts f. We can similarly define $reg_M^{\mathcal{P}^+|^p E}(f)$ for beliefs updated using prior-by-prior updating. Also, let $C_M^{reg,\mathcal{P}^+|^l E}(M')$ be the set of acts in M' that minimize the weighted expected regret $reg_M^{\mathcal{P}^+|^l E}$. If $\mathcal{P}^+|^l E$ is empty, then $C_M^{reg,\mathcal{P}^+|^l E}(M') = M'$. We can similarly define $C_M^{reg,\mathcal{P}^+|^p E}$, $C_M^{reg,\mathcal{P}|E}$ and $C_M^{reg,\Pr|E}$.

4.4 Forgone opportunities

As we have seen, when making a decision at a history h in a dynamic decision problem, the DM must decide what menu to use. In this section we focus on one choice. Take a *forgone opportunity* to be a plan that was initially available to the DM, but is no longer available due to earlier actions. As we observed in the introduction, while it may seem irrational to consider forgone opportunities, people often do. Moreover, when combined with regret, behavior that results by considering forgone opportunities may be arguably *more* rational than if forgone opportunities are not considered. Consider the following example.

Example 4.4.1. Suppose that a student has an exam in two days. She can either start studying today, play today and then study tomorrow, or just play on both days and never study. There are two states of nature: one where the exam is difficult, and one where the exam is easy. The utilities reflect a combination of the amount of pleasure that the student derives in the next two days, and her score on the exam relative to her classmates. Suppose that the first day of play gives the student $p_1 > 0$ utils, and the second day of play gives her

$p_2 > 0$ utils. Her exam score affects her utility only in the case where the exam is hard and she studies both days, in which case she gets an additional g_1 utils for doing much better than everyone else, and in the case where the exam is hard and she never studies, in which case she loses $g_2 > 0$ utils for doing much worse than everyone else. Figure 4.1 provides a graphical representation of the decision problem. Since, in this example, the available actions for the DM are independent of nature's move, for compactness, we omit nature's initial move (whether the exam is easy or hard). Instead, we describe the payoffs of the DM as a pair $[a_1, a_2]$, where a_1 is the payoff if the exam is hard, and a_2 is the payoff if the exam is easy.

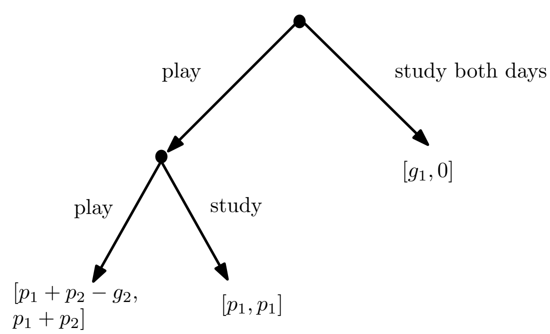

Figure 4.1: An explanation for procrastination.

Assume that $2p_1+p_2 > g_1 > p_1+p_2$ and $2p_2 > g_2 > p_2$. That is, if the test were hard, the student would be happier studying and doing well on the test than she would be if she played for two days, but not too much happier; similarly, the penalty for doing badly in the exam if the exam is hard and she does not study is greater than the utility of playing the second day, but not too much greater. Suppose that the student uses minimax regret to make her decision. On the first day, she observes that playing one day and then studying the next day has a worst-case regret of $g_1 - p_1$, while studying on both days has a worst-case regret of p_1+p_2. Therefore, she plays on the first day. On the next day, suppose that she does not consider forgone opportunities and just compares her two available options, studying and playing. Studying has a worst-case regret of p_2, while playing has a worst-case regret of $g_2 - p_2$, so, since $g_2 < 2p_2$, she plays again on the second day. On the other hand, if the student had included the forgone opportunity in the menu on the second day, then studying would have regret $g_1 - p_1$, while playing would have regret $g_1 + g_2 - p_1 - p_2$. Since $g_2 > p_2$, studying minimizes regret. □

Example 4.4.1 emphasizes the roles of the menus M and M' in $C_{M,E}(M')$. Here we took M, the menu relative to which choices were evaluated, to consist of all plans, even the ones that were no longer feasible, while M' consisted of only feasible plans. In general, to determine the menu component M of the choice function $C_{M,E(h)}$ used at a history h, we use a *menu-selection function* μ. The menu $\mu(h)$ is the menu relative to which choice are computed at h. We sometimes write $C_{\mu,h}$ rather than $C_{\mu(h),E(h)}$.

We can now formalize the notion of *no preference reversal*. Roughly speaking, this says that if a plan f is considered one of the best at history h and is still feasible at an extension h' of h, then f will still be considered one of the best plans at h'.

Definition 4.4.2 (No preference reversal). *A family of choice functions $C_{\mu,h}$ has* no preference reversals *if, for all histories h and all histories h' extending h, if $f \in C_{\mu,h}(M_h)$ and $f \in M_{h'}$, then $f \in C_{\mu,h'}(M_{h'})$.*

The fact that we do not get a preference reversal in Example 4.4.1 if we take forgone opportunities into account here is not just an artifact of this example. As we now show, as long as we do not get new information and also use a constant

	Hard		Easy	
	Short	Long	Short	Long
$\Pr_1$	1	0	0	0
$\Pr_2$	0	0.2	0.2	0.2
play-study	1	0	5	0
play-play	0	3	0	3

Table 4.1: $\alpha_{\Pr_1} = 1, \alpha_{\Pr_2} = 0.6$.

menu (i.e., by keeping all forgone opportunities in the menu), then there will be no preference reversals if we minimize (weighted) expected regret in a dynamic setting.

Proposition 4.4.3. *If, for all histories h, h', we have $E(h) = S$ and $\mu(h) = \mu(h')$, and decisions are made according to MWER (i.e., the agent has a set $\mathcal{P}^+$ of weighted probability distributions and a utility function u, and $f \in C_{\mu,h}(M_h)$ if f minimizes weighted expected regret with respect to $\mathcal{P}^+|^l E(h)$ or $\mathcal{P}^+|^p E(h)$), then no preference reversals occur.*

Proof. Suppose that $f \in C_{\mu,\langle s\rangle}$, h is a history extending $\langle s\rangle$, and $f \in M_h$. Since $E(h) = S$ and $\mu(h) = \mu(\langle s\rangle)$ by assumption, we have $C_{\mu(h),E(h)} = C_{\mu(\langle s\rangle),E(\langle s\rangle)}$. By assumption, $f \in C_{\mu(\langle s\rangle),M_{\langle s\rangle}}(M_{\langle s\rangle}) = C_{\mu(h),E(h)}(M_{\langle s\rangle})$. It is easy to check that MWER satisfies what is known in decision theory as *Sen's α axiom* [1988]: if $f \in M' \subseteq M''$ and $f \in C_{M,E}(M'')$, then $f \in C_{M,E}(M')$. That is, if f is among the most preferred acts in menu M'', if f is in the smaller menu M', then it must also be among the most preferred acts in menu M'. Because $f \in M_h \subseteq M_{\langle s\rangle}$ and $f \in C_{\mu,\langle s\rangle}(M_{\langle s\rangle})$, we have $f \in C_{\mu(h),E(h)}(M_h)$, as required. □

Proposition 4.4.3 shows that we cannot have preference reversals if the DM does not learn about the world. However, if the DM learns about the world, then we can have preference reversals. Suppose, as is depicted in Table 4.1, that in addition to being hard and easy, the exam can also be short or long. The student's beliefs are described by the set of weighted probabilities $\Pr_1$ and $\Pr_2$, with weights 1 and 0.6, respectively.

We take the option of studying on both days out of the picture by assuming that its utility is low enough for it to never be preferred, and for it to never affect the regret computations. After the first day, the student learns whether the exam will be hard or easy. One can verify that the ex ante regret of playing then studying is lower than that of playing on both days, while after the first day, the student prefers to play on the second day, regardless of whether she learns that the exam is hard or easy.

4.5 Characterizing no preference reversal

We now consider conditions under which there is no preference reversal in a more general setting, where the DM can acquire new information. While including all forgone opportunities is no longer a sufficient condition to prevent preference reversals, it is necessary, as the following example shows: Consider the two similar decision problems depicted in Figure 4.2. Note that at the node

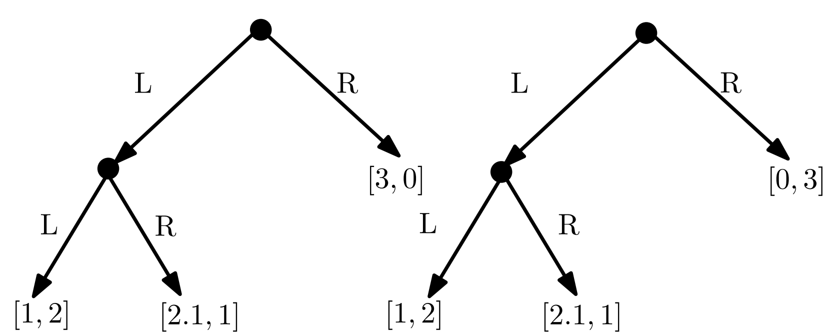

Figure 4.2: Two decision trees.

after first playing L, the utilities and available choices are identical in the two problems. If we ignore forgone opportunities, the DM necessarily makes the same decision in both cases if his beliefs are the same. However, in the tree to the left, the ex ante optimal plan is LR, while in the tree to the right, the ex ante optimal plan is LL. If the DM ignores forgone opportunities, then after the first step, she cannot tell whether she is in the decision tree on the left side, or the one on the right side. Therefore, if she follows the ex ante optimal plan in one of the trees, she necessarily is not following the ex ante optimal plan in the other tree.

In light of this example, we now consider what happens if the DM learns information over time. Our no preference reversal condition is implied by a well-studied notion called *dynamic consistency*. One way of describing dynamic consistency is that a plan considered optimal at a given point in the decision process is also optimal at any preceding point in the process, as well as any future point that is reached with positive probability [Siniscalchi 2011]. For menu-independent preferences, dynamic consistency is usually captured axiomatically by variations of an axiom called *Dynamic Consistency* (DC) or the *Sure Thing Principle* [Savage 1954]. We define a *menu-dependent* version of DC relative to events E and F using the following axiom. The second part of the axiom implies that if f is strictly preferred conditional on $E \cap F$ and at least weakly preferred on $E^c \cap F$, then f is also strictly preferred on F. An event E is *relevant to a dynamic decision problem* $\mathcal{D}$ if it is one of the events that the DM can potentially learn in D, that is, if there exists a history h such that $E(h) = E$. A dynamic decision problem $\mathcal{D} = (S, \Sigma, X, u, \mathcal{P})$ is "proper" if Σ is generated by the subsets of S relevant to $\mathcal{D}$. Given a decision problem D, we take the *measurable sets* to be the σ-algebra generated by the events relevant to D. The following axioms hold for all measurable sets E and F, menus M and M', and acts f and g.

Axiom 4.1 (DC-M). *If $f \in C_{M,E\cap F}(M') \cap C_{M,E^c\cap F}(M')$, then $f \in C_{M,F}(M')$. If, furthermore, $g \notin C_{M,E\cap F}(M')$, then $g \notin C_{M,F}(M')$.*

Axiom 4.2 (Conditional Preference). *If f and g, when viewed as acts, give the same outcome on all states in E, then $f \in C_{M,E}(M')$ iff $g \in C_{M,E}(M')$.*

The next two axioms put some weak restrictions on choice functions.

Axiom 4.3. *$C_{M,E}(M') \subseteq M'$ and $C_{M,E}(M') \neq \emptyset$ if $M' \neq \emptyset$.*

Axiom 4.4 (Sen's α). *If $f \in C_{M,E}(M')$ and $M'' \subseteq M'$, then $f \in C_{M,E}(M'')$.*

Theorem 4.5.1. *For a dynamic decision problem D, if Axiom 4.1–4.4 hold and $\mu(h) = M$ for some fixed menu M, then there will be no preference reversals in D.*

We next provide a representation theorem that characterizes when Axioms 4.1–4.4 hold for a MWER DM. The following condition says that the unconditional regret can be computed by separately computing the regrets conditional on measurable events $E \cap F$ and on $E^c \cap F$.

Definition 4.5.2 (SEP). *The weighted regret of f with respect to M and $\mathcal{P}^+$ is* separable *with respect to $|^\chi$ ($\chi \in \{p, l\}$) if for all measurable sets E and F such that $\overline{\mathcal{P}}^+(E \cap F) > 0$ and $\overline{\mathcal{P}}^+(E^c \cap F) > 0$,*

$$reg_M^{\mathcal{P}^+|^\chi F}(f) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|^\chi (E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\chi (E^c\cap F)}(f) \right),$$

and if $reg_M^{\mathcal{P}^+|^\chi (E\cap F)}(f) \neq 0$, then

$$reg_M^{\mathcal{P}^+|^\chi F}(f) > \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\chi (E^c\cap F)}(f).$$

We now show that Axioms 4.1–4.4 characterize SEP. Say that a decision problem $\mathcal{D}$ is *based on* (S, Σ) if $\mathcal{D} = (S, \Sigma, X, u, \mathcal{P})$ for some X, u, and $\mathcal{P}$. In the following results, we will also make use of an alternative interpretation of weighted probability measures. Define a *subprobability measure* p on (S, Σ) to be like a probability measure, in that it is a function mapping measurable subsets of S to $[0, 1]$ such that $p(T \cup T') = p(T) + p(T')$ for disjoint sets T and T', except that it may not satisfy the requirement that $p(S) = 1$. We can identify a weighted probability distribution $(\Pr, \alpha)$ with the subprobability measure $\alpha \Pr$. (Note that given a subprobability measure p, there is a unique pair $(\alpha, \Pr)$ such that $p = \alpha \Pr$: we simply take $\alpha = p(S)$ and $\Pr = p/\alpha$.) Given a set $\mathcal{P}^+$ of weighted probability measures, we let $C(\mathcal{P}^+) = \{p \geq \vec{0} : \exists c, \exists \Pr, (c, \Pr) \in \mathcal{P}^+ \text{ and } p \leq c \Pr\}$.

Theorem 4.5.3. *If $\mathcal{P}^+$ is a set of weighted distributions on (S, Σ) such that $C(\mathcal{P}^+)$ is closed, then the following are equivalent for $\chi \in \{p, l\}$:*

(a) For all decision problems D based on (S, Σ) and all menus M in D, Axioms 4.1–4.4 hold for the family $C_M^{reg, \mathcal{P}^+|^\chi E}$ of choice functions.

(b) For all decision problems D based on (S, Σ), states $s \in S$, and acts $f \in M_{\langle s \rangle}$, the weighted regret of f with respect to $M_{\langle s \rangle}$ and $\mathcal{P}^+$ is separable with respect to $|^\chi$.

Note that Theorem 4.5.3 says that to check that Axioms 1–4 hold, we need to check only that separability holds for initial menus $M_{\langle s \rangle}$.

It is not hard to show that SEP holds if the set $\mathcal{P}$ is a singleton. But, in general, it is not obvious when a set of probability measures is separable. We thus provide a characterization of separability, in the spirit of Epstein and LeBreton's [1993] rectangularity condition. We actually provide two conditions, one for the case of prior-by-prior updating, and another for the case of likelihood updating. These definitions use the notion of *maximum weighted expected value of* θ, defined as $\overline{E}_{\mathcal{P}^+}(\theta) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \sum_{s\in S} \alpha \Pr(s)\theta(s)$. We use $\overline{X}$ to denote the closure of a set X.

Definition 4.5.4 (χ-Rectangularity)**.** *A set $\mathcal{P}^+$ of weighted probability measures is* χ-rectangular *($\chi \in \{p, l\}$) if for all measurable sets E and F,*

(a) if $(\Pr_1, \alpha_1), (\Pr_2, \alpha_2), (\Pr_3, \alpha_3) \in \mathcal{P}^+$, $\Pr_1(E \cap F) > 0$, and $\Pr_2(E^c \cap F) > 0$, then

$$\alpha_3 \Pr_3(E\cap F)\alpha^{\chi}_{1,E\cap F}\Pr_1|(E\cap F) + \alpha_3 \Pr_3(E^c\cap F)\alpha^{\chi}_{2,E^c\cap F}\Pr_2|(E^c\cap F) \in \overline{C(\mathcal{P}^+ \mid^{\chi} F)},$$

(b) for all $\delta > 0$, if $\overline{\mathcal{P}}^+(F) > 0$, then there exists $(\Pr, \alpha) \in \mathcal{P}^+|^{\chi}F$ such that $\alpha(\delta \Pr(E \cap F) + \Pr(E^c \cap F)) > \sup_{(\Pr',\alpha')\in\mathcal{P}^+} \alpha' \Pr'(E^c \cap F)$, and

(c) for all nonnegative real vectors $\theta \in \mathbb{R}^{|S|}$,

$$\sup_{(\Pr,\alpha)\in\mathcal{P}^+|^{\chi}F} \alpha \left(\Pr(E \cap F)\overline{E}_{\mathcal{P}^+|^{\chi}(E\cap F)}(\theta) + \Pr(E^c \cap F)\overline{E}_{\mathcal{P}^+|^{\chi}(E^c\cap F)}(\theta)\right) \geq \overline{E}_{\mathcal{P}^+|^{\chi}F}(\theta).$$

Recall that Epstein and Schneider proved that rectangularity is a condition that guarantees no preference reversal in the case of MMEU [Epstein and Schneider 2003], and Hayashi proved a similar result for MER [Hayashi 2009]. With MMEU and MER, only unweighted probabilities are considered. Definition 4.5.4 essentially gives the generalization of Epstein and Schneider's condition to weighted probabilities. Part (a) of χ-rectangularity is analogous to the rectangularity condition of Epstein and Schneider. Part (b) of χ-rectangularity corresponds to the assumption that $(E \cap F)$ is non-null, which is analogous to Axiom 5 in Epstein and Schneider's axiomatization. Finally, part (c) of χ-rectangularity holds for MMEU when weights are in $\{0, 1\}$, and thus is not necessary for Epstein and Schneider. It is not hard to show that we can replace condition (a) above by the requirement that $\mathcal{P}^+$ is closed under conditioning, in the sense that if $(\Pr, \alpha) \in \mathcal{P}^+$, then so are $(\Pr|(E \cap F), \alpha)$ and $(\Pr|(E^c \cap F), \alpha)$.

As the following result shows, χ-rectangularity is indeed sufficient to give us Axioms 4.1–4.4 under prior-by-prior updating and likelihood updating.

Theorem 4.5.5. *If $C(\mathcal{P}^+)$ is closed and convex, then Axiom 4.1 holds for the family of choices $C_M^{reg,\mathcal{P}^+|_\chi E}$ if and only if $\mathcal{P}^+$ is χ-rectangular.*

The proof that χ-rectangularity implies Axiom 4.1 requires only that $C(\mathcal{P}^+)$ be closed (i.e., convexity is not required). Hayashi [2011] proves an analogue of Theorem 4.5.5 for MER using prior-by-prior updating. He also essentially assumes that the menu includes forgone opportunities, but his interpretation of forgone opportunities is quite different from ours. He also shows that if forgone opportunities are not included in the menu, then the set of probabilities representing the DM's uncertainty at all but the initial time must be a singleton. This implies that the DM must behave like a Bayesian at all but the initial time, since MER acts like expected utility maximization if the DM's uncertainty is described by a single probability measure.

Epstein and Le Breton [1993] took this direction even further and prove that, if a few axioms hold, then only Bayesian beliefs can be dynamically consistent. While Epstein and Le Breton's result was stated in a menu-free setting, if we use a constant menu throughout the decision problem, then our model fits into their framework. At first glance, their impossibility result may seem to contradict our sufficient conditions for no preference reversal. However, Epstein and Le Breton's impossibility result does not apply because one of their axioms, $P4^c$, does not hold for MER (or MWER). For ease of exposition, we give $P4^c$ for static decision problems. Given acts f and g and a set T of states, let fTg be the act that agrees with f on T and agrees with g on T^c. Given an outcome x, let x^* be the constant act that gives outcome x at all states.

Axiom 4.5 (Conditional weak comparative probability). *For all events T, A, B, with $A \cup B \subseteq T$, outcomes w, x, y, and z, and acts g, if $w^*Tg \succ x^*Tg$, $z^*Tg \succ y^*Tg$, and $(w^*Ax^*)Tg \succeq (w^*Bx^*)Tg$, then $(z^*Ay^*)Tg \succeq (z^*By^*)Tg$.*

$P4^c$ implies Savage's $P4$, and does not hold for MER and MWER in general. For a simple counterexample, let $S = \{s_1, s_2, s_3\}$, $X = \{o_1, o_5, o_7, o_{10}, o_{20}, o_{23}\}$, $A = \{s_1\}$, $B = \{s_2\}$, $T = A \cup B$, $u(o_k) = k$, g is the act such that $g(s_1) = o_{20}$, $g(s_2) = o_{23}$, and $g(s_3) = o_5$. Let $\mathcal{P} = \{p_1, p_2, p_3\}$, where

- $p_1(s_1) = 0.25$ and $p_1(s_2) = 0.75$;
- $p_2(s_3) = 1$;
- $p_3(s_1) = 0.25$ and $p_3(s_3) = 0.75$.

Let the menu $M = \{o_1^*, o_7^*, o_{10}^*, o_{20}^*, g\}$. Let $\succeq$ be the preference relation determined by MER. The regret of o_{10}^*Tg is 15 (this is the regret with respect to p_2), and the regret of o_7^*Tg is 15.25 (the regret with respect to p_1), therefore $o_{10}^*Tg \succ o_7^*Tg$. It is also easy to see that the regret of o_{20}^*Tg is 15 (the regret with respect to p_2), and the regret of o_1^*Tg is 21.25 (the regret with respect to p_1), so $o_{20}^*Tg \succ o_1^*Tg$. Moreover, the regret of $(o_{10}^*Ao_7^*)Tg$ is 15 (the regret with respect to p_2), and the regret of $(o_{10}^*Bo_1^*)Tg$ is 15 (the regret with respect to p_2), so $(o_{10}^*Ao_7^*)Tg \succeq (o_{10}^*Bo_1^*)Tg$. However, the regret of $(o_{20}^*Ao_1^*)Tg$ is 16.5 (the regret with respect to p_1), and the regret of $(o_{20}^*Bo_1^*)Tg$ is 16 (the regret with with respect to p_3), therefore $(o_{20}^*Ao_1^*)Tg \not\succeq (o_{20}^*Bo_1^*)Tg$. Thus, Axiom 4.5 does not hold (taking $y = o_1, x = o_7, w = o_{10}, z = o_{20}$).

Siniscalchi [2011, Proposition 1] proves that his notion of dynamically consistent conditional preference systems must essentially have beliefs that are updated by Baysian updating. However, his result does not apply in our case either, because it assumes consequentialism: that the conditional preference system treats identical subtrees equally, independent of the greater decision tree within which the subtrees belong. This does not happen if, for example, we take forgone opportunities into account.

There may be reasons to exclude forgone opportunities from the menu. *Consequentialism*, according to Machina [1989], is 'snipping' the decision tree at the current choice node, throwing the rest of the tree away, and calculating preferences at the current choice node by applying the original preference ordering to alternative possible continuations of the tree. With this interpretation, consequentialism implies that forgone opportunities should be removed from the menu.

Similarly, there many be reasons to exclude unachievable plans from the menu. Preferences computed with unachievable plans removed from the menu would be independent of these unachievable plans. This quality might make the preferences suitable for iterated elimination of suboptimal plans as a way of finding the optimal plan. In certain settings, it may be difficult to rank plans or find the most preferred plan among a large menu. For instance, consider the problem of deciding on a career path. In these settings, it may be relatively easy to identify bad plans, the elimination of which simplifies the problem. Conversely, computational benefits may motivate a DM to ignore unachievable plans. That is, a DM may choose to ignore unachievable plans because doing so simplifies the search for the preferred solution.

4.6 Chapter Conclusion

In dynamic decision problems, it is not clear which menu should be used to compute regret. However, if we use MWER with likelihood updating, then in order to avoid preference reversals, we need to include all initially feasible plans in the menu, as well as richness conditions on the beliefs. Another, well-studied approach to circumvent preference reversals is *sophistication*. A sophisticated agent is aware of the potential for preference reversals, and thus uses backward induction to determine the *achievable plans*, which are the plans that can actually be carried out. In the procrastination example, a sophisticated agent would know that she would not study the second day. Therefore, she knows that playing on the first day and then studying on the second day is an unachievable plan.

Siniscalchi [2011] considers a specific type of sophistication, called *consistent planning*, based on earlier definitions of Strotz [1955] and Gul and Pesendorfer [2005]. Assuming a filtration information structure, Siniscalchi axiomatizes behavior resulting from consistent planning using any menu-independent decision rule.[1] With a menu-dependent decision rule, we need to consider the choice of menu when using consistent planning. Hayashi [2009] axiomatizes sophistication using regret-based choices, including MER and the smooth model of anticipated regret, under the fixed filtration information setting. However, in his models of regret, Hayashi assumes that the menu that the DM uses to compute regret includes only the achievable plans. In other words, forgone opportunities and those plans that are not achievable are excluded from the menu. It would be

[1]Siniscalchi considers a more general information structure where the information that the DM receives can depend on her actions in an unpublished version of his paper [Siniscalchi 2006].

interesting to investigate the effect of including such in the menus of a sophisticated DM. A sophisticated DM who takes unachievable plans into account when computing regret can be understood as being "sophisticated enough" to understand that her preferences may change in the future, but not sophisticated enough to completely ignore the plans that she cannot force herself to commit to when computing regret. On the other hand, a sophisticated DM who ignores unachievable plans does not feel regret for not being able to commit to certain plans.

Finally, we have only considered "binary" menus in the sense that an act is either in the menu and affects regret computation, or it is not. A possible generalization is to give different weights to the acts in the menu, and multiply the regrets computed with respect to each act by the weight of the act. For example, with respect to forgone opportunities, "recently forgone" opportunities may warrant a higher weight than opportunities that have been forgone many timesteps ago. Such treatment of forgone opportunities will definitely affect the behavior of the DM.

CHAPTER 5

MAXIMIN SAFETY

5.1 Overview

We present a new decision rule, *maximin safety*, that seeks to maintain a large margin from the worst outcome, in much the same way minimax regret seeks to minimize distance from the best. We argue that maximin safety is valuable both descriptively and normatively. Descriptively, maximin safety explains the well-known *decoy effect*, in which the introduction of a dominated option changes preferences among the other options. Normatively, we provide an axiomatization that characterizes preferences induced by maximin safety, and show that maximin safety shares much of the same behavioral basis with minimax regret.

5.2 Introduction

Representing uncertainty using a probability distribution, and making decisions by maximizing expected utility, is widely accepted, founded on formal mathematical principles, and satisfies intuitive notions of rationality such as independence of irrelevant alternatives and the sure thing principle [**?**]. However, enforcing seemingly appealing concepts of rationality can ultimately lead to decisions inconsistent with what real humans consider reasonable. For example, observed behavior under unquantified (Knightian [Knight 1921]/ strict [Larbi, Konieczny, and Marquis 2010]) uncertainty, such as that in the Ellsberg paradox [Ellsberg 1961], demonstrates how appealing concepts of rationality can lead to inconsistency with human choices. Alternative decision rules, such as maximin utility [Wald 1950] and minimax regret [Savage 1954; Luce and Raiffa 1957] provide rationally plausible decisions in ambiguous situations and can be used to resolve such paradoxes, but still fail to explain some human behavioral patterns. A particularly illustrative example of such behavior is called the *decoy effect* [Huber, Payne, and Puto 1982], in which the introduction of a *dominated* option changes the preference among the *undominated* ones. While the decoy effect has been investigated in the psychology [Doyle, O'Connor, Reynolds, and Bottomley 1999; Wedell 1991] and economics literature [Bateman, Munro, and Poe 2008; Simonson and Tversky 1992], we are unaware of any axiomatic treatment of it. To address this, we introduce a criterion called *safety* as the basis for a *maximin safety* decision rule.[1] Safety serves as a dual to regret that quantifies distance from a worst outcome, much as regret quantifies proximity to a best outcome. Maximin safety also satisfies familiar properties common to maximin utility and minimax regret, and hence also resolves the Ellsberg paradox. Moreover, maximin safety accommodates observed preferences that are incompatible with minimax regret and maximin utility. We demonstrate how safety-seeking behavior can produce the decoy effect, and show how maximin safety can explain it. We also extend Stoye's [2011b] axiomatizations of standard decision rules to include maximin safety, thus allowing a comparison between maximin safety and state-of-the-art decision rules.

[1]This decision rule has been mentioned in passing, inside a proof by Hayashi [Hayashi 2008b], where it was referred to as 'maximin joy'. We use the term 'safety' rather than 'joy' to avoid confusion with the concept called 'joy of winning' in [Hayashi 2008b].

5.2.1 Relative Preferences and Regret

It is not hard to imagine situations in which performance *relative* to other possible outcomes is more important than *absolute* performance. Consider, for example, a group of duck hunters surprised by a hungry bear [Crawley 1993; Chao, Hanley, Burch, Dahlberg, and Turner 2000]. The hunters all attempt to escape by running in the same direction while the slowest one despairs: "this is hopeless, we can never outrun the bear." The hunter in front of him snickers, "I don't need to outrun the bear, I just need to outrun *you*." Whether the prospect is being picked from a group of peers for a date [Ariely 2008], winning a gold medal, or obtaining an 'A' in a class, success is often measured by relative performance, rather than by an absolute standard. One such preference for relative performance is embodied in the well-known decision theoretical concept of *regret* [Savage 1954; Luce and Raiffa 1957]. While psychological literature on regret focuses on the bad feelings that occur *after* a choice leads to an inferior outcome, some also considers that anticipation of such negative emotions may influence the choice itself [Simonson and Tversky 1992; Larrick and Boles 1995; Ritov 1996].

As before, we assume that uncertainty is captured by a set of possible worlds, one of which is the true state of the world. Regret is a measure of distance between the value of a considered outcome and the value of the best possible outcome, under a given state. This leads to an important property that is always true for regret – the introduction of a *dominated* option does not change the regrets of the existing options. We will refer to this property as *independence of dominated alternatives* (IDA). Those who believe in regret avoidance may think that this property is perfectly reasonable. For example, suppose you have a \$10 bill and you can either buy a \$10 lottery ticket, or two \$5 lottery tickets. Most would agree that your choice should not be affected by a dominated third option, "burning the \$10 bill". Other standard decision rules, such as expected utility maximization, have even stronger independence guarantees. The ranking of two choices under expected utility maximization is *menu-independent*, i.e., completely independent of the set of feasible choices (the *menu*). Menu-independence implies IDA. In contrast, regret-based preferences are menu-dependent, but since they conform to IDA, they are not compatible with observed biases sensitive to dominated options [Ariely 2008]. While IDA seems intuitively appealing, there is a great deal of empirical evidence that human preferences are indeed affected by dominated options in measurable and sometimes profound ways.

5.2.2 The Decoy Effect in Decision Theory

Suppose you are offered \$6 in cash, and the option of trading it for a Cross pen. The pen is nice, but you have plenty of pens, so decide to keep the cash. Right before you walk away, you are offered an alternative pen in exchange for the \$6. You see the new pen and find it hideous. A smile comes to your face as you turn around and say, "you know, I'll take that original Cross pen after all."

This story dramatizes an actual experiment [Huber, Payne, and Puto 1982]. When the first choice was offered to 106 people, 64% took the cash, 36% took the pen. When the second pen was added to the offer to 115 other subjects, 52% took the cash, 46% took the Cross pen, and 2% took the decoy. Generally, a *decoy* is an option that is designed to be inferior to another option in every way (i.e., it is a dominated option). Despite the intuitive appeal of IDA, the presence of a dominated option drove selection of the Cross pen from 36% to 46%. In this chapter, we focus on a particular class of *decoy effect*, called *asymmetric dominance*, which occurs when the decoy is dominated by one existing alternative, but not by another. Empirical studies show that the decoy is rarely chosen, but its addition to a set of choices consistently drives DMs toward the *dominating choice*.

Numerous empirical studies have also shown decoy effects in class action settlements [Zimmerman 2010], recreational land management [Bateman, Munro, and Poe 2008], choice of healthcare plans and political candidates [Hedgcock, Rao, and Chen 1999], purchase of consumer goods such as cameras and personal computers [Simonson and Tversky 1992], restaurant choices [Huber, Payne, and Puto 1982], and even romantic attraction [Ariely 2008]. Surprisingly, a decoy effect can occur even if the decoy is not actually an option, but merely a recent memory of an option (a phantom decoy [Farquhar and Pratkanis 1993; Doyle, O'Connor, Reynolds, and Bottomley 1999]). Furthermore, the decoy effect is not limited to humans, but is also observed in honeybees and grey jays [Shafir, Waite, and Smith 2002].

In an attempt to explain the decoy effect, experts in the behavioral sciences have offered a variety of domain-specific analyses, including "perceptual framing" [Huber, Payne, and Puto 1982], "value-shift" [Wedell 1991], "extremeness aversion" [Simonson and Tversky 1992], and "contrast bias" [Simonson and Tversky 1992; Zimmerman 2010]. All of these explanations focus on valuing the discrepancy between the decoy and the dominating alternative. Intuitively, this provides a compelling example of preferring the margin of safety from the worst outcome. As we are not aware of any formalization in decision theory

	Wet Road	Dry Road
Sprint	1	9
Hustle	3	6
Jog	2	2

Table 5.1: Hunters running from a bear.

that is consistent with the intuitive preference for “margin of safety”, we offer one here.

To illustrate our new decision rule, recall the example of the unfortunate duck hunters. As they run from the bear, they approach a blind curve and have no idea what is around it: it could be wet or dry. If it is dry they will cover the most ground if they try to run faster, however if it is wet (thus slippery) they will be better off if they slow down and maintain balance. The options and the distance traveled under each circumstance are summarized in Table 5.1. In general, exerting excessive effort on a wet road leads to slipping and less distance covered; exerting effort on a dry road leads to more distance covered.

If the probability of the road conditions is unknown, and only the first two options are available (sprint and hustle), there is no intuitively preferred choice and we may assume there are enough hunters such that at least one will pick each option. However, if we add a new option, jog, something interesting happens. As jog is dominated by hustle, IDA requires that its availability should not change the preferences among the other options. However, regardless of whether the road is wet or dry, hustle is never the worst alternative: if the road is wet, hustle (3) is faster than sprint (1), and if the road is dry, hustle (6) is faster than jog (2). In either case, selecting hustle prevents the hunter from being the slowest and getting caught by the bear.

	s_1 Safari	s_2 World Cup
a_1: Travel	4	4
a_2: Sports	2	6
a_3: Decoy	3	3

Table 5.2: Utilities in the camera purchase example.

While it may be callous, it seems perfectly reasonable for a hunter to decide to run just fast enough to make sure there is someone behind him. In other words, the most sensible decision might be to run just fast enough to guarantee the maximum possible margin between himself and the slowest runner, in the worst scenario. This margin between the hunter and his slowest compatriot can be considered a measure of $safety$, which is at the heart of this chapter.

The rest of the chapter proceeds as follows. Section 5.3 provides a formalization of the decoy paradox along with basic decision-theoretical notation. Section 5.4 describes the relationship between minimax regret and *maximin safety* and shows how maximin safety resolves the decoy paradox. Section 5.5 provide an axiomatic characterization of maximin safety. Section 5.6 suggests a unification of utility, regret, and safety using *anchoring functions*, and also considers a generalization to qualitative relative preferences.

5.3 The Formal Framework

Just like regret, safety can depend on the menu. We will only consider finite menus, from which randomized strategies can be chosen.

Consider the problem of a DM contemplating a camera purchase, summarized in Table 5.2. The DM has a choice between buying a rugged travel camera

(a_1) that takes decent pictures in a wide variety of circumstances, and buying a delicate sports camera with higher speed and image quality (a_2). Each state characterizes the possible situations that a purchaser may experience during the useful life of the camera (Will the DM experience harsh conditions? Or win tickets to the World Cup?) The utility $U(a, s)$ of act a under state s represents an abstract net value to the DM if the true world is state s.

If the DM ends up going on a safari (s_1), then act a_1 results in moderate quality pictures of exciting wildlife ($U(a_1, s_1) = 4$), but act a_2 results in a few exquisite shots and many missed opportunities ($U(a_2, s_1) = 2$). On the other hand, if the DM goes to the World Cup (s_2), then act a_2 results in many great pictures in a safe environment ($U(a_2, s_2) = 6$), while act a_1 provides only moderate quality pictures ($U(a_1, s_2) = 4$).

If the DM can assign probabilities $P(s_1)$ and $P(s_2)$ to the states, she can calculate an expected utility $E[U(a_i)] = \sum_{s \in S} P(s)U(a_i, s)$, and simply select the act that maximizes expected utility. However, if the state probabilities are unavailable, we have unquantified uncertainty. In such cases, the DM must find another method for aggregating the utility of each act across states in order to assign a *value* to each camera. Here we will focus on the methods of maximax utility, minimax utility, and minimax regret. To understand minimax regret, we need to define the notion of regret. For a menu M and act $a \in M$, the regret of a with respect to M and decision problem (S, X, U) is

$$\max_{s \in S}(\max_{a' \in M} U(a', s) - U(a, s)).$$

We denote this as $reg_M^{(S,X,U)}(a)$, and usually omit the superscript (S, X, U).

When comparing decision rules, it is often convenient to define a *value function* that assigns a numeric value to each act, for the purpose of ranking the acts.

Decision Rule	Value of an act a	Decision rule description	Best
maximax utility	$V(a) = \max_{s \in S} U(a, s)$	Optimize the best-case outcome.	a_2
maximin utility	$V(a) = \min_{s \in S} U(a, s)$	Optimize the worst-case outcome.	a_1
minimax regret	$V(a, M) = -reg_M(a)$	Pick an act to minimize the worst-case distance from the best outcome.	a_1, a_2

Table 5.3: Standard decision rules and most valued acts in the camera example.

Formally, for a decision problem (S, X, U), a value function is a function

$$V^{(S,X,U)}(a, M) : X^S \times 2^A \to \mathbb{R}.$$

We will usually omit the superscript (S, X, U) and just write $V(a, M)$, or $V(a)$ if the value function is menu-independent.

We say that the value function V *represents* the family of preference relations $\succ_{V,M}$, if for all menus M and all $a, a' \in M$,

$$a \succ_{V,M} a' \Leftrightarrow V(a, M) > V(a', M).$$

In other words, act a is (strictly) preferred to act a' with respect to menu M if and only if $V(a, M) > V(a', M)$. The value functions and preferences of several standard decision rules are given in Table 5.3.

Now, perhaps the camera vendor would like to sell more travel cameras, so the vender puts an obsolete travel camera a_3 next to a_1 as a *decoy*. Camera a_3 has the same price as a_1, but fewer features and lower picture quality. The vendor hopes to make a_1 more appealing by contrast with a_3. Table 5.2 illustrates the decision problem when a_3 is added to the menu. The ranking between a_1 and a_2 according to each of the decision rules in Table 5.3 is unaffected by the introduction of a_3 to the menu. The addition of a_3 also illustrates the concept of dominance. We say that an act a *dominates* a', if for all $s \in S$, $U(a,s) > U(a',s)$.

5.4 Maximin Safety

While minimax regret seeks to minimize separation from best outcomes, *maximin safety* is a conceptual dual that seeks to maximize separation from the worst outcomes. For a menu M and act $a \in M$, the safety of a in state s is defined as:

$$safety_M^{(S,X,U)}(a,s) = U(a,s) - \min_{a' \in M}(U(a',s)),$$

and in keeping with the convention for regret, the safety of an act is defined as:

$$safety_M^{(S,X,U)}(a) = \min_{s \in S}(safety_M^{(S,X,U)}(a,s)).$$

We will often omit the superscript (S,X,U).

The family of maximin safety preferences $\succ_{saf,M}$ represented by the $safety$ value function satisfies, for all M and $a, a' \in M$,

$$a \succ_{saf,M} a' \Leftrightarrow safety_M(a) > safety_M(a').$$

	Utility		Safety (no decoy)		Safety (w. decoy)	
	s_1	s_2	s_1	s_2	s_1	s_2
a_1: travel	4	4	2	0	2	1
a_2: sports	2	6	0	2	0	3
a_3: decoy	3	3			1	0

Table 5.4: Camera purchase with and without decoy.

Now we reconsider the camera example using safety (Table 5.4). Without the decoy, both acts have the same safety of 0, since each act has the lowest utility

	Utility		Regret		Safety		Optimal for
	s_1	s_2	s_1	s_2	s_1	s_2	
a_1	1	9	3	0	0	5	maximax utility
a_2	3	6	1	3	2	2	maximin safety
a_3	2	7	2	2	1	3	minimax regret
a_4	4	4	0	5	3	0	maximin utility

Table 5.5: Different decision rules select different acts for the same problem.

in some state; so there is no clear safety preference. However, when the decoy is present, the act a_1 never has the lowest utility at any state, and thus it has a strictly positive safety. In this case, $safety_{\{a_1,a_2,a_3\}}(a_1)$ is the unique maximum among the acts $\{a_1, a_2, a_3\}$, and therefore a_1 is the preferred choice. The relative increase in the safety of an act due to the addition of the dominated act is an essential element in solving the decoy paradox. Intuitively, this may correspond to a sense that even if a particular act gets low utility in the realized state, the DM may think that "I'm better off than the fools who bought the worse camera", or in a more positive light, "I must be getting a steal with this better camera for the same price". In competitive survival games (such as the reality game show Survivor), the notion of maximizing safety may also embody a preference to maintain a maximal distance from the lowest performer, which reduces the chance of elimination. Table 5.5 compactly demonstrates how choices based on maximin safety differs from the other standard decision rules.

In the camera example, the addition of a decoy created a strict preference between two acts that were initially tied. The introduction of a dominated act can actually *reverse* preferences between acts. Table 5.6 shows a menu of three acts: $M = \{a_1, a_2, a_3\}$. Act a_2 has a minimum safety of 1, while both a_1 and a_3 have the lowest utility for some state, so each has minimum safety of 0. Consequently, a_2 is the most preferred choice under the safety preference. When a

	Utility			Safety(no decoy)			Safety (w. decoy $a4$)		
	s_1	s_2	s_3	s_1	s_2	s_3	s_1	s_2	s_3
a_1	9	2	6	5	0	0	8	0	0
a_2	5	3	7	1	1	1	4	1	1
a_3	4	8	8	0	6	2	3	6	2
a_4	1	5	6				0	3	0

Table 5.6: Without a_4, $M = \{a_1, a_2, a_3\}$, and $a_2 \succ_{saf,M} a_3$. Adding a_4 (dominated by a_3) reverses the maximin safety preference between a_2 and a_3.

new choice a_4 is added, act a_4 is dominated by a_3, but it has higher utility than the other acts in some states. This situation is known as *asymmetric dominance*, which is typically associated with decoy effects. In this example, asymmetric dominance guarantees that a_3 is never one of the worst choices, and thus has a strictly positive safety value. In other words, the addition of a_4 to the menu M does not affect the safety of a_1 or a_2, but increases the safety of a_3 to make $a_3 \succ_{saf,M\cup\{a_4\}} a_2$.

5.5 Axiomatic Analysis

To provide an axiomatic characterization of maximin safety, we employ the standard *Anscombe-Aumann* (AA) framework [Anscombe and Aumann 1963], where outcomes are restricted to lotteries. Maximin safety is characterized by modifying one of the axioms in an existing characterization of minimax regret provided by Stoye [2011b].

Given a set Y of *prizes*, a *lottery* over Y is just a probability with finite support on Y. As in the AA framework, we let the set of *outcomes* be $\Delta(Y)$, the set of all lotteries over Y. Thus, *acts* are functions from S to $\Delta(Y)$. We can think of a lottery as modeling objective, quantified uncertainty, while the states model unquantified uncertainty. The technical advantage of considering such a set of outcomes is that we can consider convex combinations of acts. If f and g are acts, define the act $\alpha f + (1-\alpha)g$ to be the act that maps a state s to the lottery $\alpha f(s) + (1-\alpha)g(s)$. For simplicity, we follow Stoye [Stoye 2011b] and restrict to menus that are the convex hull of a finite number of acts, so that if f and g are acts in M, then so is $pf + (1-p)g$ for all $p \in [0,1]$.

In this setting, we assume that there is a utility function U on prizes in Y, and that there are at least two prizes y_1 and y_2 in Y, with different utilities. Note that $l(y)$ is the probability of getting prize y if lottery l is played. We will use l^* to denote a constant act that maps all states to l. The utility of a lottery l is just the expected utility of the prizes obtained, that is, $u(l) = \sum_{\{y \in Y\,:\, l(y)>0\}} l(y)U(y)$. The expected utility of an act f with respect to a probability $\Pr$ is then just $u(f) = \sum_{s \in S} \Pr(s)u(f(s))$, as usual. Given a set Y of prizes, a utility U on the prizes, and a state space S, we have a family $\succeq^{S,\Delta(Y),u}_{saf,M}$ of preference orders on acts determined by maximin safety, where u is the utility function on lotteries as determined by U.[2] For convenience, from here on we will write $\succeq^{S,Y,U}_{M}$ rather than $\succeq^{S,\Delta(Y),u}_{saf,M}$. We will state the axioms in a way such that they can be compared to standard axioms and those for minimax regret in [Stoye 2011b]. The axioms are universally quantified over acts f, g, and h, menus M and M', and $p \in (0,1)$. Whenever we write $f \succeq_M g$ we assume that $f, g \in M$.

Axiom 5.1. *(Monotonicity)* $f \succeq_M g$ *if* $(f(s))^* \succeq_{\{(f(s))^*,(g(s))^*\}} (g(s))^*, \forall s \in S$.

Axiom 5.2. *(Completeness)* $f \succeq_M g$ *or* $g \succeq_M f$.

Axiom 5.3. *(Nontriviality)* $f \succ_M g$ *for some acts* f *and* g *and menu* M.

[2]We let $f \succeq^{S,\Delta(Y),u}_{saf,M} g$ iff $g \not\succ^{S,\Delta(Y),u}_{saf,M} f$, and $f \sim_M g$ iff $f \succeq_M g$ and $g \succeq_M f$.

Axiom 5.4. *(Mixture Continuity) If* $f \succ_M g \succ_M h$*, then there exists* $q, r \in (0,1)$ *such that* $qf + (1-q)h \succ_M g \succ_M rf + (1-r)h$.

Axiom 5.5. *(Transitivity)* $f \succeq_M g \succeq_M h \Rightarrow f \succeq_M h$.

Menu-independent versions of Axioms 5.1 to 5.5 are standard in other axiomatizations, and in particular hold for maximin utility. Axiom 5.3 is used in the standard axiomatizations to get a nonconstant utility function in the representation. While maximin safety does not satisfy menu-independence, it does satisfy menu-independence when restricted to menus consisting of only constant acts. This property is captured by the following axiom.

Axiom 5.6. *(Menu independence for constant acts) If* l^* *and* $(l')^*$ *are constant acts, then* $l^* \succeq_M (l')^*$ *iff* $l^* \succeq_{M'} (l')^*$.

We also have a menu-dependent version of the von Neumann-Morgenstern (VNM) Independence axiom. Like the VNM Independence axiom, Axiom 5.7 says that ranking between two acts does not change when both acts are mixed with a third act; but unlike VNM Independence, the menu used to compare the original acts in Axiom 5.7 is different from that used to compare the mixtures. Axiom 5.7 holds for minimax regret and maximin safety, but not for maximin utility.

Axiom 5.7. *(Independence)* $f \succeq_M g \Leftrightarrow pf + (1-p)h \succeq_{pM+(1-p)h} pg + (1-p)h$.

Axiom 5.8. *(Symmetry) For a menu* M*, suppose that* $E, F \in 2^S \backslash \{\emptyset\}$ *are disjoint events such that for all* $f \in M$*,* f *is constant on* E *and on* F*. Define* f' *by*

$$f'(s) = \begin{cases} f(s') \text{ for some } s' \in E, \text{ if } s \in F \\ f(s') \text{ for some } s' \in F, \text{ if } s \in E \\ f(s) \text{ otherwise} \end{cases}$$

Let M' be the menu generated by replacing every act $f \in M$ with f'. Then

$$f \succeq_M g \Leftrightarrow f' \succeq_{M'} g'$$

Symmetry, which is one of the characterizing axioms for minimax regret in [Stoye 2011b], captures the intuition that no state can be considered more or less likely than another. Therefore Symmetry helps distinguish the probability-free decision rules maximin utility, minimax regret, and maximin safety, from their probabilistic counterparts [Gilboa and Schmeidler 1989b; Stoye 2011a].

Axiom 5.9. *(Ambiguity Aversion) $f \sim_M g \Rightarrow pf + (1-p)g \succeq_M g$.*

Axiom 5.9 says that the DM weakly prefers to hedge her bets. Axioms 1-5.9 are all part of the characterization in [Stoye 2011b] of minimax regret (which consists of Axioms 1-5.9 and Symmetry). Axioms 1-5 and 5.9 are also sound for the maximin decision rule [Stoye 2011b].

In [Stoye 2011b], one of the axioms characterizing minimax regret is Independence of Never Strictly Optimal alternatives (INA), which states that adding or removing acts that are not strictly potentially optimal in the menu does not affect the ordering of acts. [3] By varying this INA axiom, we obtain a characterization for maximin safety. We say that an act a is *never strictly worst relative to M* if, for all states $s \in S$, there is some $a' \in M$ such that $a(s) \succeq a'(s)$.

Axiom 5.10. *(Independence of Never Strictly Worst Alternatives (INWA)) If an act a is never strictly worst relative to M, then $f \succeq_M g$* iff *$f \succeq_{M \cup \{a\}} g$.*

Although adding acts to the menu, in general, can affect minimax regret preferences, INA implies the Independence of Dominated Alternatives property

[3] An act h is *never strictly optimal relative to M* if, for all states $s \in S$, there is some $f \in M$ such that $(f(s))^* \succeq (h(s))^*$.

that we used earlier when discussing the decoy effect. Thus, INA guarantees that minimax regret can never be compatible with the decoy effect.

Theorem 5.5.1. *For all Y, U, S, the family of maximin safety preference orders $\succeq^{S,Y,U}_{saf,M}$ induced by a decision problem $(S, \Delta(Y), u)$ satisfies Axioms 5.1–5.10. Conversely, if the family of preference orders $\succeq_M$ on the acts in $\Delta(Y)^S$ satisfies Axioms 5.1–5.10, then there exists a utility function U on Y that determines a utility u on $\Delta(Y)$ such that $\succeq_M = \succeq^{S,Y,u}_{saf,M}$. Moreover, U is unique up to affine transformations.*

Proof. The soundness of the axioms are straightforwardly verified, so we show only the completeness of the axioms. We will use the same general sequence of arguments that Stoye uses in [Stoye 2011b]. First, we establish a nonconstant utility function U, where constant acts are ranked by their expected utilities. Since we have the standard axioms ($1-5$), we get U from standard arguments, and it is unique up to affine transformations. Next, we observe the following lemma:

Lemma 5.5.2. *Suppose the family $\succeq_M$ satisfies Axioms 1-10, and $\succeq_{M^+}$ is representable by maximin safety, where M^+ is the menu of all acts with nonnegative utilities. Then the family $\succeq_M$ is representable by maximin safety.*

Lemma 5.5.2 follows from an argument analogous to that for regret in [Stoye 2011b]. The next step is to establish that the axioms on $\succeq_M$ restrict $\succeq_{M^+}$ to satisfy the axioms of ambiguity aversion, monotonicity, completeness, transitivity, non-triviality, and symmetry. It is a straightforward verification that will not be reproduced here. Theorem 1 (iii) of [Stoye 2011b] then implies that $\succeq_{M^+}$ is the maximin utility ordering. Next, let g_M be an act such that $u \circ g_M(s) = -\min_{h \in M} u(h, s)$, so that we have

$$
\begin{aligned}
f \succeq_M g \quad &\Leftrightarrow \tfrac{1}{2}f + \tfrac{1}{2}g_M \succeq_{M^+} \tfrac{1}{2}g + \tfrac{1}{2}g_M \\
&\Leftrightarrow \min_{s\in S} u(\tfrac{1}{2}f + \tfrac{1}{2}g_M, s) \geq \min_{s\in S} u(\tfrac{1}{2}g + \tfrac{1}{2}g_M, s) \\
&\Leftrightarrow \min_{s\in S}(\frac{1}{2}(u(f,s) - \min_{h\in M} u(h,s))) \geq \min_{s\in S}(\frac{1}{2}(u(g,s) - \min_{h\in M} u(h,s))).
\end{aligned}
$$

□ □

The characterizing axioms serve as a justification for maximin safety in the sense that behaving as a safety maximizer is equivalent to accepting the axioms. Axioms 1-7 are standard and broadly accepted to be reasonable, while symmetry and ambiguity aversion are implied by both maximin utility and minimax regret. Whether the INA axiom (for regret) or the INWA axiom (for safety) is more reasonable would depend on the individual and the nature of the decision problem. Thus, we believe that the reasonableness of the maximin safety decision rule is comparable to that of minimax regret.

Individual necessity of the axioms can be established, as is commonly done [Hayashi 2008b; Stoye 2011b], by giving examples of preferences that satisfy all the axioms except for the one whose necessity is being shown. For the axioms shared with minimax regret, the same examples found in [Stoye 2011b] shows their individual necessity. For the INWA axiom, the required example is minimax regret. Indeed, a decision rule equivalent to maximin safety was used by Hayashi [Hayashi 2008b] as an example to justify minimax regret's entailment of INA.

Clearly, just as minimax regret is readily generalized to *minimax expected regret* when uncertainty is represented by a set of probability distributions over the state space, maximin safety can be readily extended to *maximin expected safety* in the same manner. As one would expect, given an axiomatization of minimax expected regret [Stoye 2011a], the modification of the INA axiom to INWA results in an axiomatization for maximin *expected* safety.

5.6 Discussion, Generalizations, and Future Work

Both minimax regret and maximin safety embody preferences based on *relative*, rather than *absolute* utility. In Table 5.5, the act preferred by safety has a lower minimum utility than the act preferred by maximin utility, just as the act picked by minimax regret neglects a higher maximum utility in order to minimize the *margin* to each state's maximum utility. The shared preference for relative over absolute performance is reflected in a striking similarity in the structure of the value functions for regret and safety. In comparison, minimax regret can be expressed for all acts a, b as:

$$a \succ_{reg,M} b \text{ iff } \quad \min_{s \in S}(U(a,s) - \max_{a' \in M} U(a',s)) > \min_{s \in S}(U(b,s) - \max_{a' \in M} U(a',s)).$$

Similarly, maximin safety is represented for all acts a, b as

$$a \succ_{saf,M} b \text{ iff } \quad \min_{s \in S}(U(a,s) - \min_{a' \in M} U(a',s)) > \min_{s \in S}(U(b,s) - \min_{a' \in M} U(a',s)).$$

The structural resemblance suggests a common form for the value function. By defining a menu-dependent *anchoring function* $t : S \times 2^A \to \mathbb{R}$, we can represent several previously discussed value functions as:

$$V_t(a, M) = \min_{s \in S} U'(a, s, M, t),$$

where $U'(a, s, M, t) = U(a, s) - t(s, M)$ can be viewed as an *anchored effective utility*. One can see that V_t represents maximin utility if $t(s, M) = 0$; minimax regret if $t(s, M) = \max_{a' \in M} U(a', s)$; and maximin safety if $t(s, M) = \min_{a' \in M} U(a', s)$.

Note that by varying just the anchoring function t, we can obtain all the mentioned decision rules, and more. While we focus only on maximin safety in this chapter, other forms for $t(s, M)$ maximize the positive margin from a state-dependent average, median, or some other characteristic of interest to a DM. For example, college students might seek to conservatively maximize their margin above a desired quantile, in order to achieve a particular grade.

The present work is motivated by behavioral observations of the decoy effect that are typically described in empirical quantities such as distance, price and volume, and thus is most intuitive in a *quantitative* framework. However, the key observation is that safety, like regret, is a notion of relative performance with respect to a set of outcomes, rather than absolute performance. As absolute quantitative utility $U : X \to \mathbb{R}$ can be generalized to a qualitative framework by replacing the U with a mapping $X \to L$ for some ordered set L, relative utility may be made qualitative by considering the mapping with $2^X \times X \to L$. In the case of safety and regret, the particular element of 2^X is the set of all possible outcomes in a state, given a menu of acts. Aggregation of N state-specific orderings into an ordering over acts can be accomplished by an aggregation function $M : L^N \to L$ [Marichal 2001]. This generalization can be readily applied to various characterizations of uncertainty, including probability, plausibility, and the strict uncertainty used in this chapter [Larbi, Konieczny, and Marquis 2010]. While the authors expect that the present quantitative axiomatization can be adapted to a qualitative framework (see, e.g. [Dubois, Fargier, and Perny 2003]), it is beyond the scope of the current chapter.

CHAPTER 6

CONCLUSION

We studied regret-minimization and maxmin expected utility when beliefs are represented by a set of weighted probability distributions. We find that using weighted probability distributions with these decision rules result in well-behaved generalizations of their multiple-prior counterparts. In both cases, allowing for weights on the probability distributions result in the weakening of the axioms characterizing the decisions (i.e., an axiom is dropped). For regret, we also consider the extension of MWER and likelihood updating to dynamic decision problems, and find that both measure-by-measure updating, as well as likelihood updating, can be made dynamically consistent by imposing restrictions on what sets of distributions or weighted distributions are allowed as beliefs. Finally, we study and axiomatize safety-maximization, the “duo” of regret-minimization.

APPENDIX A

PROOFS FOR CHAPTER 2

A.1 Proof of Lemma 2.5.2

A.1.1 Defining a functional on utility acts

Stoye [2011a] also started his proof of a representation theorem for MER by reducing to a single preference order $\succeq_{M^*}$. He then noted that, the expected regret of an act f with respect to a probability $\Pr$ and menu M^* is just the negative of the expected utility of f. Thus, the worst-case expected regret of f with respect to a set $\mathcal{P}$ of probability measures is the negative of the worst-case expected utility of f with respect to $\mathcal{P}$. Thus, it sufficed for Stoye to show that $\succeq_{M^*}$ had an MMEU representation, which he did by showing that $\succeq_{M^*}$ satisfied Gilboa and Schmeidler's [1989a] axioms for MMEU, and then appealing to their representation theorem.

This argument does not quite work for us, because now $\succeq$ does not satisfy the C-independence axiom. (This is because our preference order is based on *weighted* regret, not regret.) However, we can get a representation theorem for weighted regret by using some of the techniques used by Gilboa and Schmeidler to get a representation theorem for MMEU, appropriately modified to deal with lack of C-independence. Specifically, like Gilboa and Schmeidler, we define a functional I on utility acts such that the preference order on utility acts is determined by their value according to I (see Lemma B.3.4). Using I, we can then determine the weight of each probability in $\Delta(S)$, and prove the desired representation theorem.

By standard results, u represents $\succeq$ on constant acts, and $\succeq$ depends only on the utility achieved in each state (as opposed to the actual outcomes) of the acts. The space of all utility acts is the Banach space $\mathcal{B}$ of real-valued functions on S. Let $\mathcal{B}^-$ be the set of nonpositive functions in $\mathcal{B}$, where the function b is nonpositive if $b(s) \leq 0$ for all $s \in S$.

We now define a functional I on utility acts in $\mathcal{B}^-$ such that for all f, g with $b_f, b_g \in \mathcal{B}^-$, we have $I(b_f) \geq I(b_g)$ iff $f \succeq g$. Let

$$R_f = \{\alpha' : l^*_{\alpha'} \succeq f\}.$$

If $0^* \geq b \geq (-1)^*$, then f_b exists, and we define

$$I(b) = \inf(R_{f_b}).$$

For the remaining $b \in \mathcal{B}^-$, we extend I by homogeneity. Let $||b|| = |\min_{s \in S} b(s)|$. Note that if $b \in \mathcal{B}^-$, then $0^* \geq b/||b|| \geq (-1)^*$, so we define

$$I(b) = ||b|| I(b/||b||).$$

Lemma A.1.1. *If $b_f \in \mathcal{B}^-$, then $f \sim l^*_{I(b_f)}$.*

Proof. Suppose that $b_f \in \mathcal{B}^-$ and, by way of contradiction, that $l^*_{I(b_f)} \prec f$. If $f \sim l^*_0$, then it must be the case that $I(b_f) = 0$, since $I(b_f) \leq 0$ by definition of inf, and $f \sim l^*_0 \succ l^*_\epsilon$ for all $\epsilon < 0$ by Lemma A.2.3, so $I(b_f) > \epsilon$ for all $\epsilon < 0$. Therefore, $f \sim l^*_{I(b_f)}$. Otherwise, since $b_f \in \mathcal{B}^-$, by monotonicity, we must have $l^*_0 \succ f$, and thus $l^*_0 \succ f \succ l^*_{I(b_f)}$. By mixture continuity, there is some $q \in (0, 1)$ such that $q \cdot l^*_0 + (1-q) \cdot l^*_{I(b_f)} \sim l_{(1-q)I(b_f)} \prec f$, contradicting the fact that $I(b)$ is the greatest lower bound of R_f.

If, on the other hand, $l^*_{I(b_f)} \succ f$, then $l^*_{I(b_f)} \succ f \succeq l^*_{\underline{c}}$ for some $\underline{c} \in \mathbb{R}$. If $f \sim l^*_{\underline{c}}$ then it must be the case that $I(b_f) = \underline{c}$. $I(b_f) \leq \underline{c}$ since $l^*_{\underline{c}} \succeq l^*_{\underline{c}}$, and $I(b_f) \geq \underline{c}$ since for all $c' < \underline{c}$, $l^*_{c'} \prec f \sim l^*_{\underline{c}}$.

Otherwise, $l^*_{I(b_f)} \succ f \succ l^*_{\underline{c}}$, and by mixture continuity, there is some $q \in (0,1)$ such that $q \cdot l^*_{I(b_f)} + (1-q)l^*_{\underline{c}} \succ f$. Since $qI(b_f) + (1-q)\underline{c} < I(b_f)$, this contradicts the fact that $I(b_f)$ is a lower bound of R_f. Therefore, it must be the case that $l^*_{I(b_f)} \sim f$. □

We can now show that I has the required property.

Lemma A.1.2. *For all acts f, g such that $b_f, b_g \in \mathcal{B}^-$, $f \succeq g$ iff $I(b_f) \geq I(b_g)$.*

Proof. Suppose that $b_f, b_g \in \mathcal{B}^-$. By Lemma B.3.3, $l^*_{I(b_f)} \sim f$ and $g \sim l^*_{I(b_g)}$. Thus, $f \succeq g$ iff $l^*_{I(b_f)} \succeq l^*_{I(b_g)}$, and by Lemma A.2.3, $l^*_{I(b_f)} \succeq l^*_{I(b_g)}$ iff $I(b_f) \geq I(b_g)$. □

In order to invoke a standard separation result for Banach spaces, we extend the definition of I to the Banach space $\mathcal{B}$. We extend I to $\mathcal{B}$ by taking $I(b) = I(b^-)$ for $b \in \mathcal{B} - \mathcal{B}^-$, where for all $b \in \mathcal{B}$, b^- is defined as

$$b^-(s) = \begin{cases} b(s), \text{ if } b(s) \leq 0, \\ 0, \text{ if } b(s) > 0. \end{cases}$$

Clearly $b^- \in \mathcal{B}^-$ and $b = b^-$ if $b \in \mathcal{B}^-$.

We show that the axioms guarantee that I has a number of standard properties. Since we have artificially extended I to $\mathcal{B}$, our arguments require more cases than those in [1989a]. (We remark that such an "artificial" extension seem unavoidable in our setting.) Moreover, we must work harder to get the result that we want. We need different arguments from that for MMEU [1989a], since the preference order induced by MMEU satisfies C-independence, while our preference order does not.

Lemma A.1.3. *(a) If $c \leq 0$, then $I(c^*) = c$.*

(b) I satisfies positive homogeneity: if $b \in \mathcal{B}$ and $c > 0$, then $I(cb) = cI(b)$.

(c) I is monotonic: if $b, b' \in \mathcal{B}$ and $b \geq b'$, then $I(b) \geq I(b')$.

(d) I is continuous: if $b, b_1, b_2, \ldots \in \mathcal{B}$, and $b_n \to b$, then $I(b_n) \to I(b)$.

(e) I is superadditive: if $b, b' \in \mathcal{B}$, then $I(b + b') \geq I(b) + I(b')$.

Proof. For part (a), If c is in the range of u, then it is immediate from the definition of I and Lemma A.2.3 that $I(c^*) = c$. If c is not in the range of u, then since $[-1, 0]$ is a subset of the range of u, we must have $c < -1$, and by definition of I, we have $I(c^*) = |c|I(c^*/|c|) = c$.

For part (b), first suppose that $||b|| \leq 1$ and $b \in \mathcal{B}^-$ (i.e., $0^* \geq b \geq (-1)^*$). Then there exists an act f such that $b_f = b$. By Lemma B.3.3, $f \sim l^*_{I(b)}$. We now need to consider the case that $c \leq 1$ and $c > 1$ separately. If $c \leq 1$, by Independence, $cf_b + (1-c)l^*_0 \sim cl^*_{I(b)} + (1-c)l^*_0$. By Lemma B.3.4, $I(b_{cf_b+(1-c)l^*_0}) = I(b_{cl^*_{I(b)}+(1-c)l^*_0})$. It is easy to check that $b_{cf_b+(1-c)l^*_0} = cb$, and $b_{cl^*_{I(b)}} + (1-c)l^*_0 = cI(b)^*$. Thus, $I(cb) = I(cI(b)^*)$. By part (a), $I(cI(b)^*) = cI(b)$. Thus, $I(cb) = cI(b)$, as desired.

If $c > 1$, there are two subcases. If $||cb|| \leq 1$, since $1/c < 1$, by what we have just shown $I(b) = I(\frac{1}{c}(cb)) = \frac{1}{c}I(cb)$. Crossmultiplying, we have that $I(cb) = cI(b)$, as desired. And if $||cb|| > 1$, by definition, $I(cb) = ||cb||I(bc/||cb||) = c||b||I(b/||b||)$ (since $bc/||cb|| = b/||b||$). Since $||b|| \leq 1$, by what we have shows $I(b) = I(||b||(b/||b||) = ||b||I(b/||b||)$, so $I(b/||b||) = \frac{1}{||b||}I(b)$. Again, it follows that $I(cb) = cI(b)$.

Now suppose that $||b|| > 1$. Then $I(b) = ||b||I(b/||b||)$. Again, we have two subcases. If $||cb|| > 1$, then

$$I(cb) = ||cb||I(cb/||cb||) = c||b||I(b/||b||) = cI(b).$$

And if $||cb|| \leq 1$, by what we have shown for the case $||b|| \leq 1$,

$$I(b) = I(\frac{1}{c}(cb)) = \frac{1}{c}I(cb),$$

so again $I(cb) = cI(b)$.

For part (c), first note that if $b, b' \in \mathcal{B}^-$. If $||b|| \leq 1$ and $||b'|| \leq 1$, then the acts f_b and $f_{b'}$ exist. Moreover, since $b \geq b'$, we must have $(f_b(s))^* \succeq (f_{b'})^*(s)$ for all states $s \in S$. Thus, by Monotonicity, $f_b \succeq f_{b'}$. If either $||b|| > 1$ or $||b'|| > 1$, let $n = \max(||b||, ||b'||)$. Then $||b/n|| \leq 1$ and $||b'/n|| \leq 1$. Thus, $I(b/n) \geq I(b'/n)$, by what we have just shown. By part (b), $I(b) \geq I(b')$. Finally, if either $b \in \mathcal{B} - \mathcal{B}^-$ or $b' \in \mathcal{B} - \mathcal{B}^-$, note that if $b \geq b'$, then $b^- \geq (b')^-$. By definition, $I(b) = I(b^-)$ and $I(b') = I(b')^-$; moreover, $b^-, (b')^- \in \mathcal{B}^-$. Thus, by the argument above, $I(b) \geq I(b^-)$.

For part (d), note that if $b_n \to b$, then for all k, there exists n_k such that $b_n - (1/k)^* \leq b_n \leq b_n + (1/k)^*$ for all $n \geq n_k$. Moreover, by the monotonicity of I (part (c)), we have that $I(b - (1/k)^*) \leq I(b_n) \leq I(b + (1/k)^*)$. Thus, it suffices to show that $I(b - (1/k)^*) \to I(b)$ and that $I(b + (1/k)^*) \to I(b)$.

To show that $I(b - (1/k)^*) \to I(b)$, we must show that for all $\epsilon > 0$, there exists k such that $I(b - (1/k)^*) \geq I(b) - \epsilon$. By positive homogeneity (part (b)), we can assume without loss of generality that $||b - (1/2)^*|| \leq 1$ and that $||b|| \leq 1$. Fix $\epsilon > 0$. If $I(b - (1/2)^*) \geq I(b) - \epsilon$, then we are done. If not, then $I(b) > I(b) - \epsilon > I(b - (1/2)^*)$. Since $||b|| \leq 1$ and $||b - (1/2)^*|| \leq 1$, f_b and $f_{b-(1/2)^*}$ exist. Moreover, by Lemma B.3.4, $f_b \succ f_{(I(b)-\epsilon)^*} \succ f_{b-(1/2)^*}$. By mixture continuity, for some $p \in (0,1)$, we have $pf_b + (1-p)f_{(b-(1/2)^*} \succ f_{(I(b)-\epsilon)^*}$. It is easy to check that $b_{pf_b+(1-p)f_{b-(1/2)^*}} = b - (1-p)(1/2)^*$. Thus, by Lemma B.3.4, $f_{b-(1-p)(1/2)^*} \succeq f_{(I(b)-\epsilon)^*}$, and $I(b-(1-p)1/2)^*) > I(b)-\epsilon$. Choose k such that $1/k < (1-p)(1/2)$. Then

$$I(b-(1/k)^*) \geq I(b-(1-p)1/2)^*) > I(b) - \epsilon,$$

as desired.

The argument that $I(b+(1/k)^*) \to I(b)$ is similar and left to the reader.

For part (e), first suppose that $b, b' \in \mathcal{B}^-$. If $||b||, ||b^-|| \leq 1$, and $I(b), I(b') \neq 0$, consider $\frac{b}{-I(b)}$ and $\frac{b'}{-I(b')}$. Since $I(\frac{b}{-I(b)}) = I(\frac{b'}{-I(b')}) = -1$, it follows from Lemma B.3.3 that $f_{\frac{b}{-I(b)}} \sim f_{\frac{b'}{-I(b')}}$. By ambiguity aversion, for all $p \in (0,1]$, $pf_{\frac{b}{-I(b)}} + (1-p)f_{\frac{b'}{-I(b')}} \succeq f_{\frac{b}{-I(b)}}$. Taking $p = I(b)/(I(b)+I(b'))$, we have that $(I(b)/(I(b)+I(b'))f_{b/I(b)} + (I(b')/(I(b)+I(b'))f_{b'/I(b')} \succeq f_{b/I(b)}$. Therefore, we have

$$I\left(\frac{-I(b)}{-I(b)-I(b')}\frac{b}{-I(b)} + \frac{-I(b')}{-I(b)-I(b')}\frac{b'}{-I(b')}\right) \geq I(\frac{b}{-I(b)}) = -1.$$

Simplifying, we have

$$I\left(\frac{-1}{I(b)+I(b')}b + \frac{-1}{I(b)+I(b')}b'\right) \geq -1,$$

which, together with positive homogeneity of I (part (b)), implies $I(b+b') \geq I(b)+I(b')$, as required.

If $b, b^- \in \mathcal{B}^-$ and either $||b|| > 1$ or $||b'|| > 1$, and both $I(b) \neq 0$ and $I(b') \neq 0$, then the result easily follows by positive homogeneity (property (b)).

If $b, b^- \in \mathcal{B}-$ and either $I(b) = 0$ or $I(b') = 0$, let $b_n = b - \frac{1}{n}^*$ and $b'_n = b' - \frac{1}{n}^*$. Clearly $||b_n|| > 0$, $||b'_n|| > 0$, $b_n \to b$, and $b'_n \to b'_n$. By our argument above, $I(b_n + b'_n) \geq I(b_n) + I(b'_n)$ for all $n \geq 1$. The result now follows from continuity.

Finally, if either $b \in \mathcal{B} - \mathcal{B}^-$ or $b' \in \mathcal{B} - \mathcal{B}^-$, observe that

$$(b+b')^-(s) \begin{cases} = b^-(s) + b'^-(s), \text{ if } b(s) \leq 0, b'(s) \leq 0 \\ = b^-(s) + b'^-(s), \text{ if } b(s) \geq 0, b'(s) \geq 0 \\ \geq b^-(s) + b'^-(s), \text{ if } b(s) > 0, b'(s) \leq 0 \\ \geq b^-(s) + b'^-(s), \text{ if } b(s) \leq 0, b'(s) > 0. \end{cases}$$

Therefore, $(b+b')^- \geq b^- + b'^-$. Thus, $I(b+b') = I((b+b')^-) \geq I(b^- + b'^-)$ by the monotonicity of I, and $I(b^- + b'^-) \geq I(b^-) + I(b'^-)$ by superadditivity of I on $\mathcal{B}^-$. Therefore, $I(b+b') \geq I(b) + I(b')$. □

A.1.2 Defining the weights

In this section, we use I to define a weight $\alpha_{\Pr}$ for each probability $\Pr \in \Delta(S)$. The heart of the proof involves showing that the resulting set $\mathcal{P}^+$ so determined gives us the desired representation.

Given a set $\mathcal{P}^+$ of weighted probability measures, for $b \in \mathcal{B}^-$, define

$$NWREG(b) = \inf_{\Pr \in \mathcal{P}} \alpha_{\Pr}(\sum_{s \in S} b(s) \Pr(s)).$$

Note that $NWREG$ is the negative of the weighted regret when the menu is $\mathcal{B}^-$. Define

$$NREG(b) = \inf_{\Pr \in \mathcal{P}} \sum_{s \in S} b(s) \Pr(s).$$

and

$$NREG_{\Pr}(b) = \sum_{s \in S} b(s) \Pr(s) = E_{\Pr} b.$$

For each probability $\Pr \in \Delta(S)$, define

$$\alpha_{\Pr} = \sup\{\alpha \in \mathbb{R} : \alpha NREG_{\Pr}(b) \geq I(b) \text{ for all } b \in \mathcal{B}^-\}. \quad \text{(A.1)}$$

Note that $\alpha_{\Pr} \geq 0$ for all distributions $\Pr \in \Delta(S)$, since $0 \geq I(b)$ for $b \in \mathcal{B}^-$ (by monotonicity); and $\alpha_{\Pr} \leq 1$, since $NREG_{\Pr}((-1)^*) = I((-1)^*) = -1$ for all distributions $\Pr$. Thus, $\alpha_{\Pr} \in [0, 1]$. Moreover, it is immediate from the definition of $\alpha_{\Pr}$ that $\alpha_{\Pr} NREG_{\Pr}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. The next lemma shows that there exists a probability $\Pr$ where we have equality.

Lemma A.1.4. *(a) For some distribution* $\Pr$*, we have* $\alpha_{\Pr} = 1$*.*

(b) For all $b \in \mathcal{B}^-$*, there exists* $\Pr$ *such that* $\alpha_{\Pr} NREG_{\Pr}(b) = I(b)$*.*

Proof. The proofs of both part (a) and (b) use a standard separation result: If U is an open convex subset of $\mathcal{B}$, and $b \notin U$, then there is a linear functional λ that separates U from b, that is, $\lambda(b') > \lambda(b)$ for all $b' \in U$. We proceed as follows

For part (a), we must show that for some $\Pr$, for all $b \in \mathcal{B}^-$, $NREG_{\Pr}(b) \geq I(b)$. Since $NREG_{\Pr}(b) = E_{\Pr} b$, it suffices to show that $E_{\Pr}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$.

Let $U = \{b' \in \mathcal{B} : I(b') > -1\}$. U is open (by continuity of I), and convex (by positive homogeneity and superadditivity of I), and $(-1)^* \notin U$. Thus, there exists a linear functional λ such that $\lambda(b') > \lambda((-1)^*)$ for $b' \in U$. We want to show that λ is a positive linear functional, that is, that $\lambda(b) \geq 0$ if $b \geq 0^*$. Since $0^* \in U$, and $\lambda(0^*) = 0$, it follows that $\lambda((-1)^*) < 0$. Since λ is linear, we can assume without loss of generality that $\lambda((-1)^*) = -1$. Thus, for all $b' \in \mathcal{B}^-$, $I(b') > -1$ implies $\lambda(b') > -1$. Suppose that $c > 0$ and $b' \geq 0^*$. From the definition of I, it follows that $I(cb') = I(0^*) = 0 > -1$. So $c\lambda(b') = \lambda(cb') > -1$, so $\lambda(b') > -1/c$. (The fact that $I(cb') = I(0^*)$ follows from the definition of I on elements in $\mathcal{B} - \mathcal{B}^-$.) Since this is true for all $c > 0$, it must be the case that $\lambda(b') \geq 0$. Thus, λ is a positive functional.

Define the probability distribution $\Pr$ on S by taking $\Pr(s) = \lambda(1_s)$. To see that $\Pr$ is indeed a probability distribution, note that since $1_s \geq 0$ and λ is positive, we must have $\lambda(1_s) \geq 0$. Moreover, $\sum_{s \in S} \Pr(s) = \lambda(1^*) = 1$. In addition, for all $b' \in \mathcal{B}$, we have

$$\lambda(b') = \sum_{s \in S} \lambda(1_s) b'(s) = \sum_{s \in S} \Pr(s) b'(s) = E_{\Pr}(b').$$

Next note that, for $b \in \mathcal{B}^-$,

$$\text{for all } c < 0, \text{ if } I(b) > c, \text{ then } \lambda(b) > c. \tag{A.2}$$

For if $I(b) > c$, then $I(b/|c|) > -1$ by positive homogeneity, so $\lambda(b/|c|) > -1$ and $\lambda(b) > c$. The result now follows. For if $b \in \mathcal{B}^-$, then $I(b) \leq I(0^*) = 0$ by monotonicity. Thus, if $c < I(b)$, then $c < 0$, so, by (A.2), $\lambda(b) > c$. Since $\lambda(b) > c$ whenever $I(b) > c$, it follows that $E_{\Pr}(b) = \lambda(b) \geq I(b)$, as desired.

The proof of part (b) is similar to that of part (a). We want to show that, given $b \in \mathcal{B}^-$, there exists $\Pr$ such that $\alpha_{\Pr} NREG_{\Pr}(b) = I(b)$. First suppose that $||b|| \leq 1$. If $I(b) = 0$, then there must exist some s such that $b(s) = 0$, for otherwise there exists $c < 0$ such that $b \leq c^*$, so $I(b) \leq c$. If $b(s) = 0$, let $\Pr_s$ be such that $\Pr_s(s) = 1$. Then $NREG_{\Pr_s}(b) = 0$, so (b) holds in this case.

If $||b|| \leq 1$ and $I(b) < 0$, let $U = \{b' : I(b') > I(b)\}$. Again, U is open and convex, and $b \notin U$, so there exists a linear functional λ such that $\lambda(b') > \lambda(b)$ for $b' \in U$. Since $0^* \in U$ and $\lambda(0^*) = 0$, we must have $\lambda(b) < 0$. Since $(-1)^* \leq b$, $(-1)^*$ is not in U, and therefore we also have $\lambda((-1)^*) < 0$. Thus, we can assume without loss of generality that $\lambda((-1)^*) = -1$, and hence $\lambda((1)^*) = 1$. The same argument as above shows that λ is positive: for all $c > 0$ and $b' \geq 0^*$, $I(cb') = 0$ as before. Since $I(b) < 0$, it follows that $I(cb') > I(b)$, so $cb' \in U$ and

$\lambda(cb') > \lambda(b) \geq \lambda((-1)^*) = -1$. Thus, as before, for all $c > 0$, $b' \geq 0^*$, $\lambda(b') > \frac{-1}{c}$, so λ is a positive functional.

Therefore, λ determines a probability distribution Pr such that, for all $b' \in \mathcal{B}^-$, we have $\lambda(b') = E_{\mathrm{Pr}}(b')$. This, of course, will turn out to be the desired distribution. To show this, we need to show that $\alpha_{\mathrm{Pr}} = I(b)/NREG_{\mathrm{Pr}}(b)$. Clearly $\alpha_{\mathrm{Pr}} \leq I(b)/NREG_{\mathrm{Pr}}(b)$, since if $\alpha > I(b)/NREG_{\mathrm{Pr}}(b)$, then $\alpha NREG_{\mathrm{Pr}}(b) < I(b)$ (since $NREG_{\mathrm{Pr}}(b) = \lambda(b) < 0$). To show that $\alpha_{\mathrm{Pr}} \geq I(b)/NREG_{\mathrm{Pr}}b$, we must show that $(I(b)/NREG_{\mathrm{Pr}}(b))NREG_{\mathrm{Pr}}(b') \geq I(b')$ for all $b' \in \mathcal{B}^-$. Equivalently, we must show that $I(b)\lambda(b')/\lambda(b) \geq I(b')$ for all $b' \in \mathcal{B}^-$.

Essentially the same argument used to prove (A.2) also shows

$$\text{for all } c > 0, \text{ if } I(b') > cI(b), \text{ then } \lambda(b') > c\lambda(b).$$

In particular, if $I(b') > cI(b)$, then by positive homogeneity, $\frac{I(b')}{c} > I(b)$, so $\frac{b'}{c} \in U$, and $\lambda(\frac{b'}{c}) > \lambda(b)$ and hence $\lambda(b') > c\lambda(b)$.

Thus, if $I(b')/(-I(b)) > c$ and $c < 0$, then $I(b') > -cI(b)$, and hence $\lambda(b')/(-\lambda(b)) > c$. It follows that $\lambda(b')/(-\lambda(b)) \geq I(b')/(-I(b))$ for all $b' \in \mathcal{B}^-$. Thus, $I(b)\lambda(b')/\lambda(b) \geq I(b')$ for all $b' \in \mathcal{B}^-$, as required.

Finally, if $||b|| > 1$, let $b' = b/||b||$. By the argument above, there exists a probability measure Pr such that $\alpha_{\mathrm{Pr}} NREG_{\mathrm{Pr}}(b/||b||) = I(b/||b||)$. Since $NREG_{\mathrm{Pr}}(b/||b||) = NREG_{\mathrm{Pr}}(b)/||b||$, and $I(b/||b||) = I(b)/||b||$, we must have that $\alpha_{\mathrm{Pr}} NREG_{\mathrm{Pr}}(b) = I(b)$. □

We can now complete the proof of Lemma 2.5.2. By Lemma B.3.6 and the definition of $\alpha_{\Pr}$, for all $b \in \mathcal{B}^-$,

$$\begin{aligned} I(b) &= \inf_{\Pr \in \Delta(S)} \alpha_{\Pr} NREG(b) \\ &= \inf_{\Pr \in \Delta(S)} \left(\alpha_{\Pr} \sum_{s \in S} b(s) \Pr(s) \right) \\ &= \sup_{\Pr \in \Delta(S)} \left(-\alpha_{\Pr} \sum_{s \in S} b(s) \Pr(s) \right). \end{aligned} \tag{A.3}$$

Recall that, by Lemma B.3.4, for all acts f, g such that $b_f, b_g \in \mathcal{B}^-$, $f \succeq g$ iff $I(b_f) \geq I(b_g)$. Thus, $f \succeq g$ iff

$$\sup_{\Pr \in \Delta(S)} \left(-\alpha_{\Pr} \sum_{s \in S} u(f(s)) \Pr(s) \right) \leq \sup_{\Pr \in \Delta(S)} \left(-\alpha_{\Pr} \sum_{s \in S} u(g(s)) \Pr(s) \right).$$

Note that, for $f \in M^* = \mathcal{B}^-$, we have $reg_{M^*,\Pr}(f) = \sup(-u(f(s)) \Pr(s)$, since 0^* dominates all acts in M^*. Thus, $\succeq = \succeq^{S,Y,U}_{M^*,\mathcal{P}^+}$, where $\mathcal{P}^+ = \{(\Pr, \alpha_{\Pr} : \Pr \in \Delta(S)\}$.

We have already observed that U is unique up to affine transformations, so it remains to show that $\mathcal{P}^+$ is maximal. This follows from the definition of $\alpha_{\Pr}$. If $\succeq_M = \succeq^{S,Y,U}_{M,(\mathcal{P}')^+}$, and $(\alpha', \Pr) \in (\mathcal{P}')^+$, then we claim that $\alpha' \in \{\alpha \in \mathbb{R} : \alpha NREG_{\Pr}(b) \geq I(b)$ for all $b \in \mathcal{B}^-\}$. If not, there would be some $b \in \mathcal{B}^-$ with $||b|| \leq \frac{1}{2}$, such that $\alpha' NREG_{\Pr}(b) < I(b)$, which, by the definition of $\prec^{S,Y,U}_{M^*,(\mathcal{P}')^*}$, means that $l^*_{-1} \prec^{S,Y,U}_{M^*,(\mathcal{P}')^+} f_b \prec^{S,Y,U}_{M^*,(\mathcal{P}')^+} l^*_{I(b)}$. Recall that $I(b_f) = \inf\{\gamma : l^*_\gamma \succeq_{M^*} f\}$. Moreover, since $\prec^{S,Y,U}_{M^*,(\mathcal{P}')^+}$ satisfies mixture continuity, there exists some $p \in (0,1)$ such that $f_b \prec^{S,Y,U}_{M^*,(\mathcal{P}')^+} pl^*_{-1} + (1-p)l^*_{I(b)} \prec^{S,Y,U}_{M^*,(\mathcal{P}')^+} \prec^{S,Y,U}_{M^*,(\mathcal{P}')^+} l^*_{I(b)}$. This contradicts the definition of $I(b)$. Therefore, $\alpha' \in \{\alpha \in \mathbb{R} : \alpha NREG_{\Pr}(b) \geq I(b)$ for all $b \in \mathcal{B}^-\}$, and hence $\alpha' \leq \alpha_{\Pr}$.

A.1.3 Uniqueness of Representation

In this section, we show that the canonical set of weighted probabilities we constructed, when viewed as a set of subnormal probability measures, is regular and includes at least one proper probability measure. Moreover, this set of sub-probability measures is the only regular set that induces the preference order $\succeq$ on nonpositive acts. Our uniqueness result is analogous to the uniqueness results of Gilboa and Schmeidler [1989b], who show that the convex, closed, and non-empty set of probability measures in their representation theorem for MMEU is unique. The argument is based on two lemmas: Lemma A.1.5 says that the canonical set of sub-probability measures is regular; and Lemma A.1.6 says that a set of sub-probability measures representing $\succeq$ over nonpositive acts that is regular and contains at least one proper probability measure is unique. The proof of this second lemma, like the proof of uniqueness in Gilboa and Schmeidler [1989b], uses a separating hyperplane theorem to show the existence of acts on which two different representations must 'disagree'. However, a slightly different argument is required in our case, since our acts must have utilities corresponding to nonpositive vectors in $\mathbb{R}^{|S|}$.

Lemma A.1.5. *Let $\mathcal{P}^+$ be the canonical set of weighted probability measures representing $\succeq$. The set $C(\mathcal{P}^+)$ of sub-probability measures is regular.*

Proof. It is useful to note that, by definition, $\mathbf{p} \in C(\mathcal{P}^+)$ if and only if

$$E_{\mathbf{p}}(b) \geq I(b) \text{ for all } b \in \mathcal{B}^-$$

(where expectation with respect to a subnormal probability measure is defined in the obvious way).

Recall that a set is regular if it is convex, closed, and downward-closed. We first show that $C(\mathcal{P}^+)$ is downward-closed. Suppose that $\mathbf{p} \in C(\mathcal{P}^+)$ and $\mathbf{q} \leq \mathbf{p}$ (i.e., $\mathbf{q}(s) \leq \alpha \Pr(s)$ for all $s \in S$. Since $\mathbf{p} \in C(\mathcal{P}^+)$, $E_{\mathbf{p}}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. Since $\mathbf{q} \leq \mathbf{p}$ and, if $b \in \mathcal{B}^-$, we have $b \leq 0^*$, it follows that $E_{\mathbf{q}}(b) \geq E_{\mathbf{p}}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$, and thus $\mathbf{q} \in C(\mathcal{P}^+)$.

To see that $C(\mathcal{P}^+)$ is closed, let $\mathbf{p} = \lim_{n\to\infty} \mathbf{p}_n$, where each $\mathbf{p}_n \in C(\mathcal{P}^+)$. Since $\mathbf{p}_n \in C(\mathcal{P}^+)$ it must be the case that $E_{\mathbf{p}_n}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. By the continuity of expectation, it follows that $E_{\mathbf{p}}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. Thus, $\mathbf{p} \in C(\mathcal{P}^+)$.

To show that $C(\mathcal{P}^+)$ is convex, suppose that $\mathbf{p}, \mathbf{q} \in C(\mathcal{P}^+)$. Then $E_{\mathbf{p}}(b) \geq I(b)$ and $E_{\mathbf{q}}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. It easily follows that for all $a \in (0,1)$, $E_{a\mathbf{p}+(1-a)\mathbf{q}}(b) \geq I(b)$ for all $b \in \mathcal{B}^-$. Thus, $a\mathbf{p} + (1-a)\mathbf{q} \in C(\mathcal{P}^+)$. □

Lemma A.1.6. *A set of sub-probability measures representing $\succeq$ over nonpositive acts that is regular, and has at least one proper probability measure is unique.*

Proof. Suppose for contradiction that there exists two regular sets of subnormal probability distributions, C_1 and C_2, that represent $\succeq$ and have at least one proper probability measure.

First, without loss of generality, let $\mathbf{q} \in C_2 \backslash C_1$. We actually look at an extension of C_1 that is downward-closed in each component to $-\infty$. Let $\overline{C}_1 = \{\mathbf{p} \in \mathbb{R}^{|S|} : \mathbf{p} \leq \mathbf{p}'\}$. Note an element $\mathbf{p}$ of $\overline{C}_1$ may not be subnormal probability measures; we do not require that $\mathbf{p}(s) \geq 0$ for all $s \in S$. Since $\overline{C}_1$ and $\{\mathbf{q}\}$ are closed, convex, and disjoint, and $\{\mathbf{q}\}$ is compact, the separating hyperplane theorem [1970] says that there exists $\theta \in \mathbb{R}^{|S|}$ and $c \in \mathbb{R}$ such that

$$\theta \cdot \mathbf{p} > c \text{ for all } \mathbf{p} \in \overline{C}_1, \text{ and } \theta \cdot \mathbf{q} < c. \tag{A.4}$$

By scaling c appropriately, we can assume that $|\theta(s)| \leq 1$ for all $s \in S$. Now we argue that it must be the case that $\theta(s) \leq 0$ for all $s \in S$ (so that θ corresponds to the utility profile of some act with nonpositive utilities). Suppose that $\theta(s') > 0$ for some $s' \in S$. By (C.4), $\theta \cdot \mathbf{p} > c$ for all $\mathbf{p} \in \overline{C}_1$. However, consider $\mathbf{p}^* \in \overline{C}_1$ defined by

$$\mathbf{p}^*(s) = \begin{cases} 0, \text{ if } s \neq s' \\ \frac{-|c|}{\theta(s)}, \text{ if } s = s'. \end{cases}$$

Clearly, $\theta \cdot \mathbf{p}^* \leq c$, contradicting (C.4). Thus it must be the case that $\theta(s) \leq 0$ for all $s \in S$.

Consider the θ given by the separating hyperplane theorem, and let f be an act such that $u \circ f = \theta$. By continuity, $f \sim l_d^*$ for some constant act l_d^*. Since C_1 and C_2 both represent $\succeq$, and C_1 and C_2 both contain a proper probability measure,

$$\min_{\mathbf{p} \in C_1} \mathbf{p} \cdot (u \circ f) = \min_{\mathbf{p} \in C_1} \mathbf{p} \cdot (u \circ l_d^*) = d = \min_{\mathbf{p} \in C_2} \mathbf{p} \cdot (u \circ f).$$

However, by (C.4),

$$\min_{\mathbf{p} \in C_1} \mathbf{p} \cdot (u \circ f) > c > \min_{\mathbf{p} \in C_2} \mathbf{p} \cdot (u \circ f),$$

which is a contradiction. □

A.2 An Axiomatic characterization of MWER with Preference Relations

We consider an axiomatization based on primitive preference orders $\succeq_M$ indexed by menus. (Because regret is menu-dependent, we cannot consider a single preference order $\succeq$.) As Stoye [2011b] points out, one disadvantage of considering such preference orders is that they are not observable. For example, suppose that $f_1 \succ_M f_2 \succ_M f_3$. By presenting the menu $M = \{f_1, f_2, f_3\}$ to the DM, we can observe that he prefers f_1. But there is no way to observe that $f_2 \succ_M f_3$. The traditional approach (seeing which of f_2 and f_3 the DM prefers when presented with the menu $M' = \{f_2, f_3\}$) will not work, because the DM's preferences with menu M' may be different from those with menu M. (Bleichrodt [**?**] studies a similar problem.)

In Section 2.5, we provided a characterization of MWER with choice functions as the primitives. Despite the fact that a regret-based preference order is not observable, an axiomatization using menu-dependent preference orders allows us to compare the axioms for weighted regret to those for other decision rules.

We state the axioms in a way that lets us clearly distinguish the axioms for SEU, MMEU, MER, and MWER. The axioms are universally quantified over acts f, g, and h, menus M and M', and $p \in (0, 1)$. We assume that $f, g \in M$ when we write $f \succeq_M g$.[1] We use l^* to denote a constant act that maps all states to l.

Axiom A.1. *(Transitivity)* $f \succeq_M g \succeq_M h \Rightarrow f \succeq_M h$.

Axiom A.2. *(Completeness)* $f \succeq_M g$ *or* $g \succeq_M f$.

Let $\mathcal{M}_B$ denote the set of all menus that are bounded above; that is, $\mathcal{M}_B = \{M : \sup_{g \in M} u(g(s))\text{is finite}\}$.

[1] Stoye [2011b] assumed that menus were convex, so that if $f, g \in M$, then so is $pf + (1-p)g$. We do not make this assumption, although our results would still hold if we did (with the axioms slightly modified to ensure that menus are convex). While it may seem reasonable to think that, if f and g are feasible for an DM, then so is $pf + (1-p)g$, this not always the case. For example, it may be difficult for the DM to randomize, or it may be infeasible for the DM to randomize with probability p for some choices of p (e.g., for p irrational).

Axiom A.3. *(Nontriviality)* $f \succ_M g$ *for some acts* f *and* g *and menu* $M \in \mathcal{M}_B$*.*

Up to now, we have taken the set of menus to be $\mathcal{M}_B$. This assumption is necessary (and sufficient) for regret to be well defined. Later, we use Axiom A.3 in the context of different classes of menus. In particular, we are interested in the set of finite menus and the set of *finitely generated* convex menus, that is, the menus M such such that there is a finite set A_M of acts such that M consists of all the convex combinations of acts in A_M. We denote these sets $\mathcal{M}_F$ and $\mathcal{M}_C$, respectively. When we use Axiom A.3 in such contexts, $\mathcal{M}_B$ in Axiom A.3 is understood to be replaced by $\mathcal{M}_C$ and $\mathcal{M}_F$, respectively. Observe that $\mathcal{M}_C \subseteq \mathcal{M}_B$ and $\mathcal{M}_F \subseteq \mathcal{M}_B$.

Axiom A.4. *(Monotonicity) If* $(f(s))^* \succeq_{\{(f(s))^*,(g(s))^*\}} (g(s))^*$ *for all* $s \in S$*, then* $f \succeq_M g$*.*

Axiom A.5. *(Mixture Continuity) If* $f \succ_M g \succ_M h$*, then there exist* $q, r \in (0,1)$ *such that*

$$qf + (1-q)h \succ_{M\cup\{qf+(1-q)h\}} g \succ_{M\cup\{rf+(1-r)h\}} rf + (1-r)h.$$

Menu-independent versions of Axioms A.1–A.5 are standard (for example, (menu-independent versions of) these axioms are in [1989b]). Clearly (menu-independent versions of) Axioms A.1, A.2, A.4, and A.5 hold for MMEU, and SEU; Axiom A.3 is assumed in all the standard axiomatizations, and is used to get a unique representation.

Axiom A.6. *(Ambiguity Aversion)*

$$f \sim_M g \Rightarrow pf + (1-p)g \succeq_{M\cup\{pf+(1-p)g\}} g.$$

Ambiguity aversion says that the DM weakly prefers to hedge her bets. It also holds for MMEU, MER, and SEU, and is assumed in the axiomatizations for MMEU and MER. It is not assumed for the axiomatization of SEU, since it follows from the Independence axiom, discussed next. Independence also holds for MWER, provided that we are careful about the menus involved. Given a menu M and an act h, let $pM + (1-p)h$ be the menu $\{pf + (1-p)h : p \in M\}$.

Axiom A.7. *(Independence)*

$$f \succeq_M g \textit{ iff } pf + (1-p)h \succeq_{pM+(1-p)h} pg + (1-p)h.$$

Independence holds in a strong sense for SEU, since we can ignore the menus. The menu-independent version of Independence is easily seen to imply ambiguity aversion. Independence does not hold for MMEU.

Although we have menu independence for SEU and MMEU, we do not have it for MER or MWER. The following two axioms are weakened versions of menu independence that do hold for MER and MWER.

Axiom A.8. *(Menu independence for constant acts) If l^* and $(l')^*$ are constant acts, then $l^* \succeq_M (l')^*$ iff $l^* \succeq_{M'} (l')^*$.*

In light of this axiom, when comparing constant acts, we omit the menu.

An act h is *never strictly optimal relative to M* if, for all states $s \in S$, there is some $f \in M$ such that $(f(s))^* \succeq (h(s))^*$.

Axiom A.9. *(Independence of Never Strictly Optimal Alternatives (INA)) If every act in M' is never strictly optimal relative to M, then $f \succeq_M g$ iff $f \succeq_{M \cup M'} g$.*

Theorem A.2.1. *For all Y, U, S, and $\mathcal{P}^+$, the family of preference orders $\succeq^{S,Y,U}_{M,\mathcal{P}^+}$ for $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_F$, $\mathcal{M}_C$) satisfies Axioms A.1–A.9. Conversely, if a family of preference orders $\succeq_M$ on the acts in $\Delta(Y)^S$ for $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_F$, $\mathcal{M}_C$) satisfies*

Axioms A.1–A.9, then there exist a utility U on Y and a weighted set $\mathcal{P}^+$ of probabilities on S such that $C(\mathcal{P}^+)$ is regular and $\succeq_M = \succeq_{M,\mathcal{P}^+}^{S,Y,U}$ for all $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_C$, $\mathcal{M}_F$). Moreover, U is unique up to affine transformations, and $C(\mathcal{P}^+)$ is unique in the sense that if $\mathcal{Q}^+$ represents $\succeq_M$, and $C(\mathcal{Q}^+)$ is regular, then $C(\mathcal{Q}^+) = C(\mathcal{P}^+)$.

Proof. Showing that $\succeq_{M,\mathcal{P}^+}^{S,Y,U}$ satisfies Axioms A.1–A.9 is fairly straightforward; we leave details to the reader. Essentially the same proof works for $\mathcal{M}_B$, $\mathcal{M}_C$, and $\mathcal{M}_F$. The proof of the converse is quite nontrivial, although it follows the lines of the proof of other representation theorems. We start by considering $\mathcal{M}_B$.

Using standard techniques, we can show that the axioms guarantee the existence of a utility function U on prizes that can be extended to lotteries in the obvious way, so that $l^* \succeq (l')^*$ iff $U(l) \geq U(l')$. We then use techniques of Stoye [2011b] to show that it suffices to get a representation theorem for a single menu, rather than all menus: the menu consisting of all acts f such that $U(f(s)) \leq 0$ for all states $s \in S$. This allows us to use techniques in the spirit of those used by by Gilboa and Schmeidler [1989a] to represent (unweighted) MMEU.

We show here that if a family of menu-dependent preferences $\succeq_M$ satisfies Axioms A.1–A.9, then $\succeq_M$ can be represented as minimizing expected regret with respect to a set of weighted probabilities and a utility function. Since the proof is somewhat lengthy and complicated, we split it into several steps, each in a separate subsection.

Simplifying the Problem. Our proof starts in much the same way as the proof by Stoye [2011b] of a representation theorem for regret. Lemma B.3.1 guarantees the existence of a utility function U on prizes that can be extended to lotteries in the obvious way, so that $l^* \succeq (l')^*$ iff $U(l) \geq U(l')$. In other words, preferences over all constant acts are represented by the maximization of U on the corresponding lotteries that the constant acts map to. Lemma B.3.1 is a consequence of standard results. Our menus are arbitrary sets of acts, as opposed to convex hulls of a finite number of acts in [2011b]; Lemma A.2.4 shows that Stoye's technique can be adapted to work when menus are arbitrary sets of acts. Finally, following Stoye [2011b], we reduce the proof of existence of a minimax weighted regret representation for the family $\succeq_M$ to the proof of existence of a minimax weighted regret representation for a single menu-independent preference order $\succeq$ (Lemma A.2.5).

Lemma A.2.2. *If Axioms 1-3, 5, 7, and 8 hold, then there exists a nonconstant function* $U : X \to \mathbb{R}$*, unique up to positive affine transformations, such that for all constant acts* l^* *and* $(l')^*$ *and menus* M*,*

$$l^* \succeq_M (l')^* \Leftrightarrow \sum_{\{y:\, l^*(y)>0\}} l(y)U(y) \geq \sum_{\{y:\, l'(y)>0\}} l'(y)U(y).$$

Proof. By menu independence for constant acts, the family of preferences $\succeq_M$ all agree when restricted to constant acts. The lemma then follows from standard results (see, e.g., [1988]), since menu-independence for constant acts, combined with independence, gives the standard independence (substitution) axiom from expected utility theory. □

As is commonly done, given U, we define $u(l) = \sum_{\{y:\, l(y)>0\}} l(y)U(y)$. Thus, $u(l)$ is the expected utility of lottery l. We extend u to constant acts by taking $u(l^*) = u(l)$. Thus, Lemma B.3.1 says that, for all menus M, $l^* \succeq (l')^*$ iff $u(l^*) \geq u((l')^*)$. If c is the utility of some lottery, let l^*_c be a constant lottery that $u(l^*_c) = c$. The following is now immediate. We state it as a lemma so that we can refer to it later.

Lemma A.2.3. *$u(l_c^*) \geq u(l_{c'}^*)$ iff $l_c^* \succeq l_{c'}^*$; similarly, $u(l_c^*) = u(l_{c'}^*)$ iff $l_c^* \sim l_{c'}^*$, and $u(l_c^*) > u(l_{c'}^*)$ iff $l_c^* \succ l_{c'}^*$.*

The key step in showing that we can reduce to a single menu is to show that, roughly speaking, for each menu, there exists a menu-dependent function g_M such that $u(g_M(s)) = -\sup_{f \in M} u(f(s))$. Stoye [2011b] proved a similar result, but he assumed that all menus were obtained by taking the convex hull of a finite set of acts. Because we allow arbitrary bounded menus, this result is not quite true for us. For example, suppose that the range of u is $(-1, \infty]$. Then there may be a menu M such that $\sup_{f \in M} u(f(s)) = 5$, so $-\sup_{f \in M} u(f(s)) = -5$. But there is no act g such that $u(g(s)) = -5$, since u is bounded below by -1. The following weakening of this result suffices for our purpose.

Lemma A.2.4. *There exists a utility function U such that for every menu M, there exists $\epsilon \in (0, 1]$ and constant act l^* such that for all $f, g \in M$, $f \succeq_M g \Leftrightarrow t(f) \succeq_{t(M)} t(g)$, where t has the form $t(f) = \epsilon f + (1-\epsilon) l^*$ and $t(M) = \{t(f) : f \in M\}$. Moreover, there exists an act $g_{t(M)}$ such that $u(g_{t(M)}(s)) = -\sup_{f \in t(M)} u(f(s))$ for all $s \in S$.*

Proof. The nontriviality and monotonicity axioms imply there must exist prizes x and y such that $U(x) > U(y)$. We consider four cases.

Case 1: The range of U is bounded above and below. Then we can rescale so that the range of U is $[-1, 1]$. Thus, there must be prizes x and y such that $U(x) = 1$ and $U(y) = -1$. For all $c \in [-1, 1]$, there must be a prize x' that is a convex combination of x and y such that $u(x') = c$, so we can clearly define a function g_M such that, for all $s \in S$, we have $u(g_M(s)) = -\sup_{f \in M} u(f(s))$. Furthermore, we know that such a g_M exists because it can be formed as an act which maps each state to an appropriate lottery over the prizes x and y. More generally, we know that an act with a certain utility profile exists if its utility for each state is within the range of U. This fact will be used in the other cases as well. Thus, in this case we can take t to be the identity (i.e., $\epsilon = 1$).

Case 2: The range of U is $(-\infty, \infty)$. Again, for all $c \in (\infty, \infty)$, there must exist a prize x such that $u(x) = c$. Since menus are assumed to be bounded above, we can again define the required function g and take $\epsilon = 1$.

Case 3: The range of U is bounded above and unbounded below. Then we can assume without loss of generality that the range is $(-\infty, 1]$, and for all c in the range, there is a prize x such that $u(x) = c$. For all menus M, $\epsilon > 0$, and acts $f, g \in M$, by Independence, we have that

$$f \succeq_M g \Leftrightarrow \epsilon f + (1-\epsilon)l_1^* \succeq_{\epsilon M + (1-\epsilon)l_1^*} \epsilon g + (1-\epsilon)l_1^*.$$

There exists an $\epsilon > 0$ such that for all $s \in S$,

$$1 \geq \sup_{f \in M} \epsilon u(f(s)) + (1-\epsilon) \geq -1.$$

Let $t(f) = \epsilon f + (1-\epsilon)l_1^*$. Clearly there exists an act $g_{t(M)}$ such that $u(g_{t(M)}(s)) = -\sup_{f \in t(M)} u(f(s))$ for all $s \in S$.

Case 4: The range of U is bounded below and unbounded above. By the upper-boundedness axiom, every menu has an upper bound on its utility range. Therefore, for every menu M, $\epsilon > 0$, and all acts f and g in M, by Independence,

$$f \succeq_M g \Leftrightarrow \epsilon f + (1-\epsilon)l_{-1}^* \succeq_{\epsilon M + (1-\epsilon)l_{-1}^*} \epsilon g + (1-\epsilon)l_{-1}^*.$$

There exists $\epsilon > 0$ such that for all $s \in S$,

$$\sup_{f \in M} \epsilon u(f(s)) + (1-\epsilon)u(l_{-1}^*(s)) \leq 1.$$

Let $t(f) = \epsilon f + (1-\epsilon)l_{-1}^*$. Again, it is easy to see that $g_{t(M)}$ exists. □

In light of Lemma A.2.4, we henceforth assume that the utility function u derived from U is such that its range is either $(-\infty, \infty)$, $[-, 1, 1]$, $(-\infty, 1]$, or $[-1, \infty)$. In any case, its range always includes $[-1, 1]$.

Before proving the key lemma, we establish some useful notation for acts and utility acts (real-valued functions on S). Given a utility act b, let f_b, the act corresponding to b, be the act such that $f_b(s) = l_{b(s)}$, if such an act exists. Conversely, let b_f, the utility act corresponding to the act f, be defined by taking $b_f(s) = u(f(s))$. Note that monotonicity implies that if $f_b = g_b$, then $f \sim_M g$ for all menus M. That is, only utility acts matter. If c is a real, we take c^* to be the constant utility act such that $c^*(s) = c$ for all $s \in S$.

Lemma A.2.5. *Let M^* be the menu consisting of all acts f such that $(-1)^* \leq b_f \leq 0^*$. Then $(U, \mathcal{P}^+)$ represents $\succeq_{M^*}$ (i.e., $\succeq_{M^*} = \succeq^{S,X,U}_{M^*,\mathcal{P}^+}$) iff $(U, \mathcal{P}^+)$ represents $\succeq_M$ for all menus M.*

Proof. Our arguments are similar in spirit to those of Stoye [2011b].

By Lemma A.2.4, there exists t such that $t(f) = \epsilon f + (1 - \epsilon)h$ for a constant function h such that

$$f \succeq_M g \text{ iff } t(f) \succeq_{t(M)} t(g);$$

moreover, for this choice of t, the act $g_{t(M)}$ defined in Lemma A.2.4 exists.

By Independence,

$$t(f) \succeq_{t(M)} t(g) \text{ iff } \frac{1}{2}t(f) + \frac{1}{2}g_{t(M)} \succeq_{\frac{1}{2}t(M)+\frac{1}{2}g_{t(M)}} \frac{1}{2}t(g) + \frac{1}{2}g_{t(M)}.$$

Let M^* be the menu that contains all acts with utilities in $[-1, 0]$. By INA, we know that for all acts f and g, and menus M for which g_M is defined, we have

$$f \succeq_M g \text{ iff } \frac{1}{2}f + \frac{1}{2}g_M \succeq_{M^*} \frac{1}{2}g + \frac{1}{2}g_M.$$

This is because acts of the form $\frac{1}{2}f + \frac{1}{2}g_M$ are never strictly optimal with respect to the menu $\frac{1}{2}M + \frac{1}{2}g_M$. At every state s there must be some act in $\frac{1}{2}M + \frac{1}{2}g_M$ that has utility 0 at s (namely, the mixture that involves an act $f \in M$ whose utility at s is maximal; that is, $u(f(s)) \geq \max_{f' \in M} u(f'(s))$. Thus,

$$f \succeq_M g \text{ iff } \frac{1}{2}t(f) + \frac{1}{2}g_{t(M)} \succeq_{M^*} \frac{1}{2}t(g) + \frac{1}{2}g_{t(M)}.$$

Since the MWER representation also satisfies Independence and INA, we know that for all menus M, and acts f and g in M,

$$f \succeq^{S,X,U}_{M,\mathcal{P}^+} g \Leftrightarrow t(f) \succeq^{S,X,U}_{t(M),\mathcal{P}^+} t(g) \Leftrightarrow \frac{1}{2}t(f) + \frac{1}{2}g_{t(M)} \succeq^{S,X,U}_{M^*,\mathcal{P}^+} \frac{1}{2}t(g) + \frac{1}{2}g_{t(M)}.$$

Therefore, to show that $\succeq_M$ has a MWER representation with respect to $(U, \mathcal{P}^+)$, it suffices to show that $\succeq_{M^*}$ has a MWER representation with respect to $(U, \mathcal{P}^+)$.

□

In the sequel, we drop the menu subscript when we refer to the family of preferences, and just write $\succeq$ (to denote $\succeq_{M^*}$); by Lemma A.2.5, it suffices to consider $\succeq_{M^*}$.

It is straightforward to check that $\succeq_{M^*}$ satisfies completeness, transitivity, nontriviality, monotonicity, mixture continuity, independence, INA, and ambiguity aversion. Therefore, by Lemma 2.5.2, there exists some $(U, \mathcal{P}^+)$ representing $\succeq_{M^*}$. By Lemma A.2.5, $(U, \mathcal{P}^+)$ represents $\succeq_M$ for all menus M, as required.

Since the axioms hold for all menus in $\mathcal{M}_B$, they clearly continue to hold if we restrict to $\mathcal{M}_F$ and $\mathcal{M}_C$. To prove the converse in the case of $\mathcal{M}_F$ we first argue that if the preference orders $\succeq_M$ for $M \in \mathcal{M}_F$ satisfy the axioms, then they uniquely determine preference orders $\succeq_M$ for menus $M \in \mathcal{M}_B$ that also satisfy the axioms. Clearly, it also then follows that the set of preference orders $\succeq_M$ for $M \in \mathcal{M}_C$ determines $\succeq_M$ for $M \in \mathcal{M}_B$. The proof immediately follows from this observation and the proof in the case of $\mathcal{M}_B$.

Consider a bounded menu M. The *utility frontier* of menu M is a function mapping each state to the maximum utility achieved in that state by any act in M. Since S is assumed to be finite, there exists a finite subset $M' \subseteq M$ such that the utility frontier of M' is the same as the utility frontier of M. Therefore, for all acts $f, g \in M$,

$$f \succeq_M g \Leftrightarrow f \succeq_{M' \cup \{f,g\}} g,$$

by Axiom 2.5. Since M' is finite, we have shown what we need.

Finally, for $\mathcal{M}_C$, suppose that M is a convex set of acts generated by the finite set A_M. Then, for all $f, g \in A_M$,

$$f \succeq_{A_M} g \Leftrightarrow f \succeq_M g,$$

by Axiom 2.5, since no interior points in M can be strictly optimal; hence, interior points can be removed from M without changing preferences. Thus, the result for $\mathcal{M}_C$ follows from the result for $\mathcal{M}_F$. □

It is instructive to compare Theorem A.2.1 to other representation results in the literature. Anscombe and Aumann [1963] showed that the menu-independent versions of axioms A.1–A.5 and A.7 characterize SEU. The presence of Axiom A.7 (menu-independent Independence) greatly simplifies things. Gilboa and Schmeidler [1989b] showed that axioms A.1–A.6 together with one

	SEU	REG	MER	MWER	MMEU
Ax. 1-6,8-10	✓	✓	✓	✓	✓
Ind	✓	✓	✓	✓	
C-Ind	✓				✓
Ax. 12	✓	✓	✓		
Symmetry	✓	✓			

Table A.1: Characterizing axioms for several decision rules.

more axiom that they call *certainty-independence* characterizes MMEU. Certainty-independence, or C-independence for short, is a weakening of independence (which, as we observed, does not hold for MMEU), where the act h is required to be a constant act. Since MMEU is menu-independent, we state it in a menu-independent way.

Axiom A.10. *(C-Independence) If h is a constant act, then $f \succeq g$ iff $pf + (1-p)h \succeq pg + (1-p)h$.*

Table A.1 describes the relationship between the axioms characterizing the decision rules.

A.3 Characterizing MWER with Likelihood Updating

We can in fact directly translate the MDC axiom into a setting with preference relations instead of choice functions.

Definition A.3.1 (Null event). *An event E is* null *if, for all $f, g \in \Delta(Y)^S$ and menus M with $fEg, g \in M$, we have $fEg \sim_M g$.*

MDC. For all non-null events E, $f \succeq_{E,M} g$ iff $fEh \succeq_{MEh} gEh$ for some $h \in M$.[2]

[2]Although we do not need this fact, it is worth noting that the MWER decision rule has the

The key feature of MDC is that it allows us to reduce all the conditional preference orders $\succeq_{E,M}$ to the unconditional order $\succeq_M$, to which we can apply Theorem A.2.1.

Theorem A.3.2. *For all Y, U, S, $\mathcal{P}^+$, and $M \in \mathcal{M}_B$, the family of preference orders $\succeq^{S,Y,U}_{\mathcal{P}^+|E,M}$ for events E such that $\overline{\mathcal{P}}^+(E) > 0$ satisfies Axioms A.1–A.9 and MDC. Conversely, if a family of preference orders $\succeq_{E,M}$ on the acts in $\Delta(Y)^S$ satisfies Axioms A.1–A.9 and MDC for $M \in \mathcal{M}_B$ (resp., $\mathcal{M}_F$, $\mathcal{M}_C$), then there exists a utility U on Y and a weighted set $\mathcal{P}^+$ of probabilities on S such that $C(\mathcal{P}^+)$ is regular, and for all non-null E, $\succeq_{E,M} = \succeq^{S,Y,U}_{\mathcal{P}^+|E,M}$. Moreover, U is unique up to affine transformations, and $C(\mathcal{P}^+)$ is unique in the sense that if $\mathcal{Q}^+$ represents $\succeq_{E,M}$, and $C(\mathcal{Q}^+)$ is regular, then $C(\mathcal{Q}^+) = C(\mathcal{P}^+)$.*

Proof. Since $\succeq_M = \succeq_{S,M}$ satisfies Axioms A.1–A.9, there must exist a weighted set $\mathcal{P}^+$ of probabilities on S and a utility function U such that $f \succeq_M g$ iff $f \succeq^{S,Y,U}_{M,\mathcal{P}^+} g$. The rest of the proof is identical to that of Theorem 2.6.2; we do not repeat it here. □

Analogues of MDC have appeared in the literature before in the context of updating preference orders. In particular, Epstein and Schneider [1993] discuss a menu-independent version of MDC, although they do not characterize updating in their framework. Ghirardato [2002] characterizes update for a menu-independent version of DC. Sinischalchi [2011] also uses an analogue of MDC in his axiomatization of measure-by-measure updating of MMEU. Like us, he starts with an axiomatization for unconditional preferences, and adds an axiom called *constant-act dynamic consistency* (CDC), somewhat analogous to MDC, to

property that $fEh \succeq_{MEh} gEh$ for some act h iff $fEh \succeq_{MEh} gEh$ for all acts h. Thus, this property follows from Axioms A.1–A.9.

extend the axiomatization of MMEU to deal with conditional preferences. CDC in the form in [2011] was first proposed by Pires [2002], based on an observation of Jaffray [1992].

APPENDIX B

PROOFS FOR CHAPTER 3

B.1 Proof of Theorem 3.5.2

Proof of Theorem 3.5.2 . We assume that t is continuous and strictly decreasing, and that there exists some $\beta > 0$ such that $[\beta, \beta/\alpha_0^*(\phi)] \in range(t)$. Recall that $\alpha_0^*(\phi) = \max\{\alpha_0, \min_{p\in\Delta(S)} \phi(p)\}$.

Let $\alpha'_0 = t^{-1}(\frac{\beta}{\alpha_0^*})$ and for all $p \in \Delta(S)$, let $\phi'(p) = t^{-1}(\frac{\beta}{\phi(p)})$. It is easy to see that, for all acts f, g,

$$\begin{aligned}
&\min_{p\in L_{\alpha_0}\phi} \frac{1}{\phi(p)} \int_S u(f)dp \geq \min_{p\in L_{\alpha_0}\phi} \frac{1}{\phi(p)} \int_S u(g)dp \\
\text{iff } &\min_{\{p:\phi(p)\geq\alpha'_0\}} \frac{\beta}{\phi(p)} \int_S u(f)dp \geq \min_{\{p:\phi(p)\geq\alpha'_0\}} \frac{\beta}{\phi(p)} \int_S u(g)dp \\
\text{iff } &\min_{p\in L_{\alpha'_0}\phi'} t(\phi'(p)) \int_S u(f)dp \geq \min_{p\in L_{\alpha'_0}\phi'} t(\phi'(p)) \int_S u(g)dp,
\end{aligned}$$

since for all $p \in L_{\alpha_0}\phi$,

$$\frac{\beta}{\phi(p)} = t(t^{-1}(\frac{\beta}{\phi(p)})) = t(\phi'(p)).$$

Now we show that if $t(1) = \beta$, then ϕ' must be a regular* fuzzy set. Since ϕ is normal, there exists p^* such that $\phi(p^*) = 1$. By definition of ϕ', $\phi'(p^*) = t^{-1}(\frac{\beta}{\phi(p^*)}) = t^{-1}(\beta) = 1$, so ϕ' is normal.

To show that ϕ' is weakly* upper semicontinuous, we must show that the set $L_\alpha\phi' = \{p \in \Delta(S) : \phi'(p) \geq \alpha\}$ is weakly* closed for $\alpha \in [0, 1]$. In other words, we have to show that the set $L_\alpha\phi'$ contains all of its limit points, for all $\alpha \in [0, 1]$. Now, for $\alpha = 0$, $L_\alpha\phi' = \Delta(S)$ and is closed. So consider the case $\alpha > 0$.

Recall from our definition of ϕ' that $\phi'(p) = t^{-1}(\frac{\beta}{\phi(p)})$ for all p. Suppose $p_n \to p$. Observe that $\frac{\beta}{\phi(p_n)}$ is in the domain of t^{-1} for all n, since $[\beta, \beta/\alpha_0^*(\phi)] \in range(t)$. Note that for all p, $\phi'(p) \geq \alpha$ if and only if

$$\begin{aligned} & t^{-1}(\frac{\beta}{\phi(p)}) \geq \alpha \\ \text{iff } & \frac{\beta}{\phi(p)} \leq t(\alpha) \\ \text{iff } & \phi(p) \geq \frac{\beta}{t(\alpha)}, \end{aligned}$$

where $t(\alpha) \geq \beta$ since $0 < \alpha \leq 1$, t is monotonically decreasing, and $t(1) = \beta$. Since ϕ is assumed to be weakly* upper semicontinuous, and $\phi(p_n) \geq \frac{\beta}{t(\alpha)}$ for all n, we have $\phi(p) \geq \frac{\beta}{t(\alpha)}$. Therefore, $\phi'(p) \geq \alpha$, as required.

Finally, to show that ϕ' is quasi-concave, let $\gamma \in [0, 1]$. Using the fact that t is strictly monotonically decreasing, we have that

$$\begin{aligned} & \phi(\gamma p_1 + (1-\gamma)p_2) \geq \min(\phi(p_1), \phi(p_2)) \\ \Rightarrow & \frac{\beta}{\phi(\gamma p_1 + (1-\gamma)p_2)} \leq \max(\frac{\beta}{\phi(p_1)}, \frac{\beta}{\phi(p_2)}) \\ \Rightarrow & t^{-1}(\frac{\beta}{\phi(\gamma p_1 + (1-\gamma)p_2)}) \geq \min(t^{-1}(\frac{\beta}{\phi(p_1)}), t^{-}(\frac{\beta}{\phi(p_2)})) \\ \Rightarrow & \phi'(\gamma p_1 + (1-\gamma)p_2) \geq \min(\phi'(p_1), \phi'(p_2)). \end{aligned}$$

For the other direction, suppose that $t(1) = \beta$ and that ϕ' is a regular* fuzzy confidence function. We want to show that ϕ defined by $\phi(p^*) = \frac{1}{t(\phi'(p^*))}$ is also regular* fuzzy. The arguments for this direction are analogous to those used to show the first direction.

□

B.2 Details of Example 3.6.2

We now show that for all $n \geq 1$, $f_n \sim_\phi^+ \tilde{1}$.

Suppose $c \in [\frac{1}{2^{2m+1}}, \frac{1}{2^{2m-1}})$. The weighted expected utility of f_n with respect to p_c is

$$2^m \left[c2^n + (1-c)\frac{2^n}{2^{2n+1}-1} \right], \text{ for } m \in \{0, 1, 2, \ldots\}.$$

If $m = n$, note that

$$2^n \left[\frac{1}{2^{2n+1}}2^n + \frac{2^{2n+1}-1}{2^{2n+1}}\frac{2^n}{2^{2n+1}-1} \right] = 1.$$

If $m < n$, then

$$\begin{aligned}
2^m \left[c2^n + (1-c)\frac{2^n}{2^{2n+1}-1} \right] &\geq 2^m \left[\frac{1}{2^{2m+1}}2^n + \frac{2^{2m-1}-1}{2^{2m-1}}\frac{2^n}{2^{2n+1}-1} \right] \\
&= \frac{2^n}{2^{m+1}} + \frac{2^{2m-1}-1}{2^{m-1}}\frac{2^n}{2^{2n+1}-1} \\
&\geq \frac{2^n}{2^{m+1}} \geq 1.
\end{aligned}$$

If $m > n$, then

$$\begin{aligned}
2^m \left[c2^n + (1-c)\frac{2^n}{2^{2n+1}-1} \right] &\geq 2^m \left[\frac{1}{2^{2m+1}}2^n + \frac{2^{2m-1}-1}{2^{2m-1}}\frac{2^n}{2^{2n+1}-1} \right] \\
&= \frac{2^n}{2^{m+1}} + \frac{2^{2m-1}-1}{2^{m-1}}\frac{2^n}{2^{2n+1}-1} \\
&\geq \frac{2^n}{2^{m+1}} + \frac{2^{2m-1}-1}{2^{m-1}}\frac{1}{2^{n+1}} \\
&\geq \frac{2^n}{2^{m+1}} + \frac{2^{2m-1}}{2^{m+n}} - \frac{1}{2^{m+n}} \\
&\geq 2^{m-n-1} \geq 1.
\end{aligned}$$

If $c \in [\frac{1}{2}, 1]$, then the weighted expected utility of f_n is

$$c2^n + (1-c)\frac{2^n}{2^{2n+1}-1} \geq c2^n \geq 2^{n-1}.$$

Therefore, for all n, the minimum weighted expected utility of f_n is 1, so $f_n \sim_{\phi}^{+} \tilde{1}$.

Now let $\tilde{m}$ be a constant act with constant utility m. The act $\frac{1}{2}f_n + \frac{1}{2\delta}\tilde{\delta}$ has utility $2^{n-1} + \frac{1}{2}$ in state s_1 and utility $\frac{2^{n-1}}{2^{2n+1}-1} + \frac{1}{2}$ in state s_2. If $c \in [\frac{1}{2^{2m+1}}, \frac{1}{2^{2m-1}})$ for $m \geq 1$, then the weighted expected utility of $\frac{1}{2}f_n + \frac{1}{2\delta}\tilde{\delta}$ with respect to p_c is

$$\begin{aligned}
&2^m \left[c\left(2^{n-1} + \frac{1}{2}\right) + (1-c)\left(\frac{2^{n-1}}{2^{2n+1}-1} + \frac{1}{2}\right)\right] \\
\geq &\frac{1}{2^{m+1}}\left(2^{n-1}\right) + \frac{2^{2m-1}-1}{2^{m-1}}\left(\frac{1}{2}\right) \\
\geq &2^{n-m-2} + 2^{m-2}.
\end{aligned}$$

Suppose that $n \geq 4 + 2\log_2 \delta$ and $\delta \geq 1$. If $n \geq m + 2 + \log_2 \delta$,

$$2^{n-m-2} + 2^{m-2} > 2^{\log_2 \delta} = \delta.$$

Otherwise, if $n < m + 2 + \log_2 \delta$, since $n \geq 4 + 2\log_2 \delta$, it follows that $m \geq \log_2 \delta + 2$, and

$$2^{n-m-2} + 2^{m-2} > 2^{\log_2 \delta} = \delta.$$

If $c \geq \frac{1}{2}$, then the weighted expected utility of $\frac{1}{2}f_n + \frac{1}{2\delta}\tilde{\delta}$ with respect to p_c is

$$\begin{aligned}
&c\left(2^{n-1} + \frac{1}{2}\right) + (1-c)\left(\frac{2^{n-1}}{2^{2n+1}-1} + \frac{1}{2}\right) \\
>&\frac{1}{2}2^{n-1} \geq \frac{1}{2}2^{3+2\log_2 \delta} \geq 2^2\delta^2 > \delta,
\end{aligned}$$

since $\delta \geq 1$. This means that if $n \geq 4 + 2\log_2 \delta$, then the minimum weighted expected utility of $\frac{1}{2}f_n + \frac{1}{2\delta}\tilde{\delta}$ is strictly greater than δ.

B.3 Proof of Theorem 4.5.3

We show here that if a family of preferences $\succeq$ satisfies Axioms 3.1–5.9, then $\succeq$ can be represented as maximizing weighted expected utility with respect to a regular confidence function and a utility function. We make use of many of the same techniques as used in [Halpern and Leung 2012]. Key differences are highlighted.

First, we establish a von-Neumann-Morgenstern expected utility function over constant acts. This part follows the CF proof, rather than the proof in [Halpern and Leung 2012].

Lemma B.3.1. *If Axioms 3.1, 3.3 and 3.5 hold, then there exists a nonconstant function $U : X \to \mathbb{R}$, unique up to positive affine transformations, such that for all constant acts l^* and $(l')^*$,*

$$l^* \succeq (l')^* \Leftrightarrow \sum_{\{y:\, l^*(y)>0\}} l(y)U(y) \geq \sum_{\{y:\, l'(y)>0\}} l'(y)U(y).$$

Proof. As noted by CF, it was shown by Herstein and Milnor [Herstein and Milnor 1953] that Axioms 3.1, 3.3 and 3.5 are sufficient to satisfy the premises of the von-Neumann-Morgenstern theorem. □

Since U is nonconstant, we can choose a U such that the minimum value that it takes on is 0 (for some constant act), and the maximum value it takes on is at least 1. If c is the utility of some lottery l_c, let l_c^* be a constant act such that $l^*(s) = l_c$, so that $u(l_c^*) = c$. The following lemma, whose proof is given in [Halpern and Leung 2012] (Lemma 2), follows from Lemma B.3.1.

Lemma B.3.2. $u(l_c^*) \geq u(l_{c'}^*)$ *iff* $l_c^* \succeq l_{c'}^*$*; similarly,* $u(l_c^*) = u(l_{c'}^*)$ *iff* $l_c^* \sim l_{c'}^*$*, and* $u(l_c^*) > u(l_{c'}^*)$ *iff* $l_c^* \succ l_{c'}^*$.

In [Halpern and Leung 2012] a slightly different continuity axiom (Axiom B.1) is used.

Axiom B.1 (Mixture Continuity). *If $f \succ g \succ h$, then there exist $q, r \in (0,1)$ such that*

$$qf + (1-q)h \succ g \succ rf + (1-r)h.$$

It is not difficult to derive Mixture Continuity from completeness (Axiom 3.1) and Axiom 3.3. Therefore, from here on, we assume that the preference order satisfies Mixture Continuity.

We establish some useful notation for acts and utility acts (real-valued functions on S). Given a utility act b, let f_b, the act corresponding to b, be the act such that $f_b(s) = l_{b(s)}$, if such an act exists. Conversely, let b_f, the utility act corresponding to the act f, be defined by taking $b_f(s) = u(f(s))$. Note that monotonicity implies that if $f_b = g_b$, then $f \sim g$. That is, only utility acts matter. If c is a real, we take c^* to be the constant utility act such that $c^*(s) = c$ for all $s \in S$.

B.3.1 Defining a functional on utility acts

Our proof uses the same technique as that used in [Halpern and Leung 2012]. Specifically, like Gilboa and Schmeidler [Gilboa and Schmeidler 1989b], we define a functional I on utility acts such that the preference order on utility acts is determined by their value according to I (see Lemma B.3.4). Using I, we can then determine the weight of each probability in $\Delta(S)$, and prove the desired representation theorem.

Recall that u represents $\succeq$ on constant acts, and that only utility acts matter to $\succeq$. The space of all nonnegative utility acts is the set $\mathcal{B}^+$ of real-valued functions b on S where $b(s) \geq 0$ for all $s \in S$. We now define a functional I on utility acts in $\mathcal{B}^+$ such that for all f, g with $b_f, b_g \in \mathcal{B}^+$, we have $I(b_f) \geq I(b_g)$ iff $f \succeq g$. Let

$$R_f = \{\alpha' : l^*_{\alpha'} \preceq f\}.$$

If $0^* \leq b \leq 1^*$, then f_b exists, and we define

$$I(b) = \sup(R_{f_b}).$$

For the remaining utility acts $b \in \mathcal{B}^+$, we extend I by homogeneity. Let $||b|| = |\max_{s \in S} b(s)|$. Note that if $b \in \mathcal{B}^+$, then $0^* \leq b/||b|| \leq 1^*$, so we define

$$I(b) = ||b|| I(b/||b||).$$

It is worth noting that while in [Halpern and Leung 2012] I was extended from the nonpostive utility acts to the entire set of real-valued acts in order to invoke a separating theorem for Banach spaces, the extension is not performed here. Consequently, we will be using a different separating hyperplane theorem than in [Halpern and Leung 2012].

Lemma B.3.3. *If $b_f \in \mathcal{B}^+$, then $f \sim l^*_{I(b_f)}$.*

Proof. Suppose that $b_f \in \mathcal{B}^+$ and, by way of contradiction, that $l^*_{I(b_f)} \prec f$. If $f \sim l^*_0$, then it must be the case that $I(b_f) = 0$, since $I(b_f) \geq 0$ by definition of sup, and $f \sim l^*_0 \prec l^*_\epsilon$ for all $\epsilon > 0$ by Lemma B.3.2, so $I(b_f) < \epsilon$ for all $\epsilon < 0$. Therefore, $f \sim l^*_{I(b_f)}$. Otherwise, since $b_f \in \mathcal{B}^+$, by monotonicity, we must have $l^*_0 \prec f$, and thus $l^*_0 \prec f \prec l^*_{I(b_f)}$. By mixture continuity, there is some $q \in (0, 1)$ such that $q \cdot l^*_0 + (1-q) \cdot l^*_{I(b_f)} \sim l_{(1-q)I(b_f)} \succ f$, contradicting the fact that $I(b)$ is the least upper bound of R_f.

If, on the other hand, $l^*_{I(b_f)} \succ f$, then $l^*_{I(b_f)} \succ f \succeq l^*_{\underline{c}}$, where the existence of $l^*_{\underline{c}}$ is guaranteed by Axiom 3.4. If $f \sim l^*_{\underline{c}}$ then it must be the case that $I(b_f) = \underline{c}$. This is because $I(b_f) \geq \underline{c}$ since $l^*_{\underline{c}} \succeq l^*_{\underline{c}}$, and $I(b_f) \leq \underline{c}$ since for all $c' > \underline{c}$, $l^*_{c'} \succ f \sim l^*_{\underline{c}}$.

Otherwise, $l^*_{I(b_f)} \succ f \succ l^*_{\underline{c}}$, and by Axiom 3.3, there is some $q \in (0,1)$ such that $q \cdot l^*_{I(b_f)} + (1-q)l^*_{\underline{c}} \prec f$. Since $qI(b_f) + (1-q)\underline{c} > I(b_f)$, this contradicts the fact that $I(b_f)$ is an upper bound of R_f. Therefore, it must be the case that $l^*_{I(b_f)} \sim f$. □

We can now show that I has the required property.

Lemma B.3.4. *For all acts f, g such that $b_f, b_g \in \mathcal{B}^+$, $f \succeq g$ iff $I(b_f) \geq I(b_g)$.*

Proof. Suppose that $b_f, b_g \in \mathcal{B}^+$. By Lemma B.3.3, $l^*_{I(b_f)} \sim f$ and $g \sim l^*_{I(b_g)}$. Thus, $f \succeq g$ iff $l^*_{I(b_f)} \succeq l^*_{I(b_g)}$, and by Lemma B.3.2, $l^*_{I(b_f)} \succeq l^*_{I(b_g)}$ iff $I(b_f) \geq I(b_g)$. □

We show that the axioms guarantee that I has a number of standard properties. The proof of each property is analogous to its counterpart in [Halpern and Leung 2012], but here we deal with nonnegative utility acts, as opposed to nonpositive utility acts.

Lemma B.3.5. *(a) If $c \geq 0$, then $I(c^*) = c$.*

(b) I satisfies positive homogeneity: if $b \in \mathcal{B}^+$ and $c > 0$, then $I(cb) = cI(b)$.

(c) I is monotonic: if $b, b' \in \mathcal{B}^+$ and $b \geq b'$, then $I(b) \geq I(b')$.

(d) I is continuous: if $b, b_1, b_2, \ldots \in \mathcal{B}^+$, and $b_n \to b$, then $I(b_n) \to I(b)$.

(e) I is superadditive: if $b, b' \in \mathcal{B}^+$, then $I(b + b') \geq I(b) + I(b')$.

Proof. For part (a), if c is in the range of u, then it is immediate from the definition of I and Lemma B.3.2 that $I(c^*) = c$. If c is not in the range of u, then since

$[0,1]$ is a subset of the range of u, we must have $c > 1$, and by definition of I, we have $I(c^*) = |c|I(c^*/|c|) = c$.

For part (b), first suppose that $||b|| \leq 1$ and $b \in \mathcal{B}^+$ (i.e., $0^* \leq b \leq 1^*$). Then there exists an act f such that $b_f = b$. By Lemma B.3.3, $f \sim l^*_{I(b)}$. We now consider the case that $c \leq 1$ and $c > 1$ separately. If $c \leq 1$, by Worst Independence, $cf_b + (1-c)l_0^* \sim cl^*_{I(b)} + (1-c)l_0^*$. By Lemma B.3.4, $I(b_{cf_b+(1-c)l_0^*}) = I(b_{cl^*_{I(b)}+(1-c)l_0^*})$. It is easy to check that $b_{cf_b+(1-c)l_0^*} = cb$, and $b_{cl^*_{I(b)}} + (1-c)l_0^* = cI(b)^*$. Thus, $I(cb) = I(cI(b)^*)$. By part (a), $I(cI(b)^*) = cI(b)$. Thus, $I(cb) = cI(b)$, as desired.

If $c > 1$, there are two subcases. If $||cb|| \leq 1$, since $1/c < 1$, by what we have just shown $I(b) = I(\frac{1}{c}(cb)) = \frac{1}{c}I(cb)$. Crossmultiplying, we have that $I(cb) = cI(b)$, as desired. And if $||cb|| > 1$, by definition, $I(cb) = ||cb||I(bc/||cb||) = c||b||I(b/||b||)$ (since $bc/||cb|| = b/||b||$). Since $||b|| \leq 1$, by the earlier argument, $I(b) = I(||b||(b/||b||) = ||b||I(b/||b||)$, so $I(b/||b||) = \frac{1}{||b||}I(b)$. Again, it follows that $I(cb) = cI(b)$.

Now suppose that $||b|| > 1$. Then $I(b) = ||b||I(b/||b||)$. Again, we have two subcases. If $||cb|| > 1$, then

$$I(cb) = ||cb||I(cb/||cb||) = c||b||I(b/||b||) = cI(b).$$

And if $||cb|| \leq 1$, by what we have shown for the case $||b|| \leq 1$,

$$I(b) = I(\frac{1}{c}(cb)) = \frac{1}{c}I(cb),$$

so again $I(cb) = cI(b)$.

For part (c), first note that for $b, b' \in \mathcal{B}^+$, if $||b|| \leq 1$ and $||b'|| \leq 1$, then the acts f_b and $f_{b'}$ exist. Moreover, since $b \geq b'$, we must have $(f_b(s))^* \succeq (f_{b'}(s))^*$ for all states $s \in S$. Thus, by Monotonicity, $f_b \succeq f_{b'}$. If either $||b|| > 1$ or $||b'|| > 1$, let

$n = \max(||b||, ||b'||)$. Then $||b/n|| \le 1$ and $||b'/n|| \le 1$. Thus, $I(b/n) \ge I(b'/n)$, by what we have just shown. By part (b), $I(b) \ge I(b')$.

For part (d), note that if $b_n \to b$, then for all k, there exists n_k such that $b_n - (1/k)^* \le b_n \le b_n + (1/k)^*$ for all $n \ge n_k$. Moreover, by the monotonicity of I (part (c)), we have that $I(b - (1/k)^*) \le I(b_n) \le I(b + (1/k)^*)$. Thus, it suffices to show that $I(b - (1/k)^*) \to I(b)$ and that $I(b + (1/k)^*) \to I(b)$.

To show that $I(b - (1/k)^*) \to I(b)$, we must show that for all $\epsilon > 0$, there exists k such that $I(b - (1/k)^*) \ge I(b) - \epsilon$. By positive homogeneity (part (b)), we can assume without loss of generality that $||b - (1/2)^*|| \le 1$ and that $||b|| \le 1$. Fix $\epsilon > 0$. If $I(b - (1/2)^*) \ge I(b) - \epsilon$, then we are done. If not, then $I(b) > I(b) - \epsilon > I(b - (1/2)^*)$. Since $||b|| \le 1$ and $||b - (1/2)^*|| \le 1$, f_b and $f_{b-(1/2)^*}$ exist. Moreover, by Lemma B.3.4, $f_b \succ f_{(I(b)-\epsilon)^*} \succ f_{b-(1/2)^*}$. By mixture continuity, for some $p \in (0,1)$, we have $pf_b + (1-p)f_{(b-(1/2)^*} \succ f_{(I(b)-\epsilon)^*}$. It is easy to check that $b_{pf_b+(1-p)f_{b-(1/2)^*}} = b - ((1-p)/2)^*$. Thus, by Lemma B.3.4, $f_{b-((1-p)/2)^*} \succeq f_{(I(b)-\epsilon)^*}$, and $I(b - ((1-p)/2)^*) > I(b) - \epsilon$. Choose k such that $1/k < (1-p)/2$. Then, by monotonicity (part (c)), $I(b - (1/k)^*) \ge I(b - ((1-p)/2)^*) > I(b) - \epsilon$, as desired.

The argument that $I(b + (1/k)^*) \to I(b)$ is similar and left to the reader.

For part (e), if $||b||, ||b'|| \le 1$, and $I(b), I(b') \ne 0$, consider $\frac{b}{I(b)}$ and $\frac{b'}{I(b')}$. Since $I(\frac{b}{I(b)}) = I(\frac{b'}{I(b')}) = 1$, it follows from Lemma B.3.3 that $f_{\frac{b}{I(b)}} \sim f_{\frac{b'}{I(b')}}$. By Ambiguity Aversion, for all $p \in (0,1]$, $pf_{\frac{b}{I(b)}} + (1-p)f_{\frac{b'}{I(b')}} \succeq f_{\frac{b}{I(b)}}$. Thus, taking $p = \frac{I(b)}{I(b)+I(b')}$, $I(\frac{b+b'}{I(b)+I(b')}) = \frac{1}{I(b)+I(b')}I(b+b') = I(\frac{I(b)}{I(b)+I(b')}\frac{b}{I(b)} + \frac{I(b')}{I(b)+I(b')}\frac{b'}{I(b')}) \ge I(\frac{b}{I(b)}) = I(\frac{b'}{I(b')}) = 1$. Hence, $I(b+b') \ge I(b) + I(b')$.

If either $||b|| > 1$ or $||b'|| > 1$, and both $I(b) \ne 0$ and $I(b') \ne 0$, then the result easily follows by positive homogeneity (property (b)).

If either $I(b) = 0$ or $I(b') = 0$, let $b_n = b + \frac{1}{n}^*$ and $b'_n = b' + \frac{1}{n}^*$. Clearly $||b_n|| > 0$, $||b'_n|| > 0$, $b_n \to b$, and $b'_n \to b'_n$. By our argument above, $I(b_n + b'_n) \geq I(b_n) + I(b'_n)$ for all $n \geq 1$. The result now follows from continuity.

□

B.3.2 Defining the confidence function

In this section, we use I to define a confidence function ϕ that maps each $p \in \Delta(S)$ to a confidence value in $[0, 1]$. The heart of the proof involves showing that the resulting function ϕ so determined gives us the desired representation.

Given a confidence function ϕ, for $b \in \mathcal{B}^+$, define

$$WE(b) = \inf_{p \in \mathcal{P}} \phi(p)(\sum_{s \in S} b(s)p(s)).$$

Define

$$\underline{E}(b) = \inf_{p \in \mathcal{P}} \sum_{s \in S} b(s)p(s).$$

and

$$E_p(b) = \sum_{s \in S} b(s)p(s).$$

For each probability $p \in \Delta(S)$, define

$$\phi_t(p) = \inf\{\alpha \in \mathbb{R} : I(b) \leq \alpha E_p(b) \text{ for all } b \in \mathcal{B}^+\}, \tag{B.1}$$

and let $\phi_t(p) = \infty$ if the $\inf$ does not exist. Note that $\phi_t(p) \geq 1$, since $E_p((c)^*) = I((c)^*) = c$ for all distributions p and $c \in \mathbb{R}$. Moreover, it is immediate from the definition of $\phi_t(p)$ that $\phi_t(p)E_p(b) \geq I(b)$ for all $b \in \mathcal{B}^+$. The next lemma shows that there exists a probability p where we have equality.

Lemma B.3.6. *(a) For some distribution p, we have $\phi_t(p) = 1$.*

(b) For all $b \in \mathcal{B}^+$, there exists p such that $\phi_t(p)E_p(b) = I(b)$.

Proof. The proofs of both part (a) and (b) use a separating hyperplane theorem. If U is a convex subset of $\mathcal{B}^+$, and $b \notin U$, then there is a linear functional λ that separates U from b, that is, $\lambda(b') < \lambda(b)$ for all $b' \in U$. We proceed as follows.

For part (a), we must show that there exists a probability measure p such that for all $b \in \mathcal{B}^+$, we have $E_p(b) \geq I(b)$. This would show that $\phi_t(p) = 1$.

Let $U = \{b' \in \mathcal{B}^+ : I(b') \geq 1\}$. U is closed (by continuity of I) and convex (by positive homogeneity and superadditivity of I), and $(0)^* \notin U$. Thus, there exists a linear functional λ such that $\lambda(b') > \lambda((0)^*) = 0$ for $b' \in U$. We can assume without lost of generality that $\lambda(1^*) = 1$.

We want to show that λ is a *positive* linear functional, that is, that $\lambda(b) \geq 0$ if $b \geq 0^*$. Clearly this holds for b' such that $I(b') \geq 1$. If $b' \geq 0^*$, $I(b') < 1$, and $I(b') > 0$, note that $cI(b') = I(cb') \geq 1$ for some $c \geq 0$. Therefore, $I(b') \geq \frac{1}{c} \geq 0$. If $b' \geq 0^*$ and $I(b') = 0$, note that for all $c > 0$, $\lambda(b' + c^*) \geq 0$ by the previous case. Thus, $\lambda(b') \geq 0$. It follows that λ is a positive functional.

Define the probability distribution p on S by taking $p(s) = \lambda(1_s)$. To see that p is indeed a probability distribution, note that since $1_s \geq 0$ and λ is positive, we must have $\lambda(1_s) \geq 0$. Moreover, $\sum_{s \in S} p(s) = \lambda(1^*) = 1$. In addition, for all $b' \in \mathcal{B}$, we have

$$\lambda(b') = \sum_{s \in S} \lambda(1_s) b'(s) = \sum_{s \in S} p(s) b'(s) = E_p(b').$$

Next, we claim that, for $b \in \mathcal{B}^+$,

$$\text{for all } c > 0, \text{ if } I(b) > c, \text{ then } \lambda(b) > c. \tag{B.2}$$

To see why the claim is true, note that if $I(b) \geq c$, then $I(b/c) \geq 1$ by positive homogeneity, so $\lambda(b/c) \geq 1$ and $\lambda(b) \geq c$. Therefore, $\lambda(b) \geq I(b)$, as desired.

The proof of part (b) is similar to that of part (a). We want to show that, given $b \in \mathcal{B}^+$, there exists p such that $\phi_t(p)E_p(b) = I(b)$. First consider the case where $||b|| \leq 1$. If $I(b) = 0$, then there must exist some s such that $b(s) = 0$, for otherwise there exists $c > 0$ such that $b \geq c^*$, so $I(b) \geq c$. If $b(s) = 0$, let p_s be such that $p_s(s) = 1$. Then $E_{p_s}(b) = 0$, so part (b) of the Lemma holds in this case.

If $||b|| \leq 1$ and $I(b) > 0$, let $U = \{b' : I(b') \geq I(b)\}$. Again, U is closed and convex, and $b \notin U$, so there exists a linear functional λ such that $\lambda(b') > \lambda(b)$ for $b' \in U$. Since $1^* \in U$ and we can assume without loss of generality $\lambda(1^*) = 1$, we must have $\lambda(b) < 1$.

The same argument as that used in the proof of (a) shows that λ is a positive functional.

Therefore, λ determines a probability distribution p such that, for all $b' \in \mathcal{B}^+$, we have $\lambda(b') = E_p(b')$. p, of course, will turn out to be the desired distribution. To show this, we need to show that $\phi_t(p) = I(b)/E_p(b)$. By definition, $\phi_t(p) \geq I(b)/E_p(b)$. To show that $\phi_t(p) \leq I(b)/E_p b$, we must show that $\frac{I(b)}{E_p(b)} \geq \frac{I(b')}{E_p b'}$ for all $b' \in \mathcal{B}^+$. Equivalently, we must show that $I(b)\lambda(b')/\lambda(b) \geq I(b')$ for all $b' \in \mathcal{B}^+$.

Essentially the same argument used to prove (B.2) also shows that

$$\text{for all } c > 0, \text{ if } \tfrac{I(b')}{I(b)} \geq c, \text{ then } \tfrac{\lambda(b')}{\lambda(b)} \geq c.$$

In particular, if $\frac{I(b')}{I(b)} \geq c$, then by positive homogeneity, $\frac{I(b')}{c} \geq I(b)$, so $\frac{b'}{c} \in U$, and $\lambda(\frac{b'}{c}) > \lambda(b)$ and hence $\frac{\lambda(b')}{\lambda(b)} \geq c$.

It follows that $\lambda(b')/(\lambda(b)) \geq I(b')/(I(b))$ for all $b' \in \mathcal{B}^+$. Thus, $I(b)\lambda(b')/\lambda(b) \geq I(b')$ for all $b' \in \mathcal{B}^+$, as required.

Finally, if $||b|| > 1$, let $b' = b/||b||$. By the argument above, there exists a probability measure p such that $\phi_t(p)E_p(b/||b||) = I(b/||b||)$. Since $E_p(b/||b||) = E_p(b)/||b||$, and $I(b/||b||) = I(b)/||b||$, we must have that $\phi_t(p)E_p(b) = I(b)$. □

We can now complete the proof of Theorem 4.5.3. By Lemma B.3.6 and the definition of $\phi_t(p)$, for all $b \in \mathcal{B}^+$,

$$I(b) = \inf_{p \in \Delta(S)} \phi_t(p)E_p(b). \tag{B.3}$$

Recall that, by Lemma B.3.4, for all acts f, g such that $b_f, b_g \in \mathcal{B}^+$, $f \succeq g$ iff $I(b_f) \geq I(b_g)$. Thus, $f \succeq g$ iff

$$\inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(f(s))p(s) \right) \geq \inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(g(s))p(s) \right).$$

To get the confidence function ϕ from ϕ_t, note that $\lim_{x \to 0^+} t(x) = \infty$ and $t(1) > 0$. We let $\phi(p) = t^{-1}(t(1)\phi_t(p))$, with the special case $\phi(p) = 0$ if $\phi_t(p) = \infty$. (Note that $t(1)\phi_t(p)$ is in the range of t^{-1}, since $\phi_t(p) \geq 1$, t is nonincreasing, and $\lim_{x \to 0^+} t(x) = \infty$.)

B.3.3 Properties of the confidence function

In this section, we show that the confidence function ϕ that we constructed satisfies the properties claimed in Theorem 4.5.3.

We first show that $t \circ \phi = \phi_t$ has convex upper support. To that end, we show that if $c_1 \geq \phi_t(p_1)$ and $c_2 \geq \phi_t(p_2)$, then for all $\alpha \in (0,1)$,

$$(\alpha c_1 p_1 + (1-\alpha)c_2 p_2)(S) \geq \phi_t\left(\frac{\alpha c_1 p_1 + (1-\alpha)c_2 p_2}{(\alpha c_1 p_1 + (1-\alpha)c_2 p_2)(S)}\right).$$

By the definition of ϕ_t, it suffices to show that for all $b \in \mathcal{B}^+$,

$$I(b) \leq (\alpha c_1 p_1 + (1-\alpha)c_2 p_2)(S) E_{\frac{\alpha c_1 p_1 + (1-\alpha)c_2 p_2}{(\alpha c_1 p_1 + (1-\alpha)c_2 p_2)(S)}}(b). \tag{B.4}$$

It is easy to see that the inequality holds. Let $b \in \mathcal{B}^+$. The right-hand side of (B.4) is equal to

$$\begin{aligned}
\sum_{s \in S} ((\alpha c_1 p_1(s) + (1-\alpha)c_2 p_2(s))b(s)) &= \alpha c_1 E_{p_1}(b) + (1-\alpha)c_2 E_{p_2}(b) \\
&\geq \alpha \phi_t(p_1) E_{p_1}(b) + (1-\alpha)\phi_t(p_2) E_{p_2}(b) \\
&\geq \alpha I(b) + (1-\alpha) I(b) \text{ (by (B.3))} \\
&\geq I(b).
\end{aligned}$$

We now show that ϕ is regular*. Since we've shown that, for some p^*, $\phi_t(p^*) = 1$, we have $\phi(p^*) = t^{-1}(t(1)1) = 1$. Therefore ϕ is normal.

Secondly, we show that ϕ is weakly* upper semicontinuous. We show that if $\{p_n\} \to p$ and $\phi(p_n) \geq \alpha$ for all n, then $\phi(p) \geq \alpha$. Suppose for the purpose of contradiction that $\phi(p) < \alpha$. Then $\phi_t(p) = t(\phi(p)) > t(\alpha)$. By continuity of t, $\phi_t(p_n) = t(\phi(p_n)) > t(\alpha)$ for all sufficiently large n, implying that $\phi(p_n) < \alpha$, contradicting the assumption that $\phi(p_n) \geq \alpha$. Therefore $\phi(p) \geq \alpha$, as required.

We now show that ϕ is quasiconcave; that is, $\phi(\beta p_1 + (1-\beta)p_2) \geq \min\{\phi(p_1), \phi(p_2)\}$ for any $\beta \in [0,1]$. Since t is strictly decreasing, so is t^{-1}. Thus, $-t^{-1}$ is strictly increasing. Moreover, if ϕ_t is quasiconvex then $-t^{-1} \circ \phi_t$ is also quasiconvex. Since the negative of a quasiconvex function is quasiconcave,

$t^{-1} \circ \phi_t$ is quasiconcave. Therefore, if we show that ϕ_t is quasiconvex, this would show that $\phi = t^{-1} \circ \phi_t$ is quasiconcave.

Recall from (B.1) that

$$\phi_t(p) = \inf\{\alpha \in \mathbb{R} : I(b) \leq \alpha E_p(b) \text{ for all } b \in \mathcal{B}^+\}.$$

If $\max\{\phi_t(p_1), \phi_t(p_2)\} \leq c$ for $c \in \mathbb{R}$, then for all $b \in \mathcal{B}^+$, we have

$$I(b) \leq cE_{p_1}(b),$$

and

$$I(b) \leq cE_{p_2}(b).$$

Therefore, for all $b \in \mathcal{B}^+$ and all $\beta \in [0,1]$, by the linearity of $E_p(b)$ with respect to the parameter p,

$$I(b) \leq cE_{\beta p_1 + (1-\beta)p_2}(b).$$

This means that $\phi_t(\beta p_1 + (1-\beta)p_2) \leq c$. Thus, $\phi_t(\beta p_1 + (1-\beta)p_2) \leq \max\{\phi_t(p_1), \phi_t(p_2)\}$. Therefore, ϕ_t is quasiconvex.

B.3.4 Uniqueness of the representation

In this section, we show that our constructed ϕ is the only regular* fuzzy confidence function such that $t \circ \phi$ has convex upper support, and such that $\succeq^+_{t,\phi} = \succeq$. Our uniqueness result is similar in spirit to the uniqueness results of Gilboa and Schmeidler [Gilboa and Schmeidler 1989b], who show that the convex, closed, and non-empty set of probability measures in their representation theorem for MMEU is unique.

The proof of this result, like the proof of uniqueness in Gilboa and Schmeidler [Gilboa and Schmeidler 1989b], uses a separating hyperplane theorem to show the existence of acts on which two different representations must 'disagree'. The proof presented here is essentially the same as that used in [Halpern and Leung 2012], with only superficial changes to accommodate our definitions and notation.

Lemma B.3.7. *For all confidence functions ϕ', if $\succeq^+_{t,\phi'} = \succeq$ and $t \circ \phi'$ has convex upper support, then $\phi = \phi'$.*

Proof. Suppose for contradiction that there exists a regular* fuzzy confidence function $\phi' \neq \phi$ such that $t \circ \phi'$ has convex upper support, and that $\succeq^+_{t,\phi'} = \succeq^+_{t,\phi}$. Consider the two upper supports $\overline{V}_{t\circ\phi}$ and $\overline{V}_{t\circ\phi'}$. $\overline{V}_{t\circ\phi}$ and $\overline{V}_{t\circ\phi'}$ are both closed. To see why, consider a sequence $\{p_n\}_{n\in\mathbb{N}}$ contained in $p_n \in \overline{V}_{t\circ\phi}$ such that $p_n \to p$. We show that $p \in \overline{V}_{t\circ\phi}$, by showing that for some $q \in \Delta(S)$, $\phi(q) > 0$ and $p \geq t(\phi(q))q$.

We first show that $p \geq t(\phi(q))q$ for some $q \in \Delta(S)$. Recall that for all n, there exists $q_n \in \Delta(S)$ such that $p_n \geq t(\phi(q_n))q_n$. Since $q_n \in \Delta(S)$, $q_{k_m} \to q$ for some subsequence $\{q_{k_m}\}$ and $q \in \Delta(S)$. Therefore, we have

$$\begin{aligned}
p &= \lim_{n\to\infty} p_n \\
&\geq \limsup_{n\to\infty} t(\phi(q_n))q_n, \text{ since } p_n \geq t(\phi(q_n))q_n \\
&= \lim_{n\to\infty} \sup_{m\geq n} t(\phi(q_{k_m}))q_{k_m} \\
&= \lim_{n\to\infty} \sup_{m\geq n} t(\phi(q_{k_m})) \lim_{m\to\infty} q_{k_m} \\
&= \lim_{n\to\infty} t(\inf_{m\geq n} \phi(q_{k_m})) \lim_{m\to\infty} q_{k_m}, \text{ since } t \text{ is nonincreasing and continuous} \\
&= t(\liminf_{m\to\infty} \phi(q_{k_m})) \lim_{m\to\infty} q_{k_m}, \text{ by continuity of } t \\
&\geq t(\phi(q))q,
\end{aligned}$$

since $\phi(q) \geq \limsup_{m\to\infty} \phi(q_{k_m}) \geq \liminf_{m\to\infty} \phi(q_{k_m})$ by upper semicontinuity of ϕ, and t is nonincreasing.

It remains to show that $\phi(q) > 0$. To that end, suppose for the purpose of contradiction that $\phi(q) = 0$. Then it must be the case that $\lim_{m\to\infty} \phi(q_{k_m}) = 0$, since if there exists an $\epsilon > 0$ such that $\lim_{m\to\infty} \phi(q_{k_m}) \geq \epsilon$, then by upper semicontinuity of ϕ it must be the case that $\phi(q) \geq \epsilon$. Since $\lim_{x\to 0^+} t(x) = \infty$, we have that $\lim_{m\to\infty} t(\phi(q_{k_m})) = \infty$. However, recall that $p_n \geq t(\phi(q_n))q_n$ for all n. Since $q_n \in \Delta(S)$ and hence does not vanish, p_n cannot be a convergent sequence. Hence it must be the case that $\phi(q) > 0$.

Therefore, $p \in \overline{V}_{t\circ\phi}$, as required, and that $\overline{V}_{t\circ\phi}$ is closed. The same argument shows that $\overline{V}_{t\circ\phi'}$ is closed.

Without loss of generality, let $q \in \overline{V}_{t\circ\phi'} \backslash \overline{V}_{t\circ\phi}$. Since $\overline{V}_{t\circ\phi}$ and $\{q\}$ are closed, convex, and disjoint, and $\{q\}$ is compact, the separating hyperplane theorem [Rockafellar 1970] says that there exists $\theta \in \mathbb{R}^{|S|}$ and $c \in \mathbb{R}$ such that

$$\theta \cdot p > c \text{ for all } p \in \overline{V}_{t\circ\phi}, \text{ and } \theta \cdot q < c. \tag{B.5}$$

By scaling c appropriately, we can assume that $|\theta(s)| \leq 1$ for all $s \in S$. Now we argue that it must be the case that $\theta(s) \geq 0$ for all $s \in S$ (so that θ corresponds to the utility profile of some act with nonnegative utilities). Suppose that $\theta(s') < 0$ for some $s' \in S$. By (C.4), $\theta \cdot p > c$ for all $p \in \overline{V}_{t\circ\phi}$. Let $p^* \in \overline{V}_{t\circ\phi}$ be any measure with $\phi(p^*) = 1$, and let $p^{**} \in \overline{V}_{t\circ\phi}$ be defined by

$$p^{**}(s) = \begin{cases} p^*(s), \text{ if } s \neq s' \\ \frac{|S| \max\{|c|, \max_{s''\in S} |p^*(s'')|\}}{|\theta(s')|}, \text{ if } s = s'. \end{cases}$$

We have defined p^{**} such that $p^{**} \geq p^*$, since for all $s \in S$, $p^{**}(s) \geq p^*(s)$. To see how, note that $p^{**}(s) = p^*(s)$ for $s \neq s'$, and $p^{**}(s) \geq \max_{s''\in S} |p^*(s'')| \geq p^*(s)$ for

$s = s'$. Therefore, p^{**} is in $\overline{V}_{t\circ\phi}$.

Our definition of p^{**} also ensures that $\theta \cdot p^{**} = \sum_{s\in S} p^{**}(s)\theta(s) \le c$, since

$$\begin{aligned}\sum_{s\in S} p^{**}(s)\theta(s) =& p^{**}(s')\theta(s') + \sum_{s\neq s'} p^{**}(s)\theta(s)\\ \le& p^{**}(s')\theta(s') + \sum_{s\neq s'} |p^{**}(s)|, \text{ since } |\theta(s)| \le 1\\ =& -|S| \max\{|c|, \max_{s''\in S} |p^*(s'')|\} + \sum_{s\neq s'} |p^{**}(s)|\\ \le& -|c| \le c.\end{aligned}$$

This contradicts (C.4), which says that $\theta \cdot p > c$ for all $p \in \overline{V}_{t\circ\phi}$. Thus it must be the case that $\theta(s) \ge 0$ for all $s \in S$.

Consider the θ given by the separating hyperplane theorem, and let f be an act such that $u \circ f = \theta$. $f \sim l_d^*$ for some constant act l_d^*. Since $\overline{V}_{t\circ\phi}$ and $\overline{V}_{t\circ\phi'}$ as sets of generalized probabilities both represent $\succeq$, and $\overline{V}_{t\circ\phi}$ and $\overline{V}_{t\circ\phi'}$ both contain a normal probability measure,

$$\min_{p\in\overline{V}_{t\circ\phi}} p \cdot (u \circ f) = \min_{p\in\overline{V}_{t\circ\phi}} p \cdot (u \circ l_d^*) = d = \min_{p\in\overline{V}_{t\circ\phi'}} p \cdot (u \circ f).$$

However, by (C.4),

$$\min_{p\in\overline{V}_{t\circ\phi}} p \cdot (u \circ f) > c > \min_{p\in\overline{V}_{t\circ\phi'}} p \cdot (u \circ f),$$

which is a contradiction. □

B.4 Proof of Theorem 3.6.7

Proof. The proof is almost the same as the proof of Theorem 4.5.3. We point out the differences, which are mostly straightforward adaptations from $\mathcal{B}^+$ to

$\mathcal{B}^-$. Lemma B.3.1 and Lemma B.3.2 hold without change. By Axiom 3.8, we can assume that the maximum value that u takes on is 0, and by Axiom 3.1 we can assume that the minimum is no greater than -1.

We now define a functional I on utility acts, as before. All occurrences of $\mathcal{B}^+$ in the proof of Theorem 4.5.3 needs to be replaced by $\mathcal{B}^-$, defined by the real-valued functions b on S where $b(s) \leq 0$ for all $s \in S$.

More specifically, let

$$R_f = \{\alpha' : l^*_{\alpha'} \preceq f\}.$$

If $0^* \geq b \geq (-1)^*$, then f_b exists, and we define

$$I(b) = \sup(R_{f_b}).$$

For the remaining utility acts $b \in \mathcal{B}^+$, we extend I by homogeneity, as before.

The analog of Lemma B.3.3 for $b_f \in \mathcal{B}^-$ follows from analogous arguments used in the original proof. The case of $l^*_{I(b_f)} \prec f$, however, is a bit simpler than for the positive case.

Lemma B.4.1. *If $b_f \in \mathcal{B}^-$, then $f \sim l^*_{I(b_f)}$.*

Proof. Suppose, by way of contradiction, that $l^*_{I(b_f)} \prec f$. If $f \sim l^*_0$, then $I(b_f) \geq 0$ by the definition of I. However, we also have $I(b_f) \leq 0$ by Lemma B.3.4, so $I(b_f) = 0$, and therefore $f \sim l^*_{I(b_f)}$, as required. Otherwise, $f \prec l^*_0$ by monotonicity, so $l^*_{I(b_f)} \prec f \prec l^*_0$, which, when taken together with mixture continuity, contradicts the definition of I. $\square$

The proof of Lemma B.3.4 still holds. The analog of Lemma B.3.5 also follows from simular arguments; we discuss some key differences below.

Lemma B.4.2. *(a) If $c \leq 0$, then $I(c^*) = c$.*

(b) I satisfies positive homogeneity: if $b \in \mathcal{B}^-$ and $c > 0$, then $I(cb) = cI(b)$.

(c) I is monotonic: if $b, b' \in \mathcal{B}^-$ and $b \geq b'$, then $I(b) \geq I(b')$.

(d) I is continuous: if $b, b_1, b_2, \ldots \in \mathcal{B}^-$, and $b_n \to b$, then $I(b_n) \to I(b)$.

(e) I is superadditive: if $b, b' \in \mathcal{B}^-$, then $I(b + b') \geq I(b) + I(b')$.

Proof. For part (b), instead of making use of Axiom 3.4 (worst independence), we use Axiom 3.8 (best independence).

For part (e), note that since $I(b)$ is nonpositive for $b \in \mathcal{B}^-$, $I(\frac{b}{I(b)})$ is not defined, unlike in the case of nonnegative utilities. We use the same proof as in [Halpern and Leung 2012]: Clearly, $I(\frac{b}{-I(b)}) = -1$. Therefore, $f_{\frac{b}{-I(b)}} \sim f_{\frac{b'}{-I(b')}} \sim l^*_{-1}$. From Axiom 5.9 (ambiguity aversion), taking $p = \frac{-I(b)}{-I(b)-I(b')}$, we have

$$I\left(\frac{-I(b)}{-I(b) - I(b')}\frac{b}{-I(b)} + \frac{-I(b')}{-I(b) - I(b')}\frac{b'}{-I(b')}\right) \geq I(\frac{b}{-I(b)}) = -1,$$

which implies that $I(b + b') \geq I(b) + I(b')$, as required. □

We now use I to define a confidence function ϕ. WE, $\underline{E}$, and E are defined as before. For each probability $p \in \Delta(S)$, define

$$\phi_t(p) = \sup\{\alpha \in \mathbb{R} : I(b) \leq \alpha E_p(b) \text{ for all } b \in \mathcal{B}^-\}.$$

Note that $\phi_t(p) \leq 1$, since $E_p((c)^*) = I((c)^*) = c$ for all distributions p and $c \in \mathbb{R}$. Moreover, $\phi_t(p) \geq 0$ for all $b \in \mathcal{B}^-$. The next lemma shows that there exists a probability p where we have equality. The proof of the lemma is similar to that of Lemma B.3.6, and is left to the reader.

Lemma B.4.3. *(a) For some distribution p, we have $\phi_t(p) = 1$.*

(b) For all $b \in \mathcal{B}^-$, there exists p such that $\phi_t(p)E_p(b) = I(b)$.

By Lemma B.4.3 and the definition of $\phi_t(p)$, for all $b \in \mathcal{B}^-$,

$$I(b) = \inf_{p \in \Delta(S)} \phi_t(p)E_p(b).$$

We have $f \succeq g$

$$\text{iff } \inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(f(s))p(s) \right) \geq \inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(g(s))p(s) \right)$$

$$\text{iff } t(1) \inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(f(s))p(s) \right) \geq t(1) \inf_{p \in \Delta(S)} \left(\phi_t(p) \sum_{s \in S} u(g(s))p(s) \right).$$

Since t is strictly increasing, $t(1) > t(0)$. Therefore, since $\phi_t(p) \in [0,1]$ and $t(0) \leq 0$, $t(1)\phi_t(p)$ is in the range of t, and we can define

$$\phi(p) = t^{-1}(t(1)\phi_t(p)).$$

We now have $f \succeq g$

$$\text{iff } \inf_{p \in \Delta(S)} \left(t(\phi(p)) \sum_{s \in S} u(f(s))p(s) \right) \geq \inf_{p \in \Delta(S)} \left(t(\phi(p)) \sum_{s \in S} u(g(s))p(s) \right).$$

Finally, uniqueness of the representation follows from arguments analogous to those for nonnegative utilities. □

APPENDIX C

PROOFS FOR CHAPTER 4

C.1 Proof of Theorem 4.5.1

We restate the theorem (and elsewhere in the appendix) for the reader's convenience.

THEOREM 4.5.1. *For a dynamic decision problem D, if and $\mu(h) = M$ for some fixed menu M, then there will be no preference reversals in D.*

Proof. Before proving the result, we need some definitions. Say that an information set I *refines* an information set I' if, for all $h \in I$, some prefix h' of h is in I'. Suppose that there is a history h such that $f, g \in M_h$ and $I(h) = I$. Let fIg denote the plan that agrees with f at all histories h' such that $I(h')$ refines I and agrees with g otherwise. As we now show, fIg gives the same outcome as f on states in $E = E(h)$ and the same outcome as g on states in E^c; moreover, $fIg \in M_h$.

Suppose that $s(h) = s$ and that $s \in E$. Since $E(h) = E$, there exists a history $h' \in I(h)$ such that $s(h') = s'$ and $R(h') = R(h)$. Since $f, g \in M_h$, there must exist some k such that $f^k(\langle s \rangle) = g^k(\langle s \rangle) = h$ (where, as usual, $f^0(\langle s \rangle) = \langle s \rangle$ and for $k' \geq 1$, $f^{k'}(\langle s \rangle) = f(f^{k'-1}(\langle s \rangle)))$. We claim that for all $k' \leq k$, $f^{k'}(\langle s' \rangle) = g^{k'}(\langle s' \rangle)$, and $f^{k'}(\langle s' \rangle)$ is in the same information set as $f^{k'}(\langle s \rangle)$. The proof is by induction on k'. If $k' = 0$, the result follows from the observation that since $\langle s \rangle$ is a prefix of h, there must be some prefix of h' in $I(\langle s \rangle)$. For the inductive step, suppose that $k' \geq 1$. We must have $f^{k'}(\langle s \rangle) = g^{k'}(\langle s \rangle)$ (otherwise g would not be in

M_h). Since $g^{k'-1}(\langle s\rangle) = f^{k'-1}(\langle s\rangle)$ and $f^{k'-1}(\langle s'\rangle) = g^{k'-1}(\langle s'\rangle)$ are in the same information set, by the inductive hypothesis, g must perform the same action at $g^{k'-1}(\langle s\rangle)$ and $g^{k'-1}(\langle s'\rangle)$, and must perform the same action at $f^{k'-1}(\langle s\rangle)$ and $f^{k'-1}(\langle s'\rangle)$. Since $g^{k'}(\langle s\rangle)$ and $f^{k'}(\langle s\rangle)$ are both prefixes of h, g and f perform the same action at $f^{k'-1}(\langle s\rangle) = g^{k'-1}(\langle s\rangle)$. It follows that f and g perform the same action at $f^{k'-1}(\langle s'\rangle) = g^{k'-1}(\langle s'\rangle)$, and so $f^{k'}(\langle s'\rangle) = g^{k'}(\langle s'\rangle)$. Thus, $g^{k'}(\langle s'\rangle)$ must be a prefix of h', and so must be in the same information set as $f^{k'}(\langle s\rangle)$. This completes the inductive proof.

Since $f^k(\langle s'\rangle) = g^k(\langle s'\rangle) = h'$, it follows that $f^k(\langle s'\rangle) = (fIg)^k(\langle s'\rangle)$. Below I, all the information sets are refinements of I, so by definition, for $k' \leq k$, we must $f^{k'}(\langle s'\rangle) = (fIg)^{k'}(\langle s'\rangle)$. Thus, f and fIg give the same outcome for s', and hence all states in E. Note it follows that $(fIg)^k(\langle s\rangle) = h$, so $fIg \in M_h$.

For $s' \notin E$ and all k', it cannot be the case that $I((fIg)^{k'}(\langle s'\rangle))$ is a refinement of I, since the first state in $(fIg)^{k'}(\langle s'\rangle))$ is s', and no history in a refinement of I has a first state of s'. Thus, $fIg^{k'}(\langle s'\rangle) = g^{k'}(\langle s'\rangle)$ for all k', so f and fIg give the same outcome for s', and hence all states in E^c.

Returning to the proof of the proposition, suppose that $f \in C_{\mu,h}(M_h)$, h' is a history extending h, and $f \in M_{h'}$. We want to show that $f \in C_{\mu,h'}(M_{h'})$. By perfect recall, $E(h') \subseteq E(h)$. Suppose, by way of contradiction, that $f \notin C_{\mu,h'}(M_{h'})$. Since $f \in C_{\mu,h'}(M_{h'})$, we cannot have $E(h') = E(h)$, so $E(h') \subset E(h)$. Choose $f' \in C_{\mu,E(h')}(M_{h'})$ and $g \in C_{\mu,E(h')^c \cap E(h)}(M_{h'})$ (note that $C_{\mu,E(h')}(M_{h'}) \neq \emptyset$ and $C_{\mu,E(h')^c \cap E(h)}(M_{h'}) \neq \emptyset$ by Axiom 3). Since $f', g \in M_{h'}$ (by Axiom 3), $f'I(h')g$ is in $M_{h'}$. Since $f'I(h')g$ and f', when viewed as acts, agree on states in $E(h')$, we must have $f'I(h')g \in C_{\mu,E(h')}(M_{h'})$ by Axiom 4.2. Similarly, since $f'I(h')g$ and g, when viewed as acts, agree on states in $E(h')^c \cap E(h)$, we must have

$f'I(h')g \in C_{\mu,E(h')^c \cap E(h)}(M_{h'})$. Therefore, by Axiom 4.1, $f'I(h')g \in C_{\mu,h}(M_{h'})$. Also by Axiom 4.1, since $f \notin C_{\mu,h'}(M_{h'})$, we must have $f \notin C_{\mu,h}(M_{h'})$. By Axiom 4.4, this implies that $f \notin C_{\mu,h}(M_h)$ (since $M_{h'} \subseteq M_h$), giving us the desired contradiction. $\square$

C.2 Proof of Theorem 4.5.3

THEOREM 4.5.3. *If $\mathcal{P}^+$ is a set of weighted distributions on (S, Σ) such that $C(\mathcal{P}^+)$ is closed, then the following are equivalent:*

(a) For all decision problems D based on (S, Σ) and all menus M in D, Axioms 4.1–4.4 hold for choice functions represented by $\mathcal{P}^+|^l E$ (resp., $\mathcal{P}^+|^p E$).

(b) For all decision problems D based on (S, Σ), states $s \in S$, and acts $f \in M_{\langle s \rangle}$, the weighted regret of f with respect to $M_{\langle s \rangle}$ and $\mathcal{P}^+$ is separable.

We actually prove the following stronger result.

Theorem C.2.1. *If $\mathcal{P}^+$ is a set of weighted distributions on (S, Σ) such that $C(\mathcal{P}^+)$ is closed, then the following are equivalent:*

(a) For all decision problems D based on (S, Σ), Axioms 4.1–4.4 hold for menus of the form $M_{\langle s \rangle}$ for choice functions represented by $\mathcal{P}^+|^l E$ (resp., $\mathcal{P}^+|^p E$).

(b) For all decision problems D based on (S, Σ) and all menus M in D, Axioms 4.1–4.4 hold for choice functions represented by $\mathcal{P}^+|^l E$ (resp., $\mathcal{P}^+|^p E$).

(c) For all decision problems D based on (S, Σ), states $s \in S$, and acts $f \in M_{\langle s \rangle}$, the weighted regret of f with respect to $M_{\langle s \rangle}$ and $\mathcal{P}^+$ is separable.

(d) *For all decision problems D based on (S, Σ), menus M in D, and acts $f \in M$, the weighted regret of f with respect to M and $\mathcal{P}^+$ is separable.*

Proof. Fix an arbitrary state space S, measurable events $E, F \subseteq S$, and a set $\mathcal{P}^+$ of weighted distributions on (S, Σ). The fact that (b) implies (a) and (d) implies (c) follows immediately. Therefore, it remains to show that (a) implies (d) and that (c) implies (b).

Since the proof is identical for prior-by-prior updating ($|^p$) and for likelihood updating ($|^l$), we use $|$ to denote the updating operator. That is, the proof can be read with $|$ denoting $|^p$, or with $|$ denoting $|^l$.

To show that (a) implies (d), we first show that Axiom 4.1 implies that for all decision problems D based on (S, Σ), menu M in D, sets $\mathcal{P}^+$ of weighted probabilities, and acts $f \in M$,

$$reg_M^{\mathcal{P}^+|F}(f) \geq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f) \right). \tag{C.1}$$

Suppose, by way of contradiction, that (C.1) does not hold. Then for some decision problem D based on (S, Σ), measurable events $E, F \subseteq S$, menu M in D, and act $f \in M$, we have that

$$reg_M^{\mathcal{P}^+|F}(f) < \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c \cap F) reg_{M,F}^{\mathcal{P}^+|(E^c\cap F)}(f) \right).$$

We define a new decision problem D' based on (S, Σ). The idea is that in D', we will have a plan $a_{f'}$ such that $a_{f'} \in C_{M',E\cap F}^{reg,\mathcal{P}^+}(M'')$ and $a_{f'} \in C_{M',E^c\cap F}^{reg,\mathcal{P}^+}(M'')$ and $a_{f'} \notin C_{M',F}^{reg,\mathcal{P}^+}(M'')$ for some $M'' \subseteq M'$, where M' is the menu at the initial decision node for the DM.

We construct D' as follows. D' is a depth-two tree; that is, nature makes a single move, and then the DM makes a single move. At the first step, nature choose a state $s \in F$. At the second step, the DM chooses from the set $\{a_g : g \in M\} \cup \{a_{f'}\}$ of actions. With a slight abuse of notation, we let a_g also denote the plan in T' that chooses the action a_g at the initial history $\langle s \rangle$. Therefore, the initial menu in decision problem D' is $M' = \{a_g : g \in M\} \cup \{a_{f'}\}$.

The utilities for the actions/plans in D' are defined as follows. For actions $\{a_g : g \in M\}$, the utility of a_g in state s is just the utility of the outcome resulting from applying plan g in state s in decision problem D. The action $a_{f'}$ has utilities

$$u(a_{f'}(s)) = \begin{cases} \sup_{g \in M} u(g(s)) - reg_M^{\mathcal{P}^+|(E \cap F)}(f) \text{ if } s \in E \cap F \\ \sup_{g \in M} u(g(s)) - reg_M^{\mathcal{P}^+|(E^c \cap F)}(f) \text{ if } s \in E^c \cap F. \end{cases}$$

For all states $s \in F$, we have that $u(a_{f'}(s)) \leq \sup_{g \in M} u(g(s))$. As a result, for all states $s \in F$, we have that

$$\sup_{g \in M} u(g(s)) = \sup_{a_g \in M'} u(a_g(s)).$$

Since the regret of a plan in state s depends only on its payoff in s and the best payoff in s, it is not hard to see that the regrets of a_g with respect to M' is the same as the regret of g with respect to M. More precisely, for all $g \in M$,

$$\begin{aligned} reg_{M'}^{\mathcal{P}^+|(E \cap F)}(a_g) &= reg_M^{\mathcal{P}^+|(E \cap F)}(g), \\ reg_{M'}^{\mathcal{P}^+|(E^c \cap F)}(a_g) &= reg_M^{\mathcal{P}^+|(E^c \cap F)}(g), \text{ and} \\ reg_{M'}^{\mathcal{P}^+|F}(a_g) &= reg_M^{\mathcal{P}^+|F}(g). \end{aligned}$$

By definition of $a_{f'}$, for each state $s \in E \cap F$, we have $reg_{M'}(a_{f'}, s) = reg_M^{\mathcal{P}^+|(E \cap F)}(f)$, and for each state $s \in E^c \cap F$, we have $reg_{M'}(a_{f'}, s) = reg_M^{\mathcal{P}^+|(E^c \cap F)}(f)$. Thus, for all $\Pr \in \mathcal{P}$, if $\Pr(E \cap F) \neq 0$, then $reg_M^{\Pr|(E \cap F)}(f) = reg_M^{\mathcal{P}^+|(E \cap F)}(f)$, and if $\Pr(E^c \cap F) \neq 0$, then $reg_M^{\Pr|(E^c \cap F)}(f) = reg_M^{\mathcal{P}^+|(E^c \cap F)}(f)$.

If for all $(\Pr,\alpha) \in \mathcal{P}^+|(E\cap F)$, $\alpha\Pr(E\cap F) = 0$, then $reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_{f'}) = reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_f) = 0$. Otherwise, since there is some measure in $\mathcal{P}^+|(E\cap F)$ that has weight 1, we must have $reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_{f'}) = reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_f)$. Similarly, $reg_{M'}^{\mathcal{P}^+|(E^c\cap F)}(a_{f'}) = reg_{M'}^{\mathcal{P}^+|(E^c\cap F)}(a_f)$. Thus,

$$\begin{aligned}
reg_{M'}^{\mathcal{P}^+|F}(a_{f'}) &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F)reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right)\\
&> reg_M^{\mathcal{P}^+|F}(f) \quad \text{[by assumption]}\\
&= reg_{M'}^{\mathcal{P}^+|F}(a_f) \quad \text{[by construction]}.
\end{aligned}$$

Therefore, we have $a_{f'} \in C_{M',E\cap F}^{reg,\mathcal{P}^+}(\{a_{f'},a_f\})$, $a_{f'} \in C_{M',E^c\cap F}^{reg,\mathcal{P}^+}(\{a_{f'},a_f\})$, and $a_{f'} \notin C_{M',F}^{reg,\mathcal{P}^+}(\{a_{f'},a_f\})$, violating Axiom 4.1.

By an analogous argument, we show that the opposite weak inequality,

$$reg_M^{\mathcal{P}^+|F}(f) \leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F)reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right), \tag{C.2}$$

is also implied by Axiom 4.1. Suppose, by way of contradiction, that (C.2) does not hold. Then for some decision problem D based on (S,Σ), measurable events $E,F\subseteq S$, menu M in D, and act $f\in M$, we have that

$$reg_M^{\mathcal{P}^+|F}(f) > \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F)reg_{M,F}^{\mathcal{P}^+|(E^c\cap F)}(f)\right).$$

We define a decision problem D' based on (S,Σ) just as in the previous case. Specifically, we have that $reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_{f'}) = reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_f)$, and that $reg_{M'}^{\mathcal{P}^+|(E^c\cap F)}(a_{f'}) = reg_{M'}^{\mathcal{P}^+|(E^c\cap F)}(a_f)$. The one difference from the previous case is that we now have

$$\begin{aligned}
reg_{M'}^{\mathcal{P}^+|F}(a_{f'}) &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F)reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right)\\
&< reg_M^{\mathcal{P}^+|F}(f) \quad \text{[by assumption]}\\
&= reg_{M'}^{\mathcal{P}^+|F}(a_f) \quad \text{[by construction]}.
\end{aligned}$$

Therefore, we have $a_f \in C^{reg,\mathcal{P}^+}_{M',E\cap F}(\{a_{f'}, a_f\})$, $a_f \in C^{reg,\mathcal{P}^+}_{M',E^c\cap F}(\{a_{f'}, a_f\})$, and $a_f \notin C^{reg,\mathcal{P}^+}_{M',F}(\{a_{f'}, a_f\})$, violating Axiom 4.1.

To complete the proof that (a) implies (d), we show that Axiom 4.1 also implies that for all decision problems D based on (S, Σ), menus M in D, sets $\mathcal{P}^+$ of weighted probabilities, and acts $f \in M$, if $reg_M^{\mathcal{P}^+|(E\cap F)}(f) > 0$, then

$$reg_M^{\mathcal{P}^+|F}(f) > \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f). \tag{C.3}$$

Suppose, by way of contradiction, that (C.3) does not hold. Then for some decision problem D based on (S, Σ), events $E, F \subseteq S$, menu M in D, and act $f \in M$ such that $reg_M^{\mathcal{P}^+|(E\cap F)}(f) > 0$ and

$$reg_M^{\mathcal{P}^+|F}(f) \leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f).$$

We now define a new decision problem D' based on (S, Σ). The idea is that in D', we have a plan a_f such that $a_f \notin C^{reg,\mathcal{P}^+}_{M,E\cap F}(M')$ but $a_f \in C^{reg,\mathcal{P}^+}_{M,F}(M')$ for some $M' \subseteq M$.

Construct D' exactly as before. That is, in the first step, nature chooses a state $s \in S$, and in the second step, the DM chooses from the set of actions/plans $M' = \{a_g : g \in M\} \cup \{a_{g'}\}$. For each $g \in M$, define the actions a_g as before. We define a new action $a_{g'}$ with utilities

$$u(a_{g'}(s)) = \begin{cases} \sup_{g\in M} u(g(s)), \text{ if } s \in E \cap F \\ \sup_{g\in M} u(g(s)) - reg_M^{\mathcal{P}^+|(E^c\cap F)}(f), \text{ if } s \in E^c \cap F. \end{cases}$$

It is almost immediate from the definition of $a_{g'}$ that we have

$$reg_{M'}^{\mathcal{P}^+|F}(a_{g'}) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right) \geq reg_{M'}^{\mathcal{P}^+|F}(a_f).$$

However, we also have

$$reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_g) = 0 < reg_{M'}^{\mathcal{P}^+|(E\cap F)}(a_f).$$

Therefore, we have $a_f \notin C^{reg,\mathcal{P}^+}_{M',E\cap F}(\{a_{g'}, a_f\})$ but $a_f \in C^{reg,\mathcal{P}^+}_{M',F}(\{a_{g'}, a_f\})$, violating Axiom 4.1.

We next show that (c) implies (b). Specifically, we show that SEP for the initial menus of all decision problems D is sufficient to guarantee that Axioms 4.1–4.4 hold for menu M and all choice sets $M' \subseteq M$. It is easy to check that Axioms 4.2–4.4 hold for MWER, so we need to check only Axiom 4.1.

Consider an arbitrary decision problem D, menu M in D, $M' \subseteq M$, and a plan f in M'. We construct a new decision problem D' such that the initial menu of D' is "equivalent" to M. Just as before, let D' be a two-stage decision problem where in the first stage, nature chooses $s \in S$, and in the second stage, the DM chooses from the set $M_0 = \{a_g : g \in M\}$, where a_g is defined as before. Again, we associate each action a_g with the plan that chooses a_g in D'. M_0 is then "equivalent" to M in the sense that

$$\begin{aligned}
reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) &= reg_{M}^{\mathcal{P}^+|(E\cap F)}(g),\\
reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g) &= reg_{M}^{\mathcal{P}^+|(E^c\cap F)}(g), \text{ and}\\
reg_{M_0}^{\mathcal{P}^+|F}(a_g) &= reg_{M}^{\mathcal{P}^+|F}(g).
\end{aligned}$$

Suppose that $f \in C^{reg,\mathcal{P}^+}_{M,E\cap F}(M')$ and $f \in C^{reg,\mathcal{P}^+}_{M,E^c\cap F}(M')$. This means that for all $g \in M'$, we have $reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) \le reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g)$ and $reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f) \le$

$reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)$. Therefore, we have

$$\begin{aligned} reg_M^{\mathcal{P}^+|F}(f) &= reg_{M_0}^{\mathcal{P}^+|F}(a_f) \\ &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f)\right) \\ &\leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)\right) \\ &= reg_M^{\mathcal{P}^+|F}(g), \end{aligned}$$

which means that $f \in C_{M,F}^{reg,\mathcal{P}^+}(M')$, as required.

Next, consider an act $g \in M'$ such that $g \notin C_{M,E\cap F}^{reg,\mathcal{P}^+}(M')$. This means that $reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) < reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g)$ and $reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f) \leq reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)$. Let $(\alpha_{\Pr^*}, \Pr^*) \in C(\mathcal{P}^+)$ be such that

$$\begin{aligned} & \alpha_{\Pr^*}(\Pr^*(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) + \Pr^*(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g) \\ = \;& \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)\right). \end{aligned}$$

Such a pair $(\alpha_{\Pr^*}, \Pr^*)$ exists, since we have assumed that $C(\mathcal{P}^+)$ is closed. If $\alpha_{\Pr^*}\Pr^*(E\cap F) = 0$, then $reg_{M_0}^{\mathcal{P}^+|F}(a_g) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)\right)$. By separability, it must be the case that $reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) = 0$, contradicting our assumption that $0 \leq reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) < reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g)$. Therefore, it must be that $\alpha_{\Pr^*}\Pr^*(E\cap F) > 0$, and

$$\begin{aligned} reg_M^{\mathcal{P}^+|F}(f) &= reg_{M_0}^{\mathcal{P}^+|F}(a_f) \\ &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f)\right) \\ &< \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)\right) \\ &= reg_M^{\mathcal{P}^+|F}(g), \end{aligned}$$

which means that $g \notin C_{M,F}^{reg,\mathcal{P}^+}(M')$. □

C.3 Proof of Theorem 4.5.5

To prove Theorem 4.5.5, we need the following lemma.

Lemma C.3.1. *For all utility functions u, sets $\mathcal{P}^+$ of weighted probabilities, acts f, and menus M containing f, $reg_M^{\mathcal{P}^+}(f) = reg_M^{C(\mathcal{P}^+)}(f)$.*

Proof. Simply observe that

$$\begin{aligned} reg_M^{\mathcal{P}^+}(f) &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+}\left(\alpha\sum_{s\in S}\Pr(s)reg_M(f,s)\right) \\ &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+}\left(\sum_{s\in S}\alpha\Pr(s)reg_M(f,s)\right) \\ &= \sup_{\{p:\, p\le\alpha\Pr,(\Pr,\alpha)\in\mathcal{P}^+\}}\left(\sum_{s\in S}p(s)reg_M(f,s)\right) \\ &= reg_M^{C(\mathcal{P}^+)}(f), \end{aligned}$$

by definition. $\square$

The next lemma uses an argument almost identical to one used in Lemma 7 of [Halpern and Leung 2012].

Lemma C.3.2. *If $C(\mathcal{P}^+|^{\chi}F)$ is convex and q is a subprobability on F not in $\overline{C(\mathcal{P}^+|^{\chi}F)}$, then there exists a non-negative vector θ such that for all $(\Pr,\alpha)\in\mathcal{P}^+|^{\chi}F$, we have*

$$\sum_{s\in F}\alpha\Pr(s)\theta(s) < \sum_{s\in F}q(s)\theta(s).$$

Proof. Given a set $\mathcal{P}^+$ of weighted probabilities, let $C'(\mathcal{P}^+) = \{p : p \in \mathbb{R}^{|S|}$ and $p \le \alpha\Pr$ for some $(\Pr,\alpha)\in\mathcal{P}^+\}$. Note that an element $q \in C'(\mathcal{P}^+)$ may not be a subprobability measure, since we do not require that $q(s)\ge 0$. Since

$\overline{C'(\mathcal{P}^+|^\chi F)}$ and $\{q\}$ are closed, convex, and disjoint, and $\{q\}$ is compact, the separating hyperplane theorem [Rockafellar 1970] says that there exist $\theta \in \mathbb{R}^{|S|}$ and $c \in \mathbb{R}$ such that

$$\theta \cdot p < c \text{ for all } p \in \overline{C'(\mathcal{P}^+|^\chi F)}, \text{ and } \theta \cdot q > c. \tag{C.4}$$

Since $\{\alpha \Pr : (\Pr, \alpha) \in \mathcal{P}^+|^\chi F\} \subseteq \overline{C'(\mathcal{P}^+|^\chi F)}$, we have that for all $(\Pr, \alpha) \in \mathcal{P}^+|^\chi F$,

$$\sum_{s \in F} \alpha \Pr(s)\theta(s) < \sum_{s \in F} q(s)\theta(s).$$

Now we argue that it must be the case that $\theta(s) \geq 0$ for all $s \in F$. Suppose that $\theta(s') < 0$ for some $s' \in F$. Define p^* by setting

$$p^*(s) = \begin{cases} 0, \text{ if } s \neq s' \\ \frac{-|c|}{|\theta(s')|}, \text{ if } s = s'. \end{cases}$$

Note that $p^* \leq \vec{0}$, since for all $s \in S$, $p^*(s) \leq 0$. Therefore, $p^* \in C'(\mathcal{P}^+|^\chi F)$.

Our definition of p^* also ensures that $\theta \cdot p^* = \sum_{s \in S} p^*(s)\theta(s) = p^*(s')\theta(s') = |c| \geq c$. This contradicts (C.4), which says that $\theta \cdot p < c$ for all $p \in C'(\mathcal{P}^+|^\chi F)$. Thus it must be the case that $\theta(s) \geq 0$ for all $s \in S$. □

We are now ready to prove Theorem 4.5.5, which we restate here.

THEOREM 4.5.5. *If $C(\mathcal{P}^+)$ is closed and convex, then Axiom 4.1 holds for the family of choices $C_M^{reg,\mathcal{P}^+|^\chi E}$ if and only if $\mathcal{P}^+$ is χ-rectangular.*

We prove the two directions of implication in the theorem separately. Note that the proof that χ-rectangularity implies Axiom 4.1 does not require $C(\mathcal{P}^+)$ to be convex.

Claim C.3.3. *If $\mathcal{P}^+$ is χ-rectangular, then Axiom 4.1 holds for the family of choices $C_M^{reg,\mathcal{P}^+|^\chi E}$.*

Proof. By Theorem 4.5.3, it suffices to show that SEP holds. For the first part of SEP, we must show that

$$reg_M^{\mathcal{P}^+|^\chi F}(f) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \left(\Pr(E\cap F) reg_M^{\mathcal{P}^+|^\chi (E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\chi (E^c\cap F)}(f) \right). \tag{C.5}$$

Unwinding the definitions, (C.5) is equivalent to

$$\begin{array}{rl} & reg_M^{\mathcal{P}^+|^\chi F}(f) \\ = & \sup_{(\Pr_3,\alpha_3)\in\mathcal{P}^+|^\chi F} \alpha_{\Pr_3} \left(\Pr_3(E\cap F) \sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\chi F} \alpha^\chi_{1,E\cap F} \sum_{s\in E\cap F} \Pr_1(s|(E\cap F)) reg_M(f,s)) \right. \\ & \left. + \Pr_3(E^c\cap F) \sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\chi F} \alpha^\chi_{2,E^c\cap F} \sum_{s\in E^c\cap F} \Pr_2(s|(E^c\cap F)) reg_M(f,s)) \right). \end{array}$$

The sups in this expression are taken on by some $(\Pr_1^*,\alpha_1^*),(\Pr_2^*,\alpha_2^*),(\Pr_3^*,\alpha_3^*) \in \overline{\mathcal{P}^+|^\chi F}$. By χ-rectangularity, we have that for all $(\Pr_1,\alpha_1),(\Pr_2,\alpha_2),(\Pr_3,\alpha_3) \in \mathcal{P}^+|^\chi F$,

$$\alpha_{\Pr_3}\Pr_3(E\cap F)\alpha^\chi_{1,E\cap F}\Pr_1|(E\cap F) + \alpha_{\Pr_3}\Pr_3(E^c\cap F)\alpha^\chi_{2,E^c\cap F}\Pr_2|(E^c\cap F) \in \overline{C(\mathcal{P}^+|^\chi F)}. \tag{C.6}$$

Thus, for all $\epsilon > 0$,

$$\begin{array}{rl} & reg_M^{\mathcal{P}^+|^\chi F}(f) \\ = & reg_M^{C(\mathcal{P}^+|^\chi F)}(f) \quad \text{[by Lemma C.3.1]} \\ \geq & \alpha_3^* \left(\Pr_3^*(E\cap F)(\alpha^*_{1,E\cap F})^\chi \sum_{s\in E\cap F} \Pr_1^*(s|(E\cap F)) reg_M(f,s)) \right. \\ & \left. + \Pr_3^*(E^c\cap F)(\alpha^*_{2,E^c\cap F})^\chi \sum_{s\in E^c\cap F} \Pr_2^*(s|(E^c\cap F)) reg_M(f,s)) \right) - \epsilon \text{ [by (C.6)]}. \end{array}$$

Therefore,

$$
\begin{array}{ll}
& reg_M^{\mathcal{P}^+|^\chi F}(f) \\
\geq & \alpha_3^* \left(\Pr_3^*(E\cap F)(\alpha_{1,E\cap F}^*)^\chi \sum_{s\in E\cap F} \Pr_1^*(s|(E\cap F))reg_M(f,s)\right) \\
& \quad + \Pr_3^*(E^c\cap F)(\alpha_{2,E^c\cap F}^*)^\chi \sum_{s\in E^c\cap F} \Pr_2^*(s|(E^c\cap F))reg_M(f,s))\Big) \\
= & \sup_{(\Pr_3,\alpha_3)\in\mathcal{P}^+|^\chi F} \alpha_3 \left(\Pr_3(E\cap F)\sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\chi F} \alpha_{1,E\cap F}^\chi \sum_{s\in E\cap F} \Pr_1(s|(E\cap F))reg_M(f,s)\right) \\
& \quad + \Pr_3(E^c\cap F)\sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\chi F} \alpha_{2,E^c\cap F}^\chi \sum_{s\in E^c\cap F} \Pr_2(s|(E^c\cap F))reg_M(f,s))\Big) \\
& \quad \text{[by the choice of } (\Pr_i^*,\alpha_i^*),\ i=1,2,3\text{]} \\
= & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c\cap F)reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\right),
\end{array}
$$

as required.

It remains to show the opposite inequality in (C.5), namely, that

$$
reg_M^{\mathcal{P}^+|^\chi F}(f) \leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \left(\Pr(E\cap F)reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c\cap F)reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\right).
$$

It suffices to note that the right-hand side is equal to

$$
\begin{array}{ll}
& \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \left(\alpha\Pr(E\cap F)\sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\chi F} \alpha_{1,E\cap F}^\chi \sum_{s\in E\cap F} \Pr_1(s|E\cap F)reg_M(f,s)\right) \\
& +\alpha\Pr(E^c\cap F)\sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\chi F} \alpha_{2,E^c\cap F}^\chi \sum_{s\in E^c\cap F} \Pr_2(s|E^c\cap F)reg_M(f,s))\Big) \\
\geq & \overline{E}_{\mathcal{P}^+|^\chi F}(reg_M(f)) \quad \text{[by rectangularity]} \\
= & reg_M^{\mathcal{P}^+|^\chi F}(f).
\end{array}
$$

This completes the proof that (C.5) holds.

For the second part of SEP, suppose that $\overline{\mathcal{P}}^+(E\cap F) > 0$ and $reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) \neq 0$. If $reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f) = 0$ then, since $\overline{\mathcal{P}}^+(E\cap F) > 0$, we have that $reg_M^{\mathcal{P}^+|^\chi F}(f) > 0 = \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha Pr(E^c\cap F)reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)$, as desired. Otherwise, by part (b) of χ-rectangularity, for all $\delta > 0$, there exists $(\Pr,\alpha) \in \mathcal{P}^+|^\chi F$ such that $\alpha(\delta\Pr(E\cap F) + \Pr(E^c\cap F)) > \sup_{(\Pr',\alpha')\in\mathcal{P}^+} \alpha'\Pr'(E^c\cap F)$. Therefore, using the first part of SEP, we have

$$
\begin{array}{rl}
& reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) \\
= & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \left(\Pr(E\cap F) reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\right) \\
= & reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f) \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \left(\Pr(E\cap F)\frac{reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f)}{reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)} + \Pr(E^c\cap F)\right) \\
> & reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f) \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha Pr(E^c\cap F) \quad \text{[by part (b) of } \chi\text{-rectangularity]} \\
= & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha Pr(E^c\cap F) reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f),
\end{array}
$$

as required. □

Claim C.3.4. *If $C(\mathcal{P}^+)$ is convex and Axiom 4.1 holds for the family of choices $C_M^{reg,\mathcal{P}^+|^\chi E}$, then $\mathcal{P}^+$ is χ-rectangular.*

Proof. Suppose that χ-rectangularity does not hold. Then one of the three conditions of rectangularity must fail.

First suppose that it is (a); that is, for some $(\Pr_1,\alpha_1),(\Pr_2,\alpha_2),(\Pr_3,\alpha_3)\in\mathcal{P}^+$, we have $\Pr_1(E\cap F)>0$ and $\Pr_2(E^c\cap F)>0$ and

$$\alpha_3\Pr_3(E\cap F)\alpha^\chi_{1,E\cap F}\Pr_1|(E\cap F)+\alpha_3\Pr_3(E^c\cap F)\alpha^\chi_{2,E^c\cap F}\Pr_2|(E^c\cap F)\notin\overline{C(\mathcal{P}^+)|^\chi F}.$$

Let $p^* = \alpha_3\Pr_3(E\cap F)\alpha^\chi_{1,E\cap F}\Pr_1|(E\cap F)+\alpha_3\Pr_3(E^c\cap F)\alpha^\chi_{2,E^c\cap F}\Pr_2|(E^c\cap F)$. Since we have assumed that $C(\mathcal{P}^+)$ is convex, we have that $C(\mathcal{P}^+|^\chi F)$ is also convex. By Lemma C.3.2, there exists a non-negative vector θ such that for all $\alpha\Pr\in\overline{C(\mathcal{P}^+|^\chi F)}$, we have

$$\sum_{s\in F}\alpha\Pr(s)\theta(s)<\sum_{s\in F}p^*(s)\theta(s).$$

We construct a decision problem D based on (S,Σ). D has two stages: in the first stage, nature chooses a state $s\in S$, but only states in $F\subseteq S$ are chosen with positive probability, so when the DM plays, his beliefs are characterized by

$\mathcal{P}^+|^\chi F$. In the second stage, the DM chooses an action from the set $M = \{f, g\}$, with utilities defined as follows:

$$u(f,s) = -\theta(s), \text{ and}$$
$$u(g,s) = 0 \text{ for all s.}$$

The act f will have regret precisely $\theta(s)$ in state $s \in S$. By Lemma C.3.2,

$$\begin{array}{rl} & \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E\cap F) reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\right) \\ \geq & \alpha_{\Pr_3} \left(\Pr_3(E\cap F) reg_M^{\alpha^\chi_{1,E\cap F} \Pr_1|(E\cap F)}(f) + \Pr(E^c\cap F) reg_M^{\alpha^\chi_{2,E^c\cap F} \Pr_2|(E^c\cap F)}(f)\right) \\ = & \sum_{s\in F} p^*(s)\theta(s) \\ > & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} reg_M^{\mathcal{P}^+|^\chi F}(f), \end{array}$$

violating SEP. By Theorem 4.5.3, Axiom 4.1 cannot hold.

Now suppose that condition (b) in rectangularity does not hold. That is, for some $\delta > 0$, for all $(\alpha, \Pr) \in \mathcal{P}^+$, $\alpha(\delta \Pr(E\cap F) + \Pr(E^c \cap F)) \leq \sup_{(\Pr',\alpha')\in\mathcal{P}^+} \alpha' \Pr'(E^c \cap F)$. We construct a decision problem D based on (S, Σ). D has two stages: in the first stage, nature chooses a state $s \in S$. In the second stage, the DM chooses an action from the set $M = \{f, g\}$, with utilities defined as follows:

$$\begin{array}{ll} u(f,s) = 0 & \text{for all } s \in S, \\ u(g,s) = -\delta & \text{if } s \in E\cap F \\ u(g,s) = -1 & \text{if } s \notin E\cap F. \end{array}$$

Then we have that $reg_M^{\mathcal{P}^+|^\chi(E\cap F)}(g) = \delta$ and $reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(g) = 1$. Using SEP and the choice of δ, we must have

$$\begin{array}{rcl} reg_M^{\mathcal{P}^+|^\chi F}(g) & = & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha(\Pr(E\cap F)\delta + \Pr(E^c\cap F)) \\ & \leq & \sup_{\Pr\in\mathcal{P}^+|^\chi F} \alpha \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(g). \end{array}$$

Clearly,

$$reg_M^{\mathcal{P}^+|^\chi F}(g) \geq \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F} \alpha \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(g).$$

Thus,

$$reg_M^{\mathcal{P}^+|^\times F}(g) = \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\times F} \alpha \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\times (E^c\cap F)}(g),$$

violating the second condition of SEP. Therefore, by Theorem 4.5.3, Axiom 4.1 does not hold.

Finally, suppose that condition (c) in rectangularity does not hold. Then for some nonnegative real vector $\theta \in \mathbb{R}^{|S|}$,

$$\begin{array}{ll} & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\times F} \big(\alpha \Pr(E) \sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\times (E\cap F)} \sum_{s\in E\cap F} \alpha_1 \Pr_1(s|E)\theta(s)) \\ & +\alpha \Pr(E^c) \sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\times (E^c\cap F)} \sum_{s\in E^c\cap F} \alpha_2 \Pr_2(s|E^c)\theta(s))\big) \\ < & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\times F} \alpha \sum_{s\in F} \Pr(s)\theta(s). \end{array} \tag{C.7}$$

We construct a decision problem D based on (S,Σ). D has two stages: in the first stage, nature chooses a state $s \in S$. In the second stage, the DM chooses an action from the set $M = \{f, g\}$, with utilities defined as follows:

$$u(g,s) = -\theta(s) \quad \text{for all } s \in S.$$
$$u(f,s) = 0 \quad \text{for all } s \in S.$$

So we have

$$\begin{array}{ll} & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^p F} \alpha \Big(\Pr(E\cap F) reg_M^{\mathcal{P}^+|^p (E\cap F)}(g) + \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|^p (E^c\cap F)}(g)\Big) \\ = & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\times F} \alpha \big(\Pr(E\cap F)\overline{E}_{\mathcal{P}^+|^\times (E\cap F)}(\theta) + \Pr(E^c\cap F)\overline{E}_{\mathcal{P}^+|^\times (E^c\cap F)}(\theta)\big) \\ < & \overline{E}_{\mathcal{P}^+|^\times F}(\theta) \quad \text{[by (C.7)]} \\ = & reg_M^{\mathcal{P}^+|^p F}(g). \end{array}$$

This means that SEP, and hence Axiom 1, is violated, a contradiction. □

BIBLIOGRAPHY

Abdellaoui, M., A. Baillon, L. Placido, and P. P. Wakker (2011, April). The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *American Economic Review 101*(2), 695–723.

Al-Najjar, N. and J. Weinstein (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy 25*, 249–284.

Allais, M. (1953a). Le comportement de l'homme rationel devant le risque: critique de l'école Americaine. *Econometrica 21*, 503–546.

Allais, M. (1953b). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine. *Econometrica 21*(4), 503–546.

Anscombe, F. and R. Aumann (1963). A definition of subjective probability. *Annals of Mathematical Statistics 34*, 199–205.

Ariely, D. (2008, February). *Predictably Irrational: The Hidden Forces That Shape Our Decisions* (1 ed.). HarperCollins.

Bateman, I. J., A. Munro, and G. L. Poe (2008). Decoy effects in choice experiments and contingent valuation: Asymmetric dominance. *Land Economics 84*(1).

Campos, L. M. d. and S. Moral (1995). Independence concepts for sets of probabilities. pp. 108–115.

Cattaneo, M. E. G. V. (2007). *Statistical decisions based directly on the likeihood function*. Ph. D. thesis, ETH.

Chao, L., K. A. Hanley, C. L. Burch, C. Dahlberg, and P. E. Turner (2000). Kin selection and parasite evolution: Higher and lower virulence with hard and soft selection. *The Quarterly Review of Biology 75*(3), pp. 261–275.

Chateauneuf, A. and J. Faro (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics 45*, 535 – 558.

Chew, S. H. (1983, July). A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the allais paradox. *Econometrica 51*(4), 1065–92.

Cooman, G. de (2005). A behavioral model for vague probability assessments. *Fuzzy Sets and Systems 154*(3), 305–358.

Cousa, I., S. Moral, and P. Walley (1999). Examples of independence for imprecise probabilities. In *Proc. First International Symposium on Imprecise Probabilities and Their Applications (ISIPTA '99)*.

Crawley, M. (1993). Some sociological misconceptions, center for population biology. *Pan-European Conference on the Potential Long-Term Ecological Impact of Genetically Modified Organisms: Proceedings November 1993*, 276.

Doyle, J. R., D. J. O'Connor, G. M. Reynolds, and P. A. Bottomley (1999). The Robustness of the Asymmetrically Dominated Effect: Buying Frames, Phantom Alternatives, and In-Store Purchases. *Psychology and Marketing 16*(3).

Dubois, D., H. Fargier, and P. Perny (2003). Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence 148*(12), 219 – 260.

Dubois, D. and H. Prade (1998). Possibility measures: qualitative and quantitative aspects. In D. M. Gabbay and P. Smets (Eds.), *Quantified Representation of Uncertainty and Imprecision*, Volume 1 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, pp. 169–226. Dordrecht, Netherlands: Kluwer.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics 75*(4), 643–669.

Epstein, L. G. and M. Le Breton (1993). Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory 61*(1), 1–22.

Epstein, L. G. and M. Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory 113*(1), 1–31.

Epstein, L. G. and M. Schneider (2007). Learning under ambiguity. *Review of Economic Studies 74*(4), 1275–1303.

Farquhar, P. H. and A. R. Pratkanis (1993). Decision structuring with phantom alternatives. *Management Science 39*(10), pp. 1214–1226.

Gärdenfors, P. and N. Sahlin (1982). Unreliable probabilities, risk taking, and decision making. *Synthese 53*, 361–386.

Gärdenfors, P. and N. Sahlin (1983). Decision making with unreliable probabilities. *British Journal of Mathematical and Statistical Psychology 36*, 240–251.

Ghirardato, P. (2002). Revisiting savage in a conditional world. *Economic Theory 20*(1), 83–92.

Gilboa, I. and D. Schmeidler (1989a). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics 18*, 141–153.

Gilboa, I. and D. Schmeidler (1989b). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics 18*(2), 141–153.

Gilboa, I. and D. Schmeidler (1993). Updating ambiguous beliefs. *Journal of Economic Theory 59*, 33–49.

Good, I. J. (1980). Some history of the hierarchical Bayesian methodology. In J. M. Bernardo, M. H. DeGroot, D. Lindley, and A. Smith (Eds.), *Bayesian Statistic I*, pp. 489–504. University Press: Valencia.

Gul, F. and W. Pesendorfer (2005). The revealed preference theory of changing tastes. *Review of Economic Studies 72*(2), 429–448.

Halpern, J. Y. and S. Leung (2012). Weighted sets of probabilities and minimax weighted expected regret: new approaches for representing uncertainty and making decisions. In *Proc. Twenty-Ninth Conference on Uncertainty in Artificial Intelligence (UAI 2012)*, pp. 336–345.

Hammond, P. J. (1976). Changing tastes and coherent dynamic choice. *The Review of Economic Studies 43*(1), 159–173.

Hayashi, T. (2008a). Regret aversion and opportunity dependence. *Journal of Economic Theory 139*(1), 242–268.

Hayashi, T. (2008b). Regret aversion and opportunity dependence. *Journal of Economic Theory 139*(1), 242–268.

Hayashi, T. (2009). Stopping with anticipated regret. Unpublished manuscript.

Hayashi, T. (2011). Context dependence and consistency in dynamic choice under uncertainty: the case of anticipated regret. *Theory and Decision 70*, 399–430.

Hedgcock, W., A. R. Rao, and H. A. Chen (1999). Could Ralph Nader's Entrance and Exit Have Helped Al Gore? The Impact of Decoy Dynamics on Consumer Choice. *Journal of Marketing Research 46*(3).

Herstein, I. N. and J. Milnor (1953). An axiomatic approach to measurable utility. *Econometrica 21*(2), 291–297.

Huber, J., J. W. Payne, and C. Puto (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *JCR 9*(1).

Jaffray, J.-Y. (1992, Sep). Bayesian updating and belief functions. *Systems, Man and Cybernetics, IEEE Transactions on 22*(5), 1144–1152.

Jaffray, J.-Y. (1994). Dynamic decision making with belief functions. In R. R. Yager, J. Kacprczyk, and M. Fedrizzi (Eds.), *Advances in the Dempster-Shafer Theory of Evidence*, pp. 331–352. New York: Wiley.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Houghton Mifflin.

Kreps, D. M. (1988). *Notes on the Theory of Choice*. Boulder, Colo.: Westview Press.

Larbi, R. B., S. Konieczny, and P. Marquis (2010). A characterization of optimality criteria for decision making under complete ignorance. In *KR*.

Larrick, R. P. and T. L. Boles (1995, July). Avoiding regret in decisions with feedback: A negotiation example. *OBHDP 63*(1), 87–97.

Levi, I. (1985). Imprecision and uncertainty in probability judgment. *Philosophy of Science 52*, 390–406.

Luce, R. D. and H. Raiffa (1957). *Games and decisions : introduction and critical survey*. Wiley.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature 27*(4), 1622–68.

Marichal, J.-L. (2001). An axiomatic approach of the discrete Sugeno integral as a tool to aggregate interacting criteria in a qualitative framework. *IEEE Transactions on Fuzzy Systems 9*(1), 164–172.

Moral, S. (1992). Calculating uncertainty intervals from conditional convex sets of probabilities. In *Proc. Eighth Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 199–206.

Nau, R. F. (1992). Indeterminate probabilities on finite sets. *Annals of Statistics 40*(4), 1737–1767.

Niehans, J. (1948). Zur preisbildung bei ungewissen erwartungen. *Schweizerische Zeitschrift für Volkswirtschaft und Statistik 84*(5), 433–456.

Pires, C. P. (2002). A rule for updating ambiguous beliefs. *Theory and Decision 53*(2), 137–152.

Ritov, I. (1996, May). Probability of regret: Anticipation of uncertainty resolution in choice. *Organizational Behavior and Human Decision Processes 66*(2), 228–236.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton, N. J.: Princeton University Press.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica 76*(2), 263–305.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association 46*, 55–67.

Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.

Shafir, S., T. A. Waite, and B. H. Smith (2002). Context-Dependent violations of rational choice in honeybees (apis mellifera) and gray jays (perisoreus canadensis). *Behavioral Ecology and Sociobiology 51*(2), 180–187.

Simonson, I. and A. Tversky (1992). Choice in context: Tradeoff contrast and extremeness aversion. *Journal of Marketing Research 29*(3), 281–295.

Siniscalchi, M. (2006, December). Dynamic Choice Under Ambiguity. Discussion papers, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics 6*(3), 379–421.

Stoye, J. (2011a). Axioms for minimax regret choice correspondences. *Journal of Economic Theory 146*(6), 2226 – 2251.

Stoye, J. (2011b). Statistical decisions under ambiguity. *Theory and Decision 70*(2), 129–148.

Stoye, J. (2013). Choice theory when agents can randomize.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies 23*(3), 165–180.

Wald, A. (1950). *Statistical Decision Functions*. John Wiley, NY.

Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*, Volume 42 of *Monographs on Statistics and Applied Probability*. London: Chapman and Hall.

Walley, P. (1997). Statistical inferences based on a second-ordr possibility distribution. *International Journal of General Systems 26*(4), 337–383.

Wedell, D. H. (1991). Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology 17*(4), 767–778+.

Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems 1*, 3–28.

Zimmerman, A. S. (2010). Funding Irrationality. *Duke Law Journal 59*(1105).