Minimizing Regret in Dynamic Decision Problems*

Joseph Y. Halpern Samantha Leung
Department of Computer Science
Cornell University
Ithaca, NY 14853
halpern | samlyy@cs.cornell.edu

February 2, 2016

Abstract

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.

1 Introduction

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 [1953] and Ellsberg [1961] point out using compelling examples, sometimes intuitive choices are incompatible with maximizing expected utility. 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

*Work supported in part by NSF grants IIS-0812045, IIS-0911036, and CCF-1214844, by AFOSR grants FA9550-08-1-0438, FA9550-09-1-0266, and FA9550-12-1-0040, and by ARO grant W911NF-09-1-0281. We thank an anonymous reviewer and Spyros Galanis for useful comments.

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 2011]. Other works that use a set of probablity measures include, for example, [Campos and Moral 1995; Cousa, Moral, and Walley 1999; Gilboa and Schmeidler 1993; Levi 1985; Walley 1991].

In this paper, we consider a generalization of minimax expected regret called minimax *weighted* expected regret (MWER) that we introduced in an earlier paper [Halpern and Leung 2012]. 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 decision maker (henceforth DM). The weights work much like a "second-order" probability on the set of probability measures. However, MWER does not require the specification of a precise second-order probability distribution. Having equal weights on a set of probability measures does not mean that they are equally likely; instead it means that we do not know the relative likelihoods between these probability measures. Using sets of weighted probability measures as a representation of uncertainty has the advantage that it can represent ambiguity, while at the same time dealing well with updating. For example, consider a DM who is uncertain about the bias of a coin, thinking initially it could be anywhere between .1 and .9. This uncertainty can be represented by the obvious set of probability measures, namely, the set of all probability measures Pr on {head,tail} where $\Pr(head) \in [0.1, 0.9]$. If the coin is then tossed 100 times, and 80 heads are observed, the DM should now be much more certain that the true bias of the coin is close to .8. There is no natural way to represent this learning using unweighted probability measures; the *likelihood updating* rule introduced in our earlier paper (and discussed in Section 2) provides a natural way of doing so.

In [Halpern and Leung 2012], we provided an axiomatization of MWER and proved a representation theorem (as well as providing much more motivation for the approach). Our motivation for MWER is largely pragmatic, although it has some normative and descriptive features. It is clear from many examples that people do not always act so as to maximize expected utility; moreover, research in psychology shows that regret is a powerful motivator. While we do believe that regret is an important component of an accurate description of human decision making, so MWER does capture some features of human decision making, we are not claiming that MWER is a compleley accurate descripton of how people make decisions. Rather, we view MWER as a pragmatically useful approach for making decisions, especially in cases where agents do not feel that they can describe their uncertainty adequately by a single prior. We outline its advantages in the paper. However, we are not arguing that an agent would be irrational if he did not make decisions using MWER, so we do not view it as normative in this strong sense.

While MWER is a natural decision rule when considering weighted probability measures, it is, of course, not the only one that can be considered. The notion of using weighted 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. [2006a]; Nau [1992]) proposed approaches to decision making using their representations of uncertainty.

The focus of our earlier paper was on static decision making. 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? This question is important for situations where the DM has to handle ambiguity and learning, for example doctors choosing treatments for patients.

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 3 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. Other conditions that guarantee dynamic consistency for ambiguity-averse decision rules have also been proposed. For example, Hanany and Klibanoff [**?**, **?**] guarantee dynamic consistency for ambiguity-averse preferences by using updating rules that depend on the ex ante preferred plans and/or the set of feasible choices. Weinstein [2009] provides an overview of the approaches used in achieving dynamic consistency with ambiguity-averse decision rules.

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, 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 paper is organized as follows. Section 2 discuss preliminaries. Section 3 introduces forgone opportunities. Section 4 gives conditions under which consistent planning is not required. We conclude in Section 5. We defer most proofs to the appendix.

2 Preliminaries

2.1 Static decision setting and regret

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 paper, 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, regret can depend on the menu. 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. 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

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

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'.

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* $(S, \Sigma, X, u, \Pr)$. 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 $reg_M^{\mathcal{P}}(f)$.

As we said, in an earlier paper, we considered another representation of uncertainty, *weighted set of probability measures* [Halpern and Leung 2012]. 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 $\alpha_{\Pr} = 1$.

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

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

If $\mathcal{P}^+$ is empty, then $reg_M^{\mathcal{P}^+}$ is identically zero. Of course, we can define the choice functions $C_M^{reg,\Pr}$, $C_M^{reg,\mathcal{P}}$, and $C_M^{reg,\mathcal{P}^+}$ using $reg_M^{\Pr}$, $reg_M^{\mathcal{P}}$, and $reg_M^{\mathcal{P}^+}$, by analogy with C_M^{reg}.

2.2 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 by 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. Moreover, the DM's preference relation is assumed to depend only on her knowledge, and not on her past choices. These assumptions significantly restrict 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:

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

Put simply, prior-by-prior updating conditions each probability measure in the set on the information, and does not change any of the weights in a weighted set of probability measures. Prior-by-prior updating can produce some rather counterintuitive 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 another way of updating weighted sets of probabilities, called *likelihood updating* [Halpern and Leung 2012]. The intuition is that the weights are updated as if they were a second-order probability distribution over the probability measures. Given an event $E \subseteq S$, define $\overline{\mathcal{P}}^+(E) = \sup\{\alpha \Pr(E) : (\Pr, \alpha) \in \mathcal{P}^+\}$; if $\overline{\mathcal{P}}^+(E) > 0$, let $\alpha^l_E = \sup_{\{(\Pr', \alpha') \in \mathcal{P}^+ : \Pr'|E = \Pr|E\}} \frac{\alpha' \Pr'(E)}{\overline{\mathcal{P}}^+(E)}$. Given a measure $\Pr \in \mathcal{P}$, there may be several distinct measures $\Pr'$ in $\mathcal{P}$ such that $\Pr'|E = \Pr|E$. Thus, we take the weight of $\Pr|E$ to be the sup of the possible candidate values of α^l_E. By dividing by $\overline{\mathcal{P}}^+(E)$, we guarantee that $\alpha^l_E \in [0, 1]$, and that there is some weighted measure $(\Pr, \alpha)$ such that $\alpha^l_E = 1$, as long as there is some pair $(\Pr, \alpha) \in \mathcal{P}^+$ such that $\alpha \Pr(E) = \overline{\mathcal{P}}^+(E)$. If $\overline{\mathcal{P}}^+(E) > 0$, we take $\mathcal{P}^+|^l E$, the result of applying likelihood updating by E to $\mathcal{P}^+$, to be

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

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^l_E. 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 $C_M^{reg,\Pr|E}$.

3 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 3.1. 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 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.

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

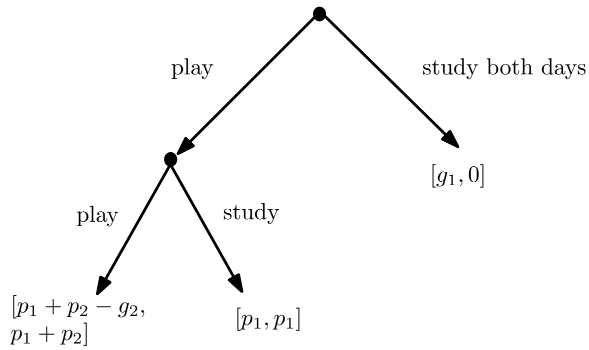

Figure 1: An explanation for procrastination.

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 3.1 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 3.2 (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 3.1 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 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 3.3. *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.*

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

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 3.3 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 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 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 2. Note that at the node 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*

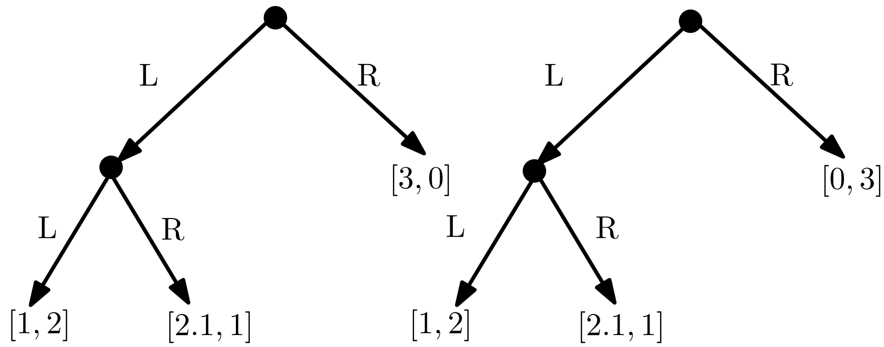

Figure 2: Two decision trees.

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 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 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 $g \in C_{M,E}(M')$.*

Axiom 1 says that if an act f would be chosen if the DM learns that the true state is in the set E, and that f would also be chosen if the DM learns that the true state is in the set E^c, then f would be chosen if the DM didn't learn either of these. Furthermore, if there is such a "strong option" f, then no other act g can be chosen unless g is at least as good as f if the true state is in the set E. Axiom 2 says that two acts that represent the same distributions over outcomes on all states must be treated by the choice function equally.

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

Axiom 3. *$C_{M,E}(M') \subseteq M'$ and $C_{M,E}(M') \neq \emptyset$ if $M' \neq \emptyset$.*

Axiom 4 (Sen's α). *If $f \in C_{M,E}(M'), f \in M''$, and $M'' \subseteq M'$, then $f \in C_{M,E}(M'')$.*

Axiom 3 says that the choice from a nonempty menu must be a nonempty subset of the menu. Axiom 4 says that if an act is chosen from a certain menu, then it must also be chosen from a subset of that menu.

Theorem 4.1. *For a dynamic decision problem D, if Axiom 1–4 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 1–4 hold for a MWER decision maker. 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 4.2 (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 1–4 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 4.3. *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 1–4 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 4.3 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 4.4 (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

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

(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]. Hayashi proved a similar result for MER [Hayashi 2009], while Maccheroni, Marinacci, and Rustichini [2006b] prove an analogous result for variational preferences. With MMEU and MER, only unweighted probabilities are considered. Definition 4.4 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 1–4 under prior-by-prior updating and likelihood updating.

Theorem 4.5. *If $C(\mathcal{P}^+)$ is closed and convex, then Axiom 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 1 requires only that $C(\mathcal{P}^+)$ be closed (i.e., convexity is not required). Hayashi [2011] proves an analogue of Theorem 4.5 for MER using prior-by-prior updating. That is, χ-rectangularity can also be thought of as a generalization of Hayashi's conditions. In addition to prior-by-prior updating, Hayashi 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 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 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 decision maker to ignore unachievable plans. That is, a decision maker may choose to ignore unachievable plans because doing so simplifies the search for the preferred solution.

5 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 interesting to investigate the effect of including such in the menus of a sophisticated DM. A sophisticated decision maker 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 decision maker 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.

[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].

A Proof of Theorem 4.1

We restate the theorem (and elsewhere in the appendix) for the reader's convenience.

Theorem 4.1. *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 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 1, $f'I(h')g \in C_{\mu,h}(M_{h'})$. Also by Axiom 1, since $f \notin C_{\mu,h'}(M_{h'})$, we must have $f \notin C_{\mu,h}(M_{h'})$. By Axiom 4, this implies that $f \notin C_{\mu,h}(M_h)$ (since $M_{h'} \subseteq M_h$), giving us the desired contradiction. □

B Proof of Theorem 4.3

THEOREM 4.3. *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 1–4 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 B.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 1–4 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 1–4 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 1 implies that for all decision problems D based on (S, Σ), menu M in D, sets $\mathcal{P}^+$ of weighted probabilities, and acts $f \in M$,

$$reg_M^{\mathcal{P}^+|F}(f) \geq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f) \right). \quad (1)$$

Suppose, by way of contradiction, that (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 $a_{f'} \in C^{reg,\mathcal{P}^+}_{M',E\cap F}(M'')$ and $a_{f'} \in C^{reg,\mathcal{P}^+}_{M',E^c\cap F}(M'')$ and $a_{f'} \notin C^{reg,\mathcal{P}^+}_{M',F}(M'')$ 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,

$$\begin{aligned} reg_{M'}^{\mathcal{P}^+|F}(a_{f'}) &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f) \right) \\ &> reg_M^{\mathcal{P}^+|F}(f) \qquad \text{[by assumption]} \\ &= reg_{M'}^{\mathcal{P}^+|F}(a_f) \qquad \text{[by construction]}. \end{aligned}$$

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 1.

By an analogous argument, we show that the opposite weak inequality,

$$reg_M^{\mathcal{P}^+|F}(f) \leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right), \qquad (2)$$

is also implied by Axiom 1. Suppose, by way of contradiction, that (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

$$\begin{aligned} reg_{M'}^{\mathcal{P}^+|F}(a_{f'}) &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_M^{\mathcal{P}^+|(E\cap F)}(f) + \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f)\right) \\ &< reg_M^{\mathcal{P}^+|F}(f) \qquad \text{[by assumption]} \\ &= reg_{M'}^{\mathcal{P}^+|F}(a_f) \qquad \text{[by construction]}. \end{aligned}$$

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 1.

To complete the proof that (a) implies (d), we show that Axiom 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

$$reg_M^{\mathcal{P}^+|F}(f) > \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha \Pr(E^c\cap F) reg_M^{\mathcal{P}^+|(E^c\cap F)}(f). \qquad (3)$$

Suppose, by way of contradiction, that (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 $a_f \notin C_{M,E\cap F}^{reg,\mathcal{P}^+}(M')$ but $a_f \in C_{M,F}^{reg,\mathcal{P}^+}(M')$ 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 $a_f \notin C_{M',E\cap F}^{reg,\mathcal{P}^+}(\{a_{g'}, a_f\})$ but $a_f \in C_{M',F}^{reg,\mathcal{P}^+}(\{a_{g'}, a_f\})$, violating Axiom 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 1–4 hold for menu M and all choice sets $M' \subseteq M$. It is easy to check that Axioms 2–4 hold for MWER, so we need to check only Axiom 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 $f \in C_{M,E\cap F}^{reg,\mathcal{P}^+}(M')$ and $f \in C_{M,E^c\cap F}^{reg,\mathcal{P}^+}(M')$. 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

$$\begin{aligned} reg_{M}^{\mathcal{P}^+|F}(f) &= reg_{M_0}^{\mathcal{P}^+|F}(a_f) \\ &= \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f)\right) \\ &\le \sup_{(\Pr,\alpha)\in\mathcal{P}^+} \alpha\left(\Pr(E\cap F) reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g) + \Pr(E^c\cap F) reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)\right) \\ &= reg_{M}^{\mathcal{P}^+|F}(g), \end{aligned}$$

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 $reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_f) \le reg_{M_0}^{\mathcal{P}^+|(E^c\cap F)}(a_g)$. 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 $0 \le reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_f) < reg_{M_0}^{\mathcal{P}^+|(E\cap F)}(a_g)$. 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 Proof of Theorem 4.5

To prove Theorem 4.5, we need the following lemma.

Lemma C.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. □

The next lemma uses an argument almost identical to one used in Lemma 7 of [Halpern and Leung 2012].

Lemma C.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 $C'(\mathcal{P}^+) = \{p : p\in\mathbb{R}^{|S|} \text{ and } p \le \alpha\Pr \text{ for some } (\Pr,\alpha)\in\mathcal{P}^+\}$. Note that an element $q\in C'(\mathcal{P}^+)$ may not be a subprobability measure, since we do not require that $q(s) \geq 0$. 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

$$\theta \cdot p < c \text{ for all } p \in \overline{C'(\mathcal{P}^+|^\chi F)}, \text{ and } \theta \cdot q > c. \tag{4}$$

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 (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, which we restate here.

THEOREM 4.5. *If $C(\mathcal{P}^+)$ is closed and convex, then Axiom 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 1 does not require $C(\mathcal{P}^+)$ to be convex.

Claim C.3. *If $\mathcal{P}^+$ is χ-rectangular, then Axiom 1 holds for the family of choices $C_M^{reg,\mathcal{P}^+|^\chi E}$.*

Proof. By Theorem 4.3, it suffices to show that SEP holds. For the first part of SEP, we must show that

$$reg_M^{\mathcal{P}^+|^\chi F}(f) = \sup_{(\Pr,\alpha) \in \mathcal{P}^+|^\chi F} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|^\chi (E \cap F)}(f) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^\chi (E^c \cap F)}(f) \right). \tag{5}$$

Unwinding the definitions, (5) is equivalent to

$$\begin{array}{rl} & reg_M^{\mathcal{P}^+|^\chi F}(f) \\ = & \sup_{(\Pr_3,\alpha_3) \in \mathcal{P}^+|^\chi F} \alpha_{\Pr_3} \Big(\Pr_3(E \cap F) \sup_{(\Pr_1,\alpha_1) \in \mathcal{P}^+|^\chi F} \alpha^\chi_{1,E \cap F} \sum_{s \in E \cap F} \Pr_1(s|(E \cap F)) reg_M(f,s)) \\ & + \Pr_3(E^c \cap F) \sup_{(\Pr_2,\alpha_2) \in \mathcal{P}^+|^\chi F} \alpha^\chi_{2,E^c \cap F} \sum_{s \in E^c \cap F} \Pr_2(s|(E^c \cap F)) reg_M(f,s)) \Big). \end{array}$$

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 $(\Pr_1,\alpha_1),(\Pr_2,\alpha_2),(\Pr_3,\alpha_3) \in \mathcal{P}^+|^\chi F$,

$$\alpha_{\Pr_3}\Pr_3(E\cap F)\alpha^\chi_{1,E\cap F}\Pr_1|(E\cap F) + \alpha_{\Pr_3}\Pr_3(E^c\cap F)\alpha^\chi_{2,E^c\cap F}\Pr_2|(E^c\cap F) \in \overline{C(\mathcal{P}^+|^\chi F)}. \quad (6)$$

Thus, for all $\epsilon > 0$,

$$\begin{array}{rl} & \mathit{reg}_M^{\mathcal{P}^+|^\chi F}(f) \\ = & \mathit{reg}_M^{C(\mathcal{P}^+|^\chi F)}(f) \quad \text{[by Lemma C.1]} \\ \geq & \alpha_3^*\Big(\Pr_3^*(E\cap F)(\alpha^*_{1,E\cap F})^\chi \sum_{s\in E\cap F}\Pr_1^*(s|(E\cap F))\mathit{reg}_M(f,s)) \\ & \quad + \Pr_3^*(E^c\cap F)(\alpha^*_{2,E^c\cap F})^\chi \sum_{s\in E^c\cap F}\Pr_2^*(s|(E^c\cap F))\mathit{reg}_M(f,s))\Big) - \epsilon \text{ [by (6)]}. \end{array}$$

Therefore,

$$\begin{array}{rl} & \mathit{reg}_M^{\mathcal{P}^+|^\chi F}(f) \\ \geq & \alpha_3^*\Big(\Pr_3^*(E\cap F)(\alpha^*_{1,E\cap F})^\chi \sum_{s\in E\cap F}\Pr_1^*(s|(E\cap F))\mathit{reg}_M(f,s)) \\ & \quad + \Pr_3^*(E^c\cap F)(\alpha^*_{2,E^c\cap F})^\chi \sum_{s\in E^c\cap F}\Pr_2^*(s|(E^c\cap F))\mathit{reg}_M(f,s))\Big) \\ = & \sup_{(\Pr_3,\alpha_3)\in\mathcal{P}^+|^\chi F}\alpha_3\Big(\Pr_3(E\cap F)\sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\chi F}\alpha^\chi_{1,E\cap F}\sum_{s\in E\cap F}\Pr_1(s|(E\cap F))\mathit{reg}_M(f,s)) \\ & \quad + \Pr_3(E^c\cap F)\sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\chi F}\alpha^\chi_{2,E^c\cap F}\sum_{s\in E^c\cap F}\Pr_2(s|(E^c\cap F))\mathit{reg}_M(f,s))\Big) \\ & \quad \text{[by the choice of } (\Pr_i^*,\alpha_i^*),\ i=1,2,3] \\ = & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F}\alpha\Big(\Pr(E\cap F)\mathit{reg}_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c\cap F)\mathit{reg}_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\Big), \end{array}$$

as required.

It remains to show the opposite inequality in (5), namely, that

$$\mathit{reg}_M^{\mathcal{P}^+|^\chi F}(f) \leq \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F}\alpha\Big(\Pr(E\cap F)\mathit{reg}_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) + \Pr(E^c\cap F)\mathit{reg}_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)\Big).$$

It suffices to note that the right-hand side is equal to

$$\begin{array}{rl} & \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F}\Big(\alpha\Pr(E\cap F)\sup_{(\Pr_1,\alpha_1)\in\mathcal{P}^+|^\chi F}\alpha^\chi_{1,E\cap F}\sum_{s\in E\cap F}\Pr_1(s|E\cap F)\mathit{reg}_M(f,s)) \\ & +\alpha\Pr(E^c\cap F)\sup_{(\Pr_2,\alpha_2)\in\mathcal{P}^+|^\chi F}\alpha^\chi_{2,E^c\cap F}\sum_{s\in E^c\cap F}\Pr_2(s|E^c\cap F)\mathit{reg}_M(f,s))\Big) \\ \geq & \overline{E}_{\mathcal{P}^+|^\chi F}(\mathit{reg}_M(f)) \quad \text{[by rectangularity]} \\ = & \mathit{reg}_M^{\mathcal{P}^+|^\chi F}(f). \end{array}$$

This completes the proof that (5) holds.

For the second part of SEP, suppose that $\overline{\mathcal{P}}^+(E\cap F) > 0$ and $\mathit{reg}_M^{\mathcal{P}^+|^\chi(E\cap F)}(f) \neq 0$. If $\mathit{reg}_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f) = 0$ then, since $\overline{\mathcal{P}}^+(E\cap F) > 0$, we have that $\mathit{reg}_M^{\mathcal{P}^+|^\chi F}(f) > 0 = \sup_{(\Pr,\alpha)\in\mathcal{P}^+|^\chi F}\alpha Pr(E^c\cap F)\mathit{reg}_M^{\mathcal{P}^+|^\chi(E^c\cap F)}(f)$, 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.4. *If $C(\mathcal{P}^+)$ is convex and Axiom 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.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:

$$\begin{array}{l} u(f,s)=-\theta(s), \text{ and} \\ u(g,s)=0 \text{ for all s.} \end{array}$$

The act f will have regret precisely $\theta(s)$ in state $s\in S$. By Lemma C.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 4.3, Axiom 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,

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

violating the second condition of SEP. Therefore, by Theorem 4.3, Axiom 1 does not hold.

Finally, suppose that condition (c) in rectangularity does not hold. Then for some nonnegative real vector $\theta \in \mathbb{R}^{|S|}$,

$$\begin{array}{rl} & \sup_{(\Pr,\alpha) \in \mathcal{P}^+|^{\chi}F} \Big(\alpha \Pr(E) \sup_{(\Pr_1,\alpha_1) \in \mathcal{P}^+|^{\chi}(E \cap F)} \sum_{s \in E \cap F} \alpha_1 \Pr_1(s|E)\theta(s)) \\ & +\alpha \Pr(E^c) \sup_{(\Pr_2,\alpha_2) \in \mathcal{P}^+|^{\chi}(E^c \cap F)} \sum_{s \in E^c \cap F} \alpha_2 \Pr_2(s|E^c)\theta(s))\Big) \\ < & \sup_{(\Pr,\alpha) \in \mathcal{P}^+|^{\chi}F} \alpha \sum_{s \in F} \Pr(s)\theta(s). \end{array} \tag{7}$$

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:

$$\begin{array}{ll} u(g,s) = -\theta(s) & \text{for all } s \in S. \\ u(f,s) = 0 & \text{for all } s \in S. \end{array}$$

So we have

$$\begin{array}{rl} & \sup_{(\Pr,\alpha) \in \mathcal{P}^+|^{p}F} \alpha \left(\Pr(E \cap F) reg_M^{\mathcal{P}^+|^{p}(E \cap F)}(g) + \Pr(E^c \cap F) reg_M^{\mathcal{P}^+|^{p}(E^c \cap F)}(g)\right) \\ = & \sup_{(\Pr,\alpha) \in \mathcal{P}^+|^{\chi}F} \alpha \left(\Pr(E \cap F)\overline{E}_{\mathcal{P}^+|^{\chi}(E \cap F)}(\theta) + \Pr(E^c \cap F)\overline{E}_{\mathcal{P}^+|^{\chi}(E^c \cap F)}(\theta)\right) \\ < & \overline{E}_{\mathcal{P}^+|^{\chi}F}(\theta) \quad \text{[by (7)]} \\ = & reg_M^{\mathcal{P}^+|^{p}F}(g). \end{array}$$

This means that SEP, and hence Axiom 1, is violated, a contradiction. □

References

Al-Najjar, N. and J. Weinstein (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy 25*, 249–284.

Allais, M. (1953). 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.

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.

Chateauneuf, A. and J. Faro (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics 45*, 535 – 558.

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)*.

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.

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.

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. (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.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). A smooth model of decision making under ambiguity. *Econometrica 73*(6), 1849–1892.

Kreps, D. M. (1988). *Notes on the Theory of Choice*. Boulder, Colo.: Westview Press.

Levi, I. (1985). Imprecision and uncertainty in probability judgment. *Philosophy of Science 52*, 390–406.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006a). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006b, May). Dynamic variational preferences. *Journal of Economic Theory 128*(1), 4–44.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature 27*(4), 1622–68.

Nau, R. F. (1992). Indeterminate probabilities on finite sets. *Annals of Statistics 40*(4), 1737–1767.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton, N. J.: Princeton University Press.

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.

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. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory 146*(6), 2226 – 2251.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies 23*(3), 165–180.

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.

Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems 1*, 3–28.