

Weighted Regret-Based Likelihood: A New Approach to Describing Uncertainty

Joseph Y. Halpern

Computer Science Department

Cornell University

Ithaca, NY 14853, USA

HALPERN@CS.CORNELL.EDU

Abstract

Recently, Halpern and Leung suggested representing uncertainty by a set of *weighted* probability measures, and suggested a way of making decisions based on this representation of uncertainty: *maximizing weighted regret*. Their paper does not answer an apparently simpler question: what it means, according to this representation of uncertainty, for an event E to be more likely than an event E' . In this paper, a notion of comparative likelihood when uncertainty is represented by a set of weighted probability measures is defined. It generalizes the ordering defined by probability (and by lower probability) in a natural way; a generalization of upper probability can also be defined. A complete axiomatic characterization of this notion of regret-based likelihood is given.

1. Introduction

Recently, Samantha Leung and I (Halpern & Leung, 2012) suggested representing uncertainty by a set of *weighted* probability measures, and suggested a way of making decisions based on this representation of uncertainty: *maximizing weighted regret*. However, we did not answer an apparently simpler question: given this representation of uncertainty, what does it mean for an event E to be more likely than an event E' ? This is what I do in this paper. To explain the issues, I start by reviewing the Halpern-Leung approach.

It has frequently been observed that there are many situations where an agent's uncertainty is not adequately described by a single probability measure. Specifically, a single measure may not be adequate for representing an agent's ignorance. For example, there seems to be a big difference between a coin known to be fair and a coin whose bias an agent does not know, yet if the agent were to use a single measure to represent her uncertainty, in both of these cases it would seem that the measure that assigns heads probability $1/2$ would be used.

One approach that has been suggested for representing ignorance is to use a set \mathcal{P} of probability measures. This idea is an old one, apparently going back to the work of Boole (1854, ch. 16–21) and Ostrogradsky (1838); some authors (e.g., Campos & Moral, 1995; Couso, Moral, & Walley, 1999; Gilboa & Schmeidler, 1993; Levi, 1985; Walley, 1991) have additionally required the set \mathcal{P} to be *convex* (so that if μ_1 and μ_2 are in \mathcal{P} , then so is $a\mu_1 + b\mu_2$, where $a, b \in [0, 1]$ and $a + b = 1$). This approach has the benefit of representing uncertainty in general, not by a single number, but by a range of numbers. This allows us to distinguish the certainty that a coin is fair (in which case the uncertainty of heads is represented by a single number, $1/2$) from knowing only that the probability of heads could be anywhere between, say, $1/3$ and $2/3$.

But this approach also has its problems. For example, consider an agent who believes that a coin may have a slight bias. Thus, although it is unlikely to be completely fair, it is close to fair. How should we represent this with a set of probability measures? Suppose that the agent is quite sure that the bias is between $1/3$ and $2/3$. We could, of course, take \mathcal{P} to consist of all the measures that give heads probability between $1/3$ and $2/3$. But how does the agent know that the possible biases are *exactly* between $1/3$ and $2/3$. Does she not consider $2/3 + \epsilon$ possible for some small ϵ ? And even if she is confident that the bias is between $1/3$ and $2/3$, this representation cannot take into account the possibility that she views biases closer to $1/2$ as more likely than biases further from $1/2$.

There is also a second well-known concern: learning. Suppose that the agent initially considers possible all the measures that gives heads probability between $1/3$ and $2/3$. She then starts tossing the coin, and sees that, of the first 20 tosses, 12 are heads. It seems that the agent should then consider a bias of greater than $1/2$ more likely than a bias of less than $1/2$. But if we use the standard approach to updating with sets of probability measures (Halpern, 2003), and condition each of the measures on the observation, since the coin tosses are viewed as independent, the agent will continue to believe that the probability of the next coin toss is between $1/3$ and $2/3$. The observation has no impact as far as learning to predict better. The set \mathcal{P} stays the same, no matter what observation is made.

There is a well-known solution to these problems: putting a measure of uncertainty on these probability measures in \mathcal{P} . This idea too has a long history. One special case is to put a “second-order” probability on these probability measures; see (Good, 1980) for discussion of this approach and further references. For example, an agent can express the fact that the bias of a coin is more likely to be close to $1/2$ than far from $1/2$. In addition, the problem of learning can be dealt with by straightforward conditioning. But this approach leads to other problems. Essentially, it seems that the ambiguity that an agent might feel about the outcome of the coin toss seems to have disappeared. For example, suppose that the agent has no idea what the bias is. The obvious second-order probability to use is the uniform probability on possible biases. While we cannot talk about the probability that the coin is heads (there is a set of probabilities, after all, not a single probability), the *expected* probability of heads is $1/2$. Why should an agent that has no idea of the bias of the coin know or believe that the expected probability of heads is $1/2$? Of course, if one had to use a single probability measure to describe uncertainty, symmetry considerations dictate that it should be the one that ascribes equal likelihood to heads and tails; similarly, if one had to put a single second-order probability on the set of possible biases, uniform probability seems like the most obvious choice. Moreover, if our interest is in making decisions, then maximizing the expected utility using the expected probability again does not take the agent’s ignorance into account. Kyburg (1988) and Pearl (1987) have even argued that there is no need for a second-order probability on probabilities; whatever can be done with a second-order probability can already be done with a basic probability.

Nevertheless, when it comes to decision-making, it does seem useful to use an approach that represents ambiguity, while still maintaining some of the features of having a second-order probability on probabilities. This idea goes back to at least Gärdenfors and Sahlin (1982, 1983). Walley (1997) suggested putting a *possibility measure* (Dubois & Prade, 1998; Zadeh, 1978) on probability measures; this was also essentially done by Cattaneo (2007), Chateaufneuf and Faro (2009), and de Cooman (2005). All of these authors and others, such

as Klibanoff et al. (2005), Maccheroni et al. (2006), and Nau (1992), proposed approaches to decision making using their representations of uncertainty.

Leung and I similarly suggested putting weights on each probability measure in \mathcal{P} . Since we assumed that the weights are normalized so that the supremum of the weights is 1, these weights can also be viewed as a possibility measure. If the set \mathcal{P} is finite, we can also normalize so as to view the weights as being second-order probabilities. As with second-order probabilities, the weights can vary over time, as more information is acquired. For example, we can start with a state of complete ignorance (modeled by assuming that all probability measures have weight 1), and update the weights after making an observation ob , we take the weight of a measure Pr to be the relative likelihood of ob if Pr were the true measure. (See Section 2 for details.) With this approach, called *likelihood updating* by Halpern and Leung (2012), if there is a true underlying measure generating the data, over time, the weight of the true measure approaches 1, while the weight of all other measures approaches 0. Thus, this approach allows learning in a natural way. If, for example, the actual bias of the coin was $5/8$ in the example above, no matter what the initial weights, as long as $5/8$ had positive weight, then its weight would almost surely converge to 1 as more observations were made, while the weight of all other measures would approach 0. This, of course, is exactly what would happen if we had a second-order probability on \mathcal{P} . The weights can also be used to represent the fact that some probabilities in the set \mathcal{P} are more likely than others.

Like essentially all others who considered a representation of uncertainty based on a set of probability with weights, Leung and I also suggested a way of using this representation to make decisions. However, our approach was different than those suggested earlier. We based our approach on *regret*, a standard approach to decision-making that was introduced (independently) by Niehans (1948) and Savage (1951). If uncertainty is represented by a set \mathcal{P} of probability measures, then regret works as follows: for each act a and each measure $\text{Pr} \in \mathcal{P}$, we can compute the expected regret of a with respect to Pr ; this is the difference between the expected utility of a and the expected utility of the act that gives the highest expected utility with respect to Pr . We can then associate with an act a its worst-case expected regret of a , over all measures $\text{Pr} \in \mathcal{P}$, and compare acts with respect to their worst-case expected regret. With weights in the picture, we modify the procedure by multiplying the expected regret associated with measure Pr by the weight of Pr , and compare acts according to their worst-case *weighted* expected regret. This approach to making decisions is very different from the others mentioned above that incorporate a likelihood on probabilities. Moreover, using the weights in the way means that we cannot simply replace a set of weighted probability measures by a single probability measure; the objections of Kyburg (1988) and Pearl (1987) do not apply.

Leung and I (Halpern & Leung, 2012) show that this approach seems to do reasonable things in a number of examples of interest, and provide an axiomatization of decision-making with this approach. Since sets of weighted probabilities are certainly intended to be a way of representing uncertainty, it seems natural to ask whether they can be used to represent relative likelihood in a direct way. Surprisingly, this is something largely not considered in earlier papers using sets of weighted probabilities, since their focus was on decision-making (although the work of Nau discussed in Section 3 is an exception).

Representing relative likelihood is straightforward if uncertainty is represented by a single probability measure: E is more likely than E' exactly if the probability of E is greater than the probability of E' . When using sets of probability measures, various approaches have been considered in the literature. The most common takes E to be more likely than E' if the *lower probability* of E is greater than the lower probability of E' , where the lower probability of E is its worst-case probability, taken over the measures in \mathcal{P} (see Section 3). We could also compare E and E' with respect to their upper probabilities (the best-case probability with respect to the measures in \mathcal{P}). Another possibility is to take E to be more likely than E' if $\Pr(E) \geq \Pr(E')$ for all measures $\Pr \in \mathcal{P}$; this gives a partial order on likelihood.¹ But what should we do if uncertainty is represented by a set of weighted probability measures?

In this paper, I define a notion of relative likelihood when uncertainty is represented by a set of weighted probability measures that generalizes the ordering defined by lower probability in a natural way; I also define a generalization of upper probability. We can then associate with an event E two numbers that are analogues of lower and upper probability. If uncertainty is represented by a single measure, then these two numbers coincide; in general, they do not. The interval can be thought of as representing the degree of ambiguity in the likelihood of E . Indeed, in the special case when all the weights are 1, the numbers are essentially just the lower and upper probability (technically, they are 1 minus the lower and upper probability, respectively). Interestingly, the approach to assigning likelihood is based on the approach to decision-making. Essentially, what I am doing is the analogue of defining probability in terms of expected utility, rather than the other way around. The approach can be viewed as generalizing both probability and lower probability, while at the same time allowing a natural approach to updating.

Why we should be interested in such a representation? If all that we ever did with probability was to use it to make decisions, then arguably this wouldn't be of much interest; my work with Leung already shows how sets of weighted probabilities can be used in decision-making. The results of this paper add nothing further to that question. However, we often talk about the likelihood of events quite independent of their use in decision-making. There are clearly many examples in physics. The issue arises in AI applications as well: a typical explanation of why we did A rather than B is that we thought some event E was more likely than F . And computations of expectation, which clearly involve a representation of uncertainty, arise in many AI applications. Thus, having an analogue of probability seems important and useful in its own right.

The rest of this paper is organized as follows. After reviewing the relevant material from (Halpern & Leung, 2012) in Section 2, I define regret-based likelihood in Section 3, and compare it to lower probability. I provide an axiomatic characterization of regret-based likelihood in Section 4, and show how it relates to the axiomatic characterization of lower probability. I conclude in Section 5.

1. There is a long tradition of considering partially ordered notions of likelihood; see (Halpern, 1997) and the references therein, and the work of Walley (1991).

2. Weighted Expected Regret: A Review

Consider the standard setup in decision theory. We have a state space S and an outcome space O . An *act* is a function from S to O ; it describes an outcome for each state. Suppose that we have a utility function u on outcomes and a set \mathcal{P}^+ of *weighted probability measures*. That is, \mathcal{P}^+ consists of pairs $(\text{Pr}, \alpha_{\text{Pr}})$, where α_{Pr} is a weight in $[0, 1]$ and Pr is a probability on S . Let $\mathcal{P} = \{\text{Pr} : \exists \alpha((\text{Pr}, \alpha) \in \mathcal{P}^+)\}$. For each $\text{Pr} \in \mathcal{P}$ there is assumed to be exactly one α , denoted α_{Pr} , such that $(\text{Pr}, \alpha) \in \mathcal{P}^+$. It is further assumed that weights have been normalized so that there is at least one measure $\text{Pr} \in \mathcal{P}$ such that $\alpha_{\text{Pr}} = 1$. Finally, \mathcal{P}^+ is assumed to be *weakly closed*, so that if $(\text{Pr}_n, \alpha_n) \in \mathcal{P}^+$ for $n = 1, 2, 3, \dots$, $(\text{Pr}_n, \alpha_n) \rightarrow (\text{Pr}, \alpha_{\text{Pr}})$, and $\alpha_{\text{Pr}} > 0$, then $(\text{Pr}, \alpha_{\text{Pr}}) \in \mathcal{P}^+$. (I discuss below why I require \mathcal{P}^+ to be just weakly closed, rather than closed.)

The assumption that at least one probability measure has a weight of 1 is convenient for comparison to other approaches; see below. However, making this assumption has no impact on the results of this paper; as long as we restrict to sets where the weight is bounded, all the results hold without change. This assumption is, of course, incompatible with the weights being probabilities. Note that the assumption that the weights are probabilities runs into difficulties if we have an infinite number of measures in \mathcal{P} ; for example, if \mathcal{P} includes all measures on heads from $1/3$ to $2/3$, as discussed in the Introduction, using a uniform probability, we would be forced to assign each individual probability measure a weight of 0, which would not work well for our later definitions.

Where are the weights in \mathcal{P}^+ coming from? In general, they can be viewed as subjective, just like the probability measures. However, as Leung and I (Halpern & Leung, 2012) observed, there is an important special case where the weights can be given a natural interpretation. Suppose that, as in the case of the biased coin in the Introduction, we make observations in a situation where the probability of making a given observation is determined by some objective source. Then we can start by giving all probability measures a weight of 1. Given an observation *ob* (e.g., sequence of coin tosses in the example in the Introduction), we can compute $\text{Pr}(ob)$ for each measure $\text{Pr} \in \mathcal{P}$; we can then update the weight of Pr to be $\text{Pr}(ob) / \sup_{\text{Pr}' \in \mathcal{P}} \text{Pr}'(ob)$. Thus, the more likely the observation is according to Pr , the higher the updated weight of Pr relative to other probability measures in \mathcal{P} .² (The denominator is just a normalization to ensure that some measure has weight 1.) With this approach to updating, if there is a true underlying measure generating the data, then as an agent makes more observations, almost surely, the weight of the true measure approaches 1, while the weight of all other measures approaches 0.³ In addition, this approach gives an agent a natural way of determining weights for each probability measure in \mathcal{P} . While, in general, this means that the agent may need to carry around a lot of information (not

2. The idea of putting a possibility on probabilities in \mathcal{P} that is determined by likelihood also appears in the work of Moral (1992), although he does not consider a general approach to dealing with sets of weighted probability measures.

3. The “almost surely” is due to the fact that, with probability approaching 0, as more and more observations are made, it is possible that an agent will make misleading observations that are not representative of the true measure. This also depends on the set of possible observations being rich enough to allow the agent to ultimately discover the true measure generating the observations; for example, an agent will never learn the distributions of outcomes of a die she never gets to observe the die when it lands 5 or 6. Since learning is not a focus of this paper, I do not make this notion of “rich enough” precise here.

only a possibly infinite set of probabilities, but a weight associated with each one), if the set \mathcal{P} has a reasonable parametric representation, then the weight can often be evaluated in terms of the parameters, so should admit a compact representation (see Example 3.2).

The weight associated with a probability Pr can be viewed as an upper bound on an agent’s confidence that Pr actually describes the situation. That is why an agent who has no idea of what is going on is modeled as starting by placing weight 1 on all probability measures. I believe that having the weights will allow agents to express nuances that they consider important, and that such weights will not be hard to elicit. Whether this is the case is really an empirical question, one which I believe deserves further exploration, but is beyond the scope of this paper.

I now review the definition of weighted regret, and introduce the notion of *absolute* (weighted) regret. I start with regret. The regret of an act a in a state $s \in S$ is the difference between the utility of the best act at state s and the utility of a at s . Typically, the act a is not compared to all acts, but to the acts in a set M , called a *menu*. Thus, the regret of a in state s relative to menu M , denoted $\text{reg}^M(a, s)$, is $\sup_{a' \in M} u(a'(s)) - u(a(s))$.⁴ There are typically some constraints put on M to ensure that $\sup_{a' \in M} u(a'(s))$ is finite—this is certainly the case if M is finite, or the convex closure of a finite set of acts, or if there is a best possible outcome in the outcome space O . The latter assumption holds in this paper, so I assume throughout that $\sup_{a' \in M} u(a'(s))$ is finite.

For simplicity, I assume that the state space S is finite. Given a probability measure Pr on S , the expected regret of an act a with respect to Pr relative to menu M is just $\text{reg}_{\text{Pr}}^M(a) = \sum_{s \in S} \text{reg}^M(a, s) \text{Pr}(s)$. The (*expected*) *regret* of a with respect to \mathcal{P} and a menu M is just the worst-case regret, that is,

$$\text{reg}_{\mathcal{P}}^M(a) = \sup_{\text{Pr} \in \mathcal{P}} \text{reg}_{\text{Pr}}^M(a).$$

Similarly, the *weighted (expected) regret* of a with respect to \mathcal{P}^+ and a menu M is just the worst-case weighted regret, that is,

$$\text{wr}_{\mathcal{P}^+}^M(a) = \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \text{reg}_{\text{Pr}}^M(a).$$

Thus, regret is just a special case of weighted regret, where all weights are 1.

Note that, as far weighted regret goes, it does not hurt to augment a set \mathcal{P}^+ of weighted probability measures by adding pairs of the form $(\text{Pr}, 0)$ for $\text{Pr} \notin \mathcal{P}$. But if we start with a set \mathcal{P} of unweighted probability measures, the set $\mathcal{P}^+ = \{(\text{Pr}, 1) : \text{Pr} \in \mathcal{P}\} \cup \{(\text{Pr}, 0) : \text{Pr} \notin \mathcal{P}\}$ is not closed in general, although it is weakly closed. There may well be a sequence $\text{Pr}_n \rightarrow \text{Pr}$, where $\text{Pr}_n \notin \mathcal{P}$ for all n , but $\text{Pr} \in \mathcal{P}$. But then we would have $(\text{Pr}_n, 0) \in \mathcal{P}^+$ converging to $(\text{Pr}, 0) \notin \mathcal{P}^+$. This is exactly why I required only weak closedness. Note for future reference that, since \mathcal{P}^+ is assumed to be weakly closed, if $\text{wr}_{\mathcal{P}^+}^M(a) > 0$, then there is some element $(\text{Pr}, \alpha_{\text{Pr}}) \in \mathcal{P}^+$ such that $\text{wr}_{\mathcal{P}^+}^M(a) = \alpha_{\text{Pr}} \text{reg}_{\text{Pr}}^M(a)$.

Weighted regret induces an obvious preference order on acts: act a is at least as good as a' with respect to \mathcal{P}^+ and M , written $a \succeq_{\mathcal{P}^+, M}^{\text{reg}} a'$, if $\text{wr}_{\mathcal{P}^+}^M(a) \leq \text{wr}_{\mathcal{P}^+}^M(a')$. As usual, I

4. Recall that if X is a set of real numbers, $\sup X$, the *supremum* of X , is the smallest real numbers that is greater than or equal to all the elements of X . If X is finite, then the sup is the same as the max. But if X is, say, the interval $(0, 1)$, then $\sup X = 1$. Similarly, $\inf X$ is the largest real number that is less than or equal to all the elements in X .

write $a \succ_{\mathcal{P}^+, M}^{reg} a'$ if $a \succeq_{\mathcal{P}^+, M}^{reg} a'$ but it is not the case that $a' \succeq_{\mathcal{P}^+, M}^{reg} a$. The standard notion of regret is the special case of weighted regret where all weights are 1. I sometimes write $a \succeq_{\mathcal{P}, M}^{reg} a'$ to denote the unweighted case (i.e., where all the weights in \mathcal{P}^+ are 1).

In this setting, using weighted regret gives an approach that allows an agent to transition smoothly from regret to expected utility. It is well known that regret generalizes expected utility in the sense that if \mathcal{P} is a singleton $\{\text{Pr}\}$, then $wr_{\mathcal{P}}^M(a) \leq wr_{\mathcal{P}}^M(a')$ iff $\text{EU}_{\text{Pr}}(a) \geq \text{EU}_{\text{Pr}}(a')$ (where $\text{EU}_{\text{Pr}}(a)$ denotes the expected utility of act a with respect to probability Pr); this follows from the observation that, given a menu M , there is a constant c_M such that, for all acts $a \in M$, $wr_{\{\text{Pr}\}}^M(a) = c_M - \text{EU}_{\text{Pr}}(a)$. (In particular, this means that if \mathcal{P} is a singleton, regret is menu independent.) If we start with all the weights being 1, then, as observed above, the weighted regret is just the standard notion of regret. As the agent makes observations, if there is a measure Pr generating the uncertainty, the weights will get closer and closer to a situation where Pr gets weight 1, with the weights of all other measures dropping off quickly to 0, so the ordering of acts will converge to the ordering given by expected utility with respect to Pr .

There is another approach with some similar properties, which again starts with uncertainty being represented by a set \mathcal{P} of (unweighted) probability measures. Define $wc_{\mathcal{P}}(a) = \inf_{\text{Pr} \in \mathcal{P}} \text{EU}_{\text{Pr}}(a)$. Thus $wc_{\mathcal{P}}(a)$ is the worst-case expected utility of a , taken over all $\text{Pr} \in \mathcal{P}$. Then define $a \succeq_{\mathcal{P}}^{mm} a'$ if $wc_{\mathcal{P}}(a) \geq wc_{\mathcal{P}}(a')$. This is the maxmin expected utility rule, quite often used in economics (Gilboa & Schmeidler, 1989). There are difficulties in getting a weighted version of maxmin expected utility (Halpern & Leung, 2012) (discussed further in Section 3); however, Epstein and Schneider (2007) propose another approach that can be combined with maxmin expected utility. They fix a parameter $\alpha \in (0, 1)$, and update \mathcal{P} after an observation ob by retaining only those measures Pr such that $\text{Pr}(ob) \geq \alpha$. For any choice of $\alpha < 1$, we again end up converging almost surely to a single measure, so again this approach converges almost surely to expected utility.

I conclude this section with a discussion of menu dependence. Maxmin expected utility is not menu dependent; the preference ordering on acts induced by regret can be, as the following example illustrates.

Example 2.1: Take the outcome space to be $\{0, 1\}$, and the utility function to be the identity, so that $u(1) = 1$ and $u(0) = 0$. As usual, if $E \subseteq S$, 1_E denotes the *indicator function* on E , where, for each state $s \in S$, we have $1_E(s) = 1$ if $s \in E$, and $1_E(s) = 0$ if $s \notin E$. Let $S = \{s_1, s_2, s_3, s_4\}$, $E_1 = \{s_1\}$, $E_2 = \{s_2\}$, $E_3 = \{s_2, s_3\}$, $M_1 = \{1_{E_1}, 1_{E_2}\}$, $M_2 = \{1_{E_1}, 1_{E_2}, 1_{E_3}\}$, and $\mathcal{P} = \{\text{Pr}_1, \text{Pr}_2\}$, where $\text{Pr}_1(s_1) = \text{Pr}_1(s_3) = \text{Pr}_1(s_4) = 1/3$, $\text{Pr}_2(s_2) = 1/4$, and $\text{Pr}_2(s_3) = 3/4$. A straightforward calculation shows that $reg_{\text{Pr}_1}^{M_1}(1_{E_1}) = 0$, $reg_{\text{Pr}_1}^{M_1}(1_{E_2}) = 1/3$, $reg_{\text{Pr}_2}^{M_1}(1_{E_1}) = 1/4$, $reg_{\text{Pr}_2}^{M_1}(1_{E_2}) = 0$, $reg_{\text{Pr}_1}^{M_2}(1_{E_1}) = 1/3$, $reg_{\text{Pr}_1}^{M_2}(1_{E_2}) = 2/3$, $reg_{\text{Pr}_2}^{M_2}(1_{E_1}) = 1$, and $reg_{\text{Pr}_2}^{M_2}(1_{E_2}) = 3/4$. Thus, $1/4 = reg_{\mathcal{P}}^{M_1}(1_{E_1}) < reg_{\mathcal{P}}^{M_1}(1_{E_2}) = 1/3$, while $1 = reg_{\mathcal{P}}^{M_2}(1_{E_1}) > reg_{\mathcal{P}}^{M_2}(1_{E_2}) = 3/4$. The preference on 1_{E_1} and 1_{E_2} depends on whether we consider the menu M_1 or the menu M_2 . ■

Suppose that there is an outcome $o^* \in O$ that gives the maximum utility; that is, $u(o^*) \geq u(o)$ for all $o \in O$. If \bar{o}^* is the constant act that gives outcomes o^* in all states, then \bar{o}^* is clearly the best act in all states. If there is such a best act, an “absolute”, menu-independent notion of weighted expected regret can be defined by always comparing

to \bar{o}^* . That is, define

$$\begin{aligned} \text{reg}(s, a) &= u(o^*) - u(a(s)); \\ \text{reg}_{\text{Pr}}(a) &= \sum_{s \in S} (u(o^*) - u(a(s)) \text{Pr}(s)) = u(o^*) - \text{EU}_{\text{Pr}}(a); \\ \text{reg}_{\mathcal{P}}(a) &= \sup_{\text{Pr} \in \mathcal{P}} \sum_{s \in S} (u(o^*) - u(a(s)) \text{Pr}(s)) = u(o^*) - \inf_{\text{Pr} \in \mathcal{P}} (\text{EU}_{\text{Pr}}(a)); \\ \text{wr}_{\mathcal{P}^+}(a) &= \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \sum_{s \in S} (u(o^*) - u(a(s)) \text{Pr}(s)) = \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} (u(o^*) - \text{EU}_{\text{Pr}}(a)). \end{aligned}$$

If there is a best act, then I write $a \succeq_{\mathcal{P}^+} a'$ if $\text{wr}_{\mathcal{P}^+}(a) \leq \text{wr}_{\mathcal{P}^+}(a')$; similarly in the unweighted case, I write $a \succeq_{\mathcal{P}} a'$ if $\text{wr}_{\mathcal{P}}(a) \leq \text{wr}_{\mathcal{P}}(a')$.

Conceptually, we can think of the agent as always being aware of the best outcome o^* , and comparing his actual utility with a to $u(o^*)$. Equivalently, the absolute notion of regret is equivalent to a menu-based notion with respect to a menu M that includes \bar{o}^* (since if the menu includes \bar{o}^* , it is the best act in every state). As we shall see, in our setting, we can always reduce menu-dependent regret to this absolute, menu-independent notion, since there is in fact a best act: 1_S .

3. Relative Ordering of Events Using Weighted Regret

In this section, I consider how a notion of comparative likelihood can be defined using sets of weighted probability measures.

As in Example 2.1, take the outcome space to be $\{0, 1\}$, the utility function to be the identity, and consider indicator functions. It is easy to see that $\text{EU}_{\text{Pr}}(1_E) = \text{Pr}(E)$, so that with this setup, we can recover probability from expected utility. Thus, if uncertainty is represented by a single probability measure Pr and we make decisions by preferring those acts that maximize expected utility, then we have $1_E \succeq_{1_{E'}} 1_{E'}$ iff $\text{Pr}(E) \geq \text{Pr}(E')$.

Consider what happens if we apply this approach to maxmin expected utility. Now we have that $1_E \succeq_{\mathcal{P}}^{mm} 1_{E'}$ iff $\inf_{\text{Pr} \in \mathcal{P}} \text{Pr}(E) \geq \inf_{\text{Pr} \in \mathcal{P}} \text{Pr}(E')$. In the literature, $\inf_{\text{Pr} \in \mathcal{P}} \text{Pr}(E)$, denoted $\mathcal{P}_*(E)$, is called the *lower probability* of E , and is a standard approach to describing likelihood. The dual *upper probability*, $\sup_{\text{Pr} \in \mathcal{P}} \text{Pr}(E)$, is denoted $\mathcal{P}^*(E)$. An easy calculation shows that

$$\mathcal{P}^*(E) = 1 - \mathcal{P}_*(\bar{E}),$$

where, as usual, \bar{E} denotes the complement of E . The interval $[\mathcal{P}_*(E), \mathcal{P}^*(E)]$ can be thought of as describing the uncertainty of E ; the larger the interval, the greater the ambiguity.

What happens if we apply this approach to regret? First consider unweighted regret. If we restrict to acts of the form 1_E , then the best act is clearly 1_S , which is just the constant function 1. Thus, we can (and do) use the absolute notion of regret here, and for the remainder of this paper. We then get that $1_E \succeq_{\mathcal{P}}^{reg} 1_{E'}$ iff $\sup_{\text{Pr} \in \mathcal{P}} (1 - \text{Pr}(E)) \leq \sup_{\text{Pr} \in \mathcal{P}} (1 - \text{Pr}(E'))$ iff $\sup_{\text{Pr} \in \mathcal{P}} \text{Pr}(\bar{E}) \leq \sup_{\text{Pr} \in \mathcal{P}} \text{Pr}(\bar{E}')$; that is,

$$1_E \succeq_{\mathcal{P}}^{reg} 1_{E'} \text{ iff } \mathcal{P}^*(\bar{E}) \leq \mathcal{P}^*(\bar{E}').$$

Moreover, easy manipulation shows that $\sup_{\text{Pr} \in \mathcal{P}} (1 - \text{Pr}(E)) = 1 - \inf_{\text{Pr} \in \mathcal{P}} \text{Pr}(E) = 1 - \mathcal{P}_*(E)$. It follows that

$$\begin{aligned} & 1_E \succeq_{\mathcal{P}}^{\text{reg}} 1_{E'} \\ \text{iff } & (1 - \mathcal{P}_*(E)) \leq (1 - \mathcal{P}_*(E')) \\ \text{iff } & \mathcal{P}_*(E) \geq \mathcal{P}_*(E') \\ \text{iff } & 1_E \succeq_{\mathcal{P}}^{\text{mm}} 1_{E'}. \end{aligned}$$

That is, both regret and maxmin expected utility put the same ordering on events.

The extension to weighted regret is immediate. Let $\mathcal{P}_{\text{reg}}^+(E)$, the (*weighted*) *regret-based likelihood of E*, be defined by taking

$$\mathcal{P}_{\text{reg}}^+(E) = \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \text{Pr}(\bar{E}).$$

If \mathcal{P}^+ is unweighted, so that all the weights are 1, I write $\mathcal{P}_{\text{reg}}(E)$ to denote $\sup_{\text{Pr} \in \mathcal{P}} \text{Pr}(\bar{E})$. Note that $\mathcal{P}_{\text{reg}}(E) = 1 - \mathcal{P}_*(E)$, so

$$\mathcal{P}_{\text{reg}}(E) \leq \mathcal{P}_{\text{reg}}(E') \text{ iff } \mathcal{P}_*(E) \geq \mathcal{P}_*(E').$$

That is, the ordering induced by \mathcal{P}_{reg} is the opposite of that induced by \mathcal{P}_* . So, for example, $\mathcal{P}_{\text{reg}}(\emptyset) = 1$ and $\mathcal{P}_{\text{reg}}(S) = 0$; smaller sets have larger regret-based likelihood. However, since an act with smaller regret is viewed as better, the ordering on acts of the form 1_E induced by regret is the same as that induced by maxmin expected utility.

Regret-based likelihood provides a way of associating a number with each event, just as probability and lower probability do. Moreover, just as lower probability gives a lower bound on uncertainty, we can think of $\mathcal{P}_{\text{reg}}^+(E)$ as giving an upper bound on the uncertainty. (It is an upper bound rather than a lower bound because larger regret means less likely, just as smaller lower probability does.) The naive corresponding lower bound is given by $\inf_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \text{Pr}(\bar{E})$. This lower bound is not terribly interesting; if there are probability measures $\text{Pr}' \in \mathcal{P}$ such that $\alpha_{\text{Pr}'}$ is close to 0, then this lower bound will be close to 0, independent of the agent's actual feeling about the likelihood of E . A more reasonable lower bound is given by the expression $\underline{\mathcal{P}}_{\text{reg}}^+(E) = 1 - \mathcal{P}_{\text{reg}}^+(\bar{E})$ (recall that the analogous expression relates upper probability and lower probability). The intuition for this choice is the following. If nature were conspiring against us, she would try to prove us wrong by making $\alpha_{\text{Pr}} \text{Pr}(\bar{E})$ as large as possible—that is, make the weighted probability of being wrong as large as possible. On the other hand, if nature were conspiring with us, she would try to make $\alpha_{\text{Pr}} \text{Pr}(E)$ as large as possible, or, equivalently, make $1 - \alpha_{\text{Pr}} \text{Pr}(E)$ as small as possible. Note that this is different from making $\alpha_{\text{Pr}} \text{Pr}(\bar{E})$ as large as possible, unless $\alpha_{\text{Pr}} = 1$ for all $\text{Pr} \in \mathcal{P}$. An easy calculation shows that

$$\begin{aligned} 1 - \mathcal{P}_{\text{reg}}^+(\bar{E}) &= 1 - \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \text{Pr}(E) \\ &= \inf_{\text{Pr} \in \mathcal{P}} (1 - \alpha_{\text{Pr}} \text{Pr}(E)). \end{aligned}$$

This motivates the definition of $\underline{\mathcal{P}}_{\text{reg}}^+$.

The following lemma clarifies the relationship between these expressions, and shows that $[\underline{\mathcal{P}}_{\text{reg}}^+(E), \mathcal{P}_{\text{reg}}^+(E)]$ really does give an interval of ambiguity.

Lemma 3.1: $\inf_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \text{Pr}(\bar{E}) \leq 1 - \mathcal{P}_{\text{reg}}^+(\bar{E}) \leq \underline{\mathcal{P}}_{\text{reg}}^+(E)$.

Proof: Clearly

$$\inf_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \Pr(\bar{E}) = \inf_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} (1 - \Pr(E)).$$

Since, as observed above,

$$1 - \mathcal{P}_{reg}^+(\bar{E}) = \inf_{\text{Pr} \in \mathcal{P}} (1 - \alpha_{\text{Pr}} \Pr(E)),$$

and for all $\text{Pr} \in \mathcal{P}$, we have

$$1 - \alpha_{\text{Pr}} \Pr(E) \geq \alpha_{\text{Pr}} (1 - \Pr(E)),$$

it follows that $\inf_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \Pr(\bar{E}) \leq 1 - \mathcal{P}_{reg}^+(\bar{E})$.

Since, by assumption, there is a probability measure $\text{Pr}' \in \mathcal{P}$ such that $\alpha_{\text{Pr}'} = 1$, it follows that

$$\begin{aligned} 1 - \mathcal{P}_{reg}^+(\bar{E}) &= 1 - \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \Pr(E) \\ &\leq 1 - \text{Pr}'(E) \\ &= \text{Pr}'(\bar{E}) \\ &\leq \sup_{\text{Pr} \in \mathcal{P}} \alpha_{\text{Pr}} \Pr(\bar{E}) \\ &\leq \mathcal{P}_{reg}^+(E). \end{aligned}$$

■

In general, equality does not hold in Lemma 3.1, as shown by the following example. The example also illustrates how the “ambiguity interval” can decrease with weighted regret, if the weights are updated as Leung and I (Halpern & Leung, 2012) suggested.

Example 3.2: Suppose that the state space consists of $\{h, t\}$ (for heads and tails); let Pr_β be the measure that puts probability β on h . Let $\mathcal{P}_0^+ = \{(\text{Pr}_\beta, 1) : 1/3 \leq \beta \leq 2/3\}$. That is, we initially consider all the measures that put probability between $1/3$ and $2/3$ on heads. We toss the coin and observe it lands heads. Intuitively, we should now consider it more likely that the probability of heads is greater than $1/2$. Indeed, applying likelihood updating, we get the set $\mathcal{P}_1^+ = \{(\text{Pr}_\beta, 3\beta/2) : 1/3 \leq \beta \leq 2/3\}$; the probability measures that give h higher probability get higher weight. In particular, the weight of $\text{Pr}_{2/3}$ is still 1, but the weight of $\text{Pr}_{1/3}$ is only $1/2$. (The weight of Pr_β is the likelihood of observing heads according to Pr_β , which is just β , normalized by the likelihood of observing heads according to the measure that gives heads the highest probability, namely $2/3$.) If the coin is tossed again and this time tails is observed, we update further to get $\mathcal{P}_2^+ = \{(\text{Pr}_\beta, 4\beta(1 - \beta)) : 1/3 \leq \beta \leq 2/3\}$. Before going on, it is worth noting here how the simple parametric form of \mathcal{P}_0^+ leads to simple parametric forms for \mathcal{P}_1^+ and \mathcal{P}_2^+ .

An easy calculation shows that $[\mathcal{P}_{0,reg}^+(h), \mathcal{P}_{0,reg}^+(h)] = [1/3, 2/3]$, $[\mathcal{P}_{1,regret}^+(h), \mathcal{P}_{1,reg}^+(h)] = [1/3, 3/8]$, and $[\mathcal{P}_{2,reg}^+(h), \mathcal{P}_{2,reg}^+(h)] = [11/27, 16/27]$. In more detail, since $\text{Pr}_\beta(h) = \beta$ and $\text{Pr}_\beta(t) = 1 - \beta$, so we have the following:

- $\mathcal{P}_{0,reg}^0(h) = \sup_{\beta \in [1/3, 2/3]} (1 - \beta) = 2/3$.
- $\mathcal{P}_{0,reg}^0(h) = \inf_{\beta \in [1/3, 2/3]} \beta = 1/3$.
- $\mathcal{P}_{1,reg}^0(h) = \sup_{\beta \in [1/3, 2/3]} (3\beta/2)(1 - \beta)$. Taking the derivative shows that $(3\beta/2)(1 - \beta)$ is maximized when $\beta = 1/2$, so $\mathcal{P}_{1,reg}^0(h) = 3/8$.

- $\underline{\mathcal{P}}_{1,reg}^0(h) = \inf_{\beta \in [1/3, 2/3]} (1 - (3\beta/2)\beta)$. Now $1 - (3\beta/2)\beta$ is minimized, when $(3\beta/2)\beta$ is maximized; for $\beta \in [1/3, 2/3]$, this happens when $\beta = 2/3$, so $\underline{\mathcal{P}}_{1,reg}^0(h) = 1/3$.
- $\mathcal{P}_{2,reg}^0(h) = \sup_{\beta \in [1/3, 2/3]} 4\beta(1-\beta)(1-\beta)$. Taking the derivative shows that $4\beta(1-\beta)^2$ is maximized when $\beta = 1/3$, in which case it is $16/27$.
- $\underline{\mathcal{P}}_{2,reg}^0(h) = \inf_{\beta \in [1/3, 2/3]} (1 - 4\beta^2(1-\beta)\beta)$. Now $1 - 4\beta^2(1-\beta)$ is minimized when $4\beta^2(1-\beta)$ is maximized; for $\beta \in [1/3, 2/3]$, this happens when $\beta = 2/3$, so $\underline{\mathcal{P}}_{2,reg}^0(h) = 11/27$.

It is also easy to see that $\inf_{\Pr} 4\beta(1-\beta) \Pr_{\beta}(t) = \inf_{\beta \in [1/3, 2/3]} 4\beta(1-\beta)^2 = 8/27$, so

$$\inf_{\Pr \in \mathcal{P}_2} 4\beta(1-\beta) \Pr_{\beta}(t) < 1 - \mathcal{P}_{2,reg}^+(t) < \mathcal{P}_{2,reg}^+(h).$$

Thus, for \mathcal{P}_2^+ , we get strict inequalities for the expressions in Lemma 3.1. ■

The width of the interval $[\underline{\mathcal{P}}_{reg}^+(E), \mathcal{P}_{reg}^+(E)]$ can be viewed as a measure of the ambiguity the agent feels about E , just as the interval $[\mathcal{P}_*(E), \mathcal{P}^*(E)]$. Indeed, if all the weights are 1, the two intervals have the same width, since $\mathcal{P}_*(E) = 1 - \mathcal{P}_{reg}^+(E)$ and $\mathcal{P}^*(E) = 1 - \underline{\mathcal{P}}_{reg}^+(E)$ in this case.

However, weighted regret has a significant advantage over upper and lower probability. If the true bias of the coin is, say $5/8$, then if the set \mathcal{P}_k^+ represents the uncertainty after k steps, as k increases, almost surely, $[\underline{\mathcal{P}}_{k,reg}^+(h), \mathcal{P}_{k,reg}^+(h)]$ will be a smaller and smaller interval containing $1 - 5/8 = 3/8$. More generally, using likelihood updated combined with weighted regret provides a natural way to model the reduction of ambiguity via learning.

It is worth at this point comparing the approach to representing likelihood taken here to the work of Nau (1992). Nau starts with a preference order on *lotteries* (functions from some finite state space S to the reals) satisfying certain axioms, and derives from that what he calls *confidence-weighted (lower and upper) probabilities*. Roughly speaking, rather than just associating with each event its lower and upper probability, Nau can associate with each event E , confidence $c \in [0, 1]$, and probability $p \in [0, 1]$ the set $\mathcal{P}_{c,p}^-$ of probabilities that give event E lower probability p with confidence at least c . If $c' \geq c$, then $\mathcal{P}_{c',p}^- \subseteq \mathcal{P}_{c,p}^-$ (every probability measures that gives E lower probability p with the higher confidence c' will also give it lower probability p with confidence c , but the converse may not hold). Similarly, we can consider the probability measures that give E upper probability p with confidence c . With a set \mathcal{P} of unweighted probabilities, an agent's uncertainty regarding an event E can be characterized by a single interval $[\mathcal{P}_*(E), \mathcal{P}^*(E)]$. In Nau's framework, an agent's uncertainty regarding E can be characterized by a family of intervals $[\mathcal{P}_c(E), \mathcal{P}^c(E)]$, indexed by the confidence c , where $\mathcal{P}_c(E)$ is the largest p such that E has lower probability with confidence c , and $\mathcal{P}^c(E)$ is defined similarly. Clearly these intervals are nested; if $c' > c$, then $[\mathcal{P}_{c'}(E), \mathcal{P}^{c'}(E)]$ contains $[\mathcal{P}_c(E), \mathcal{P}^c(E)]$. Thus, Nau's approach provides a more "fine-grained" representation of uncertainty than the single intervals $[\mathcal{P}_*(E), \mathcal{P}^*(E)]$ or $[\underline{\mathcal{P}}_{reg}^+(E), \mathcal{P}_{reg}^+(E)]$. To some extent, this distinction is due to the fact that Nau's preference order on lotteries is only a partial order; the preference order induced by $\max = \min$ expected utility regret is total. However, note that even though \mathcal{P}_{**} , \mathcal{P}^* , \mathcal{P}_{reg}^+ , and $\underline{\mathcal{P}}_{reg}^+$ all put a total order on events, when considering both \mathcal{P}_* and \mathcal{P}^* or both \mathcal{P}_{reg}^+ and $\underline{\mathcal{P}}_{reg}^+$ together,

we can also obtain a partial order on events; in particular, these approaches can express ambiguity.

One benefit of the regret-based approach is that it provides a natural way of updating. Nau does not consider updating; it would be interesting to see if an analogue of likelihood updating could be defined axiomatically in Nau's framework, perhaps in the spirit of the characterization that Leung and I (Halpern & Leung, 2012) gave for likelihood updating in the context of regret.

One concern with the use of regret has been the dependence of regret on the menu; Nau's approach, and other approaches to decision-making that are not based on regret, do not require a menu. While there is evidence from the psychology literature suggesting that people are quite sensitive to menus, it is also worth noting that when dealing with likelihood, there is a sense in which we can work with the absolute notion of weighted regret without loss of generality: if we restrict to indicator functions, then a preference relative to a menu can always be reduced to an absolute preference. Given a menu M consisting of indicator functions, let $E_M = \cup\{E : 1_E \in M\}$; that is, E_M is the union of the events for which the corresponding indicator function is in M . The following property shows that, when restrict to indicator functions, regret satisfies satisfies an axiom similar in spirit to Nau's (1992) *cancellation axiom*.

Proposition 3.3: *If M is a menu consisting of indicator functions, and $1_{E_1}, 1_{E_2} \in M$, then $1_{E_1} \succeq_{\mathcal{P}^+, M}^{reg} 1_{E_2}$ iff $1_{E_1} + 1_{\bar{E}_M} \succeq_{\mathcal{P}^+}^{reg} 1_{E_2} + 1_{\bar{E}_M}$.*

Proof: Let M' be any menu consisting of indicator functions that includes $1_{E_1} + 1_{\bar{E}_M}$, $1_{E_2} + 1_{\bar{E}_M}$, and 1_S . Recall that $1_{E_1} + 1_{\bar{E}_M} \succeq_{\mathcal{P}^+}^{reg} 1_{E_2} + 1_{\bar{E}_M}$ iff $1_{E_1} + 1_{\bar{E}_M} \succeq_{M', \mathcal{P}^+}^{reg} 1_{E_2} + 1_{\bar{E}_M}$; the absolute notion of regret is equivalent to the menu-based notion, as long as the menu includes the best act, which in this case is 1_S . It clearly suffices to show that, for all states $s \in S$ and all acts $1_E \in M$,

$$reg^M(1_E, s) = reg^{M'}(1_E + 1_{\bar{E}_M}, s).$$

This is straightforward. There are two cases, depending on whether $s \in E_M$.

If $s \in E_M$, then, by definition, there is some act $1_{E'} \in M$ such that $s \in E'$, so $\sup_{a \in M} u(a(s)) = u(1)$. Clearly $\sup_{a \in M'} u(a(s)) = u(1)$, since $1_S \in M'$. Moreover, $1_{\bar{E}_M}(s) = 0$, so $(1_E + 1_{\bar{E}_M})(s) = 1_E(s)$. Thus, for $s \in E_M$,

$$\begin{aligned} reg^M(1_E, s) &= \sup_{a \in M} u(a(s)) - u(1_E(s)) \\ &= \sup_{a \in M'} u(a(s)) - u((1_E + 1_{\bar{E}_M})(s)) \\ &= reg^{M'}(1_E + 1_{\bar{E}_M}, s). \end{aligned}$$

For $s \notin \bar{E}_M$, we have $a(s) = 0$ for all $a \in M$ and $1_E(s) = 0$, so $\sup_{a \in M} u(a(s)) - u(1_E(s)) = 0$. On the other hand, $\sup_{a \in M'} u(a(s)) = u(1)$, and $u((1_E + 1_{\bar{E}_M})(s)) = u(1)$, so again $\sup_{a \in M'} u(a(s)) - u((1_E + 1_{\bar{E}_M})(s)) = 0$. Thus, we again have $reg^M(1_E, s) = reg^{M'}(1_E + 1_{\bar{E}_M}, s)$. ■

4. Characterizing Weighted Regret-Based Likelihood

The goal of this section is to characterize weighted regret-based likelihood axiomatically. In order to do so, it is helpful to review the characterizations of probability and lower probability. For ease of exposition in this discussion, I assume that the sample space is finite and all sets are measurable.

A probability measure on a finite set S maps subsets of S to $[0, 1]$ in a way that satisfies the following three properties:

Pr1. $\Pr(S) = 1$.

Pr2. $\Pr(\emptyset) = 0$.⁵

Pr3. $\Pr(E \cup E') = \Pr(E) + \Pr(E')$ if $E \cap E' = \emptyset$.

These three properties characterize probability in the sense that any function $f : 2^S \rightarrow [0, 1]$ that satisfies these properties is a probability measure.

Lower probabilities satisfy analogues of these properties:

LP1. $\mathcal{P}_*(S) = 1$.

LP2. $\mathcal{P}_*(\emptyset) = 0$.

LP3'. $\mathcal{P}_*(E \cup E') \geq \mathcal{P}_*(E) + \mathcal{P}_*(E')$ if $E \cap E' = \emptyset$.

However, these properties do not characterize lower probability. There are functions that satisfy LP1, LP2, and LP3' that are not the lower probability corresponding to some set of probability measures. (See (Halpern & Pucella, 2002, Proposition 2.2) for an example showing that analogous properties do not characterize \mathcal{P}^* ; the same example also shows that they do not characterize \mathcal{P}_* .)

Various characterizations of \mathcal{P}_* (and \mathcal{P}^*) have been proposed in the literature (Anger & Lembcke, 1985; Giles, 1982; Huber, 1976, 1981; Lorentz, 1952; Williams, 1976; Wolf, 1977), all similar in spirit. I discuss one due to Anger and Lembcke (1985) here, since it makes the contrast between lower probability and regret particularly clear. The characterization is based on the notion of *set cover*: a set E is said to be covered n times by a multiset M if every element of E appears at least n times in M . It is important to note here that M is a multiset, not a set; its elements are not necessarily distinct. (Of course, a set is a special case of a multiset.) Let \sqcup denote multiset union; thus, if M_1 and M_2 are multisets, then $M_1 \sqcup M_2$ consists of all the elements in M_1 or M_2 , which appear with multiplicity that is the sum of the multiplicities in M_1 and M_2 . For example, using the $\{\{\dots\}\}$ notation to denote a multiset, then $\{\{1, 1, 2\}\} \sqcup \{\{1, 2, 3\}\} = \{\{1, 1, 1, 2, 2, 3\}\}$.

If $E \subseteq S$, then an (n, k) -cover of (E, S) is a multiset M that covers S k times and covers E $n + k$ times. Multiset M is an n -cover of E if M covers E n times. For example, if $S = \{1, 2, 3\}$, then $\{\{1, 1, 1, 2, 2, 3\}\}$ is a $(2, 1)$ -cover of $(\{1\}, S)$, a $(1, 1)$ -cover of $(\{1, 2\}, S)$, and a 3-cover of $\{1\}$.

We will be interested in whether a multiset of the form $\overline{E}_1 \sqcup \dots \sqcup \overline{E}_m$ is an (n, k) -cover of (\overline{E}, S) . This is perhaps best thought of in terms of indicator functions. $\overline{E}_1 \sqcup \dots \sqcup \overline{E}_m$

5. This property actually follows from the other two, using the observation that $\Pr(S \cup \emptyset) = \Pr(S) + \Pr(\emptyset)$; I include it here to ease the comparison to other approaches.

is an (n, k) -cover of (\bar{E}, S) if and only if $1_{E_1} + \dots + 1_{E_m} \geq n1_E + k1_S$. The use of equalities and inequalities involving sums of indicator functions in axiomatic characterizations of uncertainty has a long history; for example, they were used by Scott (1964) to characterize qualitative probability. Set covers are just a special case of such inequalities. Typically, such axioms make it possible to apply results from linear programming to prove characterization results. As we shall see, that will be the case here too.

Consider the following property:

LP3. For all integers m, n, k and all subsets E_1, \dots, E_m of S , if $E_1 \sqcup \dots \sqcup E_m$ is an (n, k) -cover of (E, S) , then $k + n\mathcal{P}_*(E) \geq \sum_{i=1}^m \mathcal{P}_*(E_i)$.⁶

There is an analogous property for upper probability, where \geq is replaced by \leq . It is easy to see that LP3 implies LP3' (since $E \sqcup E'$ is a $(1, 0)$ cover of $(E \cup E', S)$). It follows by a straightforward induction from LP3' that if E_1, \dots, E_m are pairwise disjoint, then $\mathcal{P}_*(E_1 \cup \dots \cup E_m) \geq \mathcal{P}_*(E_1) + \dots + \mathcal{P}_*(E_m)$. LP3 generalizes this property to allow for sets that are not necessarily disjoint. The soundness of LP3 for lower probability follows using the same techniques as given below for the soundness of the property REG3. As Anger and Lembcke (1985) show, LP3 is just the property that is needed to characterize lower probability.

Theorem 4.1: (Anger & Lembcke, 1985) *If $f : 2^S \rightarrow [0, 1]$, then there exists a set \mathcal{P} of probability measures with $f = \mathcal{P}_*$ if and only if f satisfies LP1, LP2, and LP3.*

Moving to regret-based likelihood, clearly we have

REG1. $\mathcal{P}_{reg}^+(S) = 0$.

REG2. $\mathcal{P}_{reg}^+(\emptyset) = 1$.

The whole space S has the least regret; the empty set has the greatest regret. Again, we see that regret-based likelihood inverts the standard ordering of probability; larger regret-based likelihood corresponds to probability.

In the unweighted case, since $\mathcal{P}_{reg}(E) = \mathcal{P}^*(\bar{E})$, REG1, REG2, and the following analogue of LP3 (appropriately modified for \mathcal{P}^*) clearly characterize \mathcal{P}_{reg} :

REG3'. For all integers m, n, k and all subsets E_1, \dots, E_m of S , if $\bar{E}_1 \sqcup \dots \sqcup \bar{E}_m$ is an (n, k) -cover of (\bar{E}, S) , then $k + n\mathcal{P}_{reg}(E) \leq \sum_{i=1}^m \mathcal{P}_{reg}(E_i)$.

Note that complements of sets $(\bar{E}_1, \dots, \bar{E}_m, \bar{E})$ are used here, since regret is minimized if the probability of the complement is maximized. This need to work with the complement makes the statement of the properties (and the proofs of the theorems) slightly less elegant, but seems necessary.

It is not hard to see that REG3' does not hold for weighted regret-based likelihood. For example, suppose that $S = \{a, b, c\}$ and $\mathcal{P}^+ = ((Pr_1, 2/3), (Pr_2, 2/3), (Pr_3, 1))$, where, identifying the probability Pr with the tuple $(Pr(a), Pr(b), Pr(c))$, we have

- $Pr_1 = (2/3, 0, 1/3)$;

6. Note that LP3 implies LP2, using the fact that $\emptyset \sqcup \emptyset$ is a $(1, 0)$ -cover of (\emptyset, S) .

- $\text{Pr}_2 = (1/3, 0, 2/3)$;
- $\text{Pr}_3 = (1/3, 1/3, 1/3)$.

Then $\mathcal{P}_{reg}^+(\{a, b\}) = \mathcal{P}_{reg}^+(\{b, c\}) = 4/9$, while $\mathcal{P}_{reg}^+(\{b\}) = 2/3$. Since $\{a, b\} \sqcup \{b, c\}$ is a (1,1)-cover of $(\{b\}, \{a, b, c\})$, REG3' would require that

$$\mathcal{P}_{reg}^+(\{a, b\}) + \mathcal{P}_{reg}^+(\{b, c\}) \geq 1 + \mathcal{P}_{reg}^+(\{b\}),$$

which is clearly not the case.

We must thus weaken REG3' to capture weighted regret-based likelihood. It turns out that the appropriate weakening is the following:

REG3. For all integers m, n and all subsets E_1, \dots, E_m of S , if $\bar{E}_1 \sqcup \dots \sqcup \bar{E}_m$ is an n -cover of \bar{E} , then $n\mathcal{P}_{reg}^+(E) \leq \sum_{i=1}^m \mathcal{P}_{reg}^+(E_i)$.

Although REG3 is weaker than REG3', it still has some nontrivial consequences. For example, it follows from REG3 that \mathcal{P}_{reg}^+ is anti-monotonic. If $E \subseteq E'$, then \bar{E} is a 1-cover of \bar{E}' , so by REG3, we must have $\mathcal{P}_{reg}^+(E) \geq \mathcal{P}_{reg}^+(E')$. Since $E \sqcup E'$ is trivially a 1-cover of $E \cup E'$, it also follows that $\mathcal{P}_{reg}^+(\bar{E}) + \mathcal{P}_{reg}^+(\bar{E}') \geq \mathcal{P}_{reg}^+(\overline{E \cup E'})$. REG3 also implies REG1, since $\emptyset (= \bar{S})$ is an n -cover of itself for all n .

I can now state the representation theorem. It says that a representation of uncertainty satisfies REG1, REG2, and REG3 iff it is the weighted regret-based likelihood determined by some set \mathcal{P}^+ . The set \mathcal{P}^+ is not unique, but it can be taken to be *maximal*, in the sense that if weighted regret-based likelihood with respect to some other set $(\mathcal{P}')^+$ gives the same representation, then for all pairs $(\text{Pr}, \alpha') \in (\mathcal{P}')^+$, there exists $\alpha \geq \alpha'$ such that $(\text{Pr}, \alpha) \in \mathcal{P}^+$. This (unique) maximal set \mathcal{P}^+ can be viewed as the canonical representation of uncertainty.

Theorem 4.2: *If $f : 2^S \rightarrow [0, 1]$, then there exists a weakly closed set \mathcal{P}^+ of weighted probability measures with $f = \mathcal{P}_{reg}^+$ if and only if f satisfies REG1, REG2, and REG3; moreover, \mathcal{P}^+ can be taken to be maximal.*

Proof: Clearly, given a weakly closed set \mathcal{P}^+ of weighted probability measures, the function \mathcal{P}_{reg}^+ satisfies REG1 and REG2. To see that it satisfies REG3, suppose that $\bar{E}_1 \sqcup \dots \sqcup \bar{E}_m$ is an n -cover of \bar{E} . If $\mathcal{P}_{reg}^+(E) = 0$, then REG3 trivially holds. If $\mathcal{P}_{reg}^+(E) > 0$, then since \mathcal{P}^+ is weakly closed, there must be some probability $\text{Pr} \in \mathcal{P}$ such that $\mathcal{P}_{reg}^+(E) = \alpha_{\text{Pr}} \text{Pr}(\bar{E})$. Since $\bar{E}_1 \sqcup \dots \sqcup \bar{E}_m$ is an n -cover of \bar{E} , it is easy to see that $\text{Pr}(\bar{E}_1) + \dots + \text{Pr}(\bar{E}_m) = n \text{Pr}(\bar{E})$, so $\alpha_{\text{Pr}} \text{Pr}(\bar{E}_1) + \dots + \alpha_{\text{Pr}} \text{Pr}(\bar{E}_m) = n\alpha_{\text{Pr}} \text{Pr}(\bar{E})$. But $\alpha_{\text{Pr}} \text{Pr}(\bar{E}) = \mathcal{P}_{reg}^+(E)$, by construction, and $\alpha_{\text{Pr}} \text{Pr}(\bar{E}_i) \leq \mathcal{P}_{reg}^+(E_i)$, $i = 1, \dots, m$. Thus, $n\mathcal{P}_{reg}^+(E) \leq \sum_{i=1}^m \mathcal{P}_{reg}^+(E_i)$.

For the opposite direction, suppose that $f : 2^S \rightarrow [0, 1]$ satisfies REG1, REG2, and REG3. Let $\mathcal{P} = \Delta(S)$, the set of all probability measures on S , and for $\text{Pr} \in \mathcal{P}$, define

$$\alpha_{\text{Pr}} = \sup\{\beta : \beta \text{Pr}(\bar{E}) \leq f(E) \text{ for all } E \subseteq S\}.$$

Note that, for all $\text{Pr} \in \mathcal{P}$, we have $0 \text{Pr}(\bar{E}) \leq f(E)$ for all $E \subseteq S$, since $f(E) \in [0, 1]$, and $1 \text{Pr}(\bar{\emptyset}) = f(\emptyset) = 1$. It follows that $\alpha_{\text{Pr}} \in [0, 1]$ for all $\text{Pr} \in \mathcal{P}$. Let $\mathcal{P}^+ = \{(\text{Pr}, \alpha_{\text{Pr}}) :$

$\Pr \in \Delta(S)$. It is easy to see that \mathcal{P}^+ is weakly closed. Moreover, if we can show that \mathcal{P}^+ represents f (i.e., $f = \mathcal{P}_{reg}^+$), it is immediate that \mathcal{P}^+ is maximal among all sets of weighted probability measures that represent f . Thus, it suffices to show that there exists $\Pr \in \Delta(S)$ such that (1) $\alpha_{\Pr} = 1$ (since this is one of the conditions on sets of weighted measures) and (2) $f(E) = \mathcal{P}_{reg}^+(E)$ for all $E \subseteq S$.

The proof of this result makes critical use of the following variant of Farkas' Lemma (Farkas, 1902) (see also Schrijver, 1986, pg. 89) from linear programming, where A is a matrix, b is a column vector, and x is a column vector of distinct variables:

Lemma 4.3: *If $Ax \geq b$ is unsatisfiable, then there exists a row vector β such that*

1. $\beta \geq 0$
2. $\beta A = 0$
3. $\beta b > 0$.

Intuitively, β is a “witness” of the fact that $Ax \geq b$ is unsatisfiable. This is because if there were a vector x satisfying $Ax \geq b$, then $0 = (\beta A)x = \beta(Ax) \geq \beta b > 0$, a contradiction.

To prove the first claim, suppose that $S = \{s_1, \dots, s_N\}$. I now construct a set of linear equations in the variables x_1, \dots, x_N such that a solution to the equations guarantees the existence of a probability measure $\Pr \in \Delta(S)$ such that $\alpha_{\Pr} = 1$. Intuitively, we want x_i to be $\Pr(s_i)$. Since we must have $\Pr(\overline{E}) \leq f(E)$ for all $E \subset S$,⁷ for each $E \subset S$, we have the inequality $\sum_{\{i:s_i \notin E\}} x_i \leq f(E)$. Note that since $f(\emptyset) = 1$, the equation when $E = \emptyset$ is $x_1 + \dots + x_N \leq 1$. In addition, we require that $x_i \geq 0$ for $i = 1, \dots, N$, and that $x_1 + \dots + x_N = 1$. It suffices to require that $x_1 + \dots + x_n \geq 1$, since, as I observed earlier, the equation corresponding to $E = \emptyset$ already says $x_1 + \dots + x_n \leq 1$. To apply Farkas' Lemma all the inequalities need to involve \geq , so this collection of inequalities must be rewritten as:

$$\begin{aligned} -\sum_{\{i:s_i \notin E\}} x_i &\geq -f(E), \text{ for all } E \subset S \\ x_i &\geq 0, \text{ for } i = 1, \dots, N \\ x_1 + \dots + x_N &\geq 1. \end{aligned}$$

This system of inequalities can be expressed in the form $Ax \geq b$. Note that A is a matrix all of whose entries are either -1 , 0 , or 1 , and, in the first $2^N - 1$ rows (the lines corresponding to equations for each $E \subset S$), all the entries are either 0 or -1 , while in the final $N + 1$ rows, all the entries are either 0 or 1 .

A solution of this system of inequalities provides the desired \Pr . But if this systems has no solution, then by Farkas' Lemma, there exists a nonnegative vector β such that $\beta A = 0$ and $\beta b > 0$. Since all the entries of A are either -1 , 0 , or 1 , it follows from standard observations (cf., Fagin, Halpern, & Megiddo, 1990, Lemma 2.7) we can take β to a vector of all whose entries are rational.⁸ Since we can multiply each term in β by the product

7. I use \subset to denote strict subset.

8. There is a slight subtlety here since β also has to satisfy $\beta b > 0$, and b may involve irrational numbers (since $f(E)$ may be irrational for some sets E). However, if there is a nonnegative β that satisfies $\beta A = 0$ and $\beta b > 0$, then there is a nonnegative β that satisfies $\beta A = 0$ and $\beta b' > 0$, where b' consists only of rational entries and $b' \leq b$. Thus, there is a vector β with rational entries such that $\beta A = 0$ and $\beta b' > 0$, so $\beta b > 0$.

of the denominators of the entries of β , we can assume without loss of generality that the entries of β are natural numbers.

Since A has $2^N + N$ rows, β is a vector of the form $(\beta_1, \dots, \beta_{2^N+N})$. Let A_1, \dots, A_{2^N+N} be the rows of A ; each of these is a vector of length N . Since $\beta A = 0$, that means that $\beta_1 A_1 + \dots + \beta_{2^N+N} A_{2^N+N} = 0$. Suppose for now that $\beta_{2^N}, \dots, \beta_{2^N+N-1}$ (the coefficients for the rows corresponding to the inequalities $x_i \geq 0$ for $i = 1, \dots, N$) are all 0; as I show below, this assumption can be made without loss of generality.

With this assumption, we can rewrite the equations as $\beta_1 A_1 + \dots + \beta_{2^N-1} A_{2^N-1} = -\beta_{2^N+N} A_{2^N+N}$. If E_1, \dots, E_{2^N-1} are the subsets of S that correspond to the equations for A_1, \dots, A_{2^N-1} , respectively, this equation says that β_1 copies of \overline{E}_1 , β_2 copies of \overline{E}_2 , \dots , and β_{2^N-1} copies of \overline{E}_{2^N-1} form a β_{2^N+N} -cover of S . (Recall that A_{2^N+N} is a row of all -1 's, so $-A_{2^N+N}$ corresponds to S .) Thus, by REG3, $\beta_1 f(E_1) + \dots + \beta_{2^N-1} f(E_{2^N-1}) \geq \beta_{2^N+N} f(\emptyset) = \beta_{2^N+N}$. But Farkas' Lemma requires that $\beta b > 0$, where, by construction, $b_i = -f(E_i)$ for $i = 1, \dots, 2^N - 1$, $b_i = 0$ for $i = 2^N, \dots, 2^N + N - 1$, and $b_{2^N+N} = 1$. Thus, we must have $-(\beta_1 f(E_1) + \dots + \beta_{2^N-1} f(E_{2^N-1})) > -\beta_{2^N+N}$. Clearly, this gives a contradiction. Thus, we can conclude, as desired, that the equations are solvable, and that there exists a probability measure \Pr such that $\alpha_{\Pr} = 1$.

It remains to show that we can assume without loss of generality that $\beta^{2^N}, \dots, \beta^{2^N+N-1}$ are all 0. Note that since $\beta \geq 0$, they must all be nonnegative. I prove by induction on $\beta_{2^N} + \dots + \beta_{2^N+N-1}$ that if there is a vector $\beta \geq 0$ such that $\beta A = 0$ and $\beta b > 0$, then there is such a vector with $\beta_{2^N} + \dots + \beta_{2^N+N-1} = 0$.

So suppose that there is a solution β with $\beta_{2^N} + \dots + \beta_{2^N+N-1} > 0$. Suppose without loss of generality that $\beta_{2^N} > 0$. Recall that A_{2^N} corresponds to the inequality $x_1 \geq 0$. Choose $j \in \{0, \dots, 2^N - 1\}$ such that $\beta_j > 0$ and $s_1 \notin E_j$. There must be such a j , for otherwise we would not have $\beta A = 0$. Let j' be such that $E_{j'} = E_j \cup \{s_1\}$. Define a vector β' such that $\beta'_{2^N} = \beta_{2^N} - 1$, $\beta'_j = \beta_j - 1$, $\beta'_{j'} = \beta_j + 1$, and $\beta'_i = \beta_i$ if $i \notin \{j, j', 2^N\}$. It is easy to check that $\beta' A = 0$ and that $\beta'_{2^N} + \dots + \beta'_{2^N+N-1} < \beta_{2^N} + \dots + \beta_{2^N+N-1}$. It remains to show that $\beta' b > 0$. Since $E_j \subset E_{j'}$, we must have $f(E_j) \geq f(E_{j'})$, so $\beta' b = \beta b + f(E_j) - f(E_{j'}) \geq \beta b > 0$. This completes the inductive step of the argument.

Now we must show the second required property holds, namely, that $f(E) = \mathcal{P}_{reg}^+(E)$ for all $E \subseteq S$. By construction, $\alpha_{\Pr} \Pr(\overline{E}) \leq f(E)$ for all $E \subseteq S$, so it suffices to show that there is some $\Pr \in \mathcal{P}$ such that $\alpha_{\Pr} \Pr(\overline{E}) = f(E)$. For this, it suffices to show that there exists a measure \Pr such that $\Pr(\overline{E}) = 1$, and for each $E' \subset S$, we have $f(E) \Pr(\overline{E}') \leq f(E')$, since then $\alpha_{\Pr} = f(E)$, so $\alpha_{\Pr} \Pr(\overline{E}) = f(E)$, as desired.

To show that such a measure exists, we again construct a set of linear inequalities much as above, and apply Farkas' Lemma. Using the same notation as above, suppose for simplicity that $\overline{E} = \{s_1, \dots, s_M\}$, where $M \leq N$. Now the required inequalities just involve the variables x_1, \dots, x_M :

$$\begin{aligned} & -\sum_{\{i: s_i \in \overline{E} \cap \overline{E}'\}} x_i \geq -f(E')/f(E), \text{ for all } E' \subset S \text{ such that } \overline{E} \cap \overline{E}' \neq \emptyset \\ & x_i \geq 0, \text{ for } i = 1, \dots, M \\ & x_1 + \dots + x_M \geq 1. \end{aligned}$$

Again, the requirement that $x_1 + \dots + x_M \leq 1$ follows from the equation for E .

If this system of inequalities is satisfiable, then we have the required probability measure, so suppose that it is not satisfiable. Again, writing this system of equations as $Ax \geq b$, by

Farkas' Lemma, there exists a nonnegative vector β such that $\beta A = 0$ and $\beta b > 0$. We now proceed much as before. Again, we can assume that β is a vector of natural numbers. If we assume for now that $\beta^{2^M}, \dots, \beta^{2^M+M-1}$ (the coefficients for the rows corresponding to the inequalities $x_i \geq 0$ for $i = 1, \dots, N$) are all 0, then the fact that $\beta A = 0$ means that we have β_{2^M+M} cover of E . We get a contradiction to REG3 in an almost identical way to above. This completes the argument. ■

As I said earlier, the set \mathcal{P}^+ guaranteed to exist by Theorem 4.2 is not unique, although it is canonical, in the sense of being the unique maximal set of weighted probability measures that represents f . We might wonder if we can actually get uniqueness by imposing a few extra requirements, particularly since Leung and I were able to do so in our representation theorem. The answer seems to be no. To explain why, it is helpful to review some material from (Halpern & Leung, 2012).

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 $\mathbf{p}(S) = 1$. We can identify a weighted probability distribution (Pr, α) with the sub-probability measure αPr . Conversely, given a sub-probability measure \mathbf{p} , there is a unique pair (α, Pr) such that $\mathbf{p} = \alpha \text{Pr}$: we simply take $\alpha = \mathbf{p}(S)$ and $\text{Pr} = \mathbf{p}/\alpha$. Thus, in the sequel, I identify a set of sub-probability measures with a set of weighted probability measures.

A set B of sub-probability measures is *downward-closed* if, whenever $\mathbf{p} \in B$ and $\mathbf{q} \leq \mathbf{p}$, then $\mathbf{q} \in B$.

One advantage of considering sub-probability measures is that while it is not clear what it would mean for a set of weighted probabilities to be convex (indeed, it is not obvious what should count as a convex combination of (Pr, α) and (Pr', α')), it is quite clear what counts as a convex combination of sub-probability measures. Moreover, a convex combination of sub-probability measures is itself a sub-probability measure.

Call a set of subprobability measures *regular* if it is convex, downward-closed, closed, and contains at least one proper probability measure. (The latter requirement corresponds to having $\alpha_{\text{Pr}} = 1$ for some $\text{Pr} \in \mathcal{P}^+$.) Leung and I provide a set of axioms for preference orders, and show that a family of preference orders \succeq_M indexed by menus satisfies these axioms iff there is a unique regular set of weighted probability measures \mathcal{P}^+ such that, for all $a \succeq_M b$ iff $wr_{\mathcal{P}^+}^M(a) \leq wr_{\mathcal{P}^+}^M(b)$. Thus, we might hope that we can get uniqueness by imposing a regularity requirement. It is easy to see that the canonical maximal set \mathcal{P}^+ constructed in the proof of Theorem 4.2 is regular, which lends some credence to this hope. Unfortunately, as the following example shows, regularity does not suffice for uniqueness.

Example 4.4: Let $S = \{s_1, s_2\}$, and let f be defined on 2^S by taking $f(\{s_1\}) = 1/4$ and $f(\{s_2\}) = 1$ (and $f(S) = 0$ and $f(\emptyset) = 1$). A sub-probability measure \mathbf{p} on S can be identified with the pair $(\mathbf{p}(s_1), \mathbf{p}(s_2))$, which makes it easy to think about sub-probability measures on S geometrically. A set of sub-probability measures is just a region in \mathbb{R}^2 contained in the triangle bounded by the lines $x = 0$, $y = 0$, and $y = 1 - x$. A set \mathcal{P}^+ of subprobability measures is downward closed if, whenever it contains a point (x, y) , it contains all (x', y') in the rectangle defined by the points $(0, 0)$, $(x, 0)$, $(0, y)$, and (x, y) .

With this intuition, let \mathcal{P}_0^+ be the set of subprobabilities in the quadrilateral bounded by $x = 0$, $y = 0$, $y = 1 - x$, and $y = 1/4$ (the region marked by vertical lines in Figure 1). It

is not hard to show that \mathcal{P}_0^+ is the maximal set of weighted probabilities representing f . It is clearly regular. Since it contains the subprobability $(1, 0)$, it follows that $\mathcal{P}_{0,reg}^+(\{s_2\}) = 1$. It is also easy to see that, since $(0, 1/4) \in \mathcal{P}_0^+$ and $\mathbf{p}(s_2) \leq 1/4$ for all $\mathbf{p} \in \mathcal{P}_0^+$, we have that $\mathcal{P}_{0,reg}^+(\{s_1\}) = 1/4$.

But now let \mathcal{P}_1^+ consist of all sub-probabilities in the triangle bounded by $x = 0$, $y = 0$, and $y = \frac{1-x}{4}$ (the region marked by horizontal lines in Figure 1). Clearly \mathcal{P}_1^+ is a strict subset of \mathcal{P}_0^+ , but it is clear from the figure that it is also regular. Moreover, since it contains the points $(\frac{1}{4}, 0)$ and $(0, 1)$, it also represents f . Indeed, it easily follows from the geometry of the situation that there are uncountably many regular sets of weighted probabilities representing f ; for all $z \in [0, \frac{3}{4}]$, the regular set bounded by the lines $x = 0$, $y = 0$, $y = \frac{1}{4}$, and the line from $(z, \frac{1}{4})$ to $(1, 0)$.




Figure 1: Regular sets of weighted probability measures that represent f .

Intuitively, the problem here is that a function on S does not contain enough information to uniquely determine a regular set of weighted probability measures. It is not clear whether there are natural further conditions that can be imposed that we lead to uniqueness. It seems that the closest that we can come to uniqueness is to consider the maximal set. ■

5. Conclusion

I have defined an approach for associating with an event E a numerical representation of its likelihood when uncertainty is represented by a set of weighted probability measures. The representation consists of a pair of numbers, which can be thought of as upper and lower bounds on the uncertainty. The difference between these numbers can be viewed as a measure of ambiguity. The two numbers coincide when uncertainty is represented by a single probability. Moreover, if each probability measure gets weight 1, then the two numbers can essentially be viewed as the lower and upper probabilities of E (more precisely, $1 - \mathcal{P}_*(E)$ and $1 - \mathcal{P}^*(E)$). Thus, the approach can be viewed as a generalization of lower and upper probability to the case of weighted probability measures, with regret-based likelihood corresponding to upper probability. The definitions show that there is

an interesting connection between regret-based approaches and minimization/maximization approaches when it comes to defining likelihood; this connection breaks down when it comes to more general utility calculations (Halpern & Leung, 2012).

The main technical result of the paper is a complete characterization of the likelihood in the case where the state space is finite. The notion of likelihood can easily be extended to the case of an infinite state space (of course, an integral has to be used instead of a sum to calculate expected utility). I conjecture that the characterization theorem will still hold with essentially no change, although I have not checked details carefully.

Of course, it would be useful to get a better understanding of this numerical representation, to see if it really captures an agent's feelings about both the ambiguity and the risk associated with an event, and to understand its technical properties. I leave this to future work.

Acknowledgments

I thank Samantha Leung, the reviewers of ECSQARU, and the JAIR referees for many useful comments on the paper. The work was 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.

References

- Anger, B., & Lembecke, J. (1985). Infinitely subadditive capacities as upper envelopes of measures. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 68, 403–414.
- Boole, G. (1854). *An Investigation into the Laws of Thought on Which Are Founded the Mathematical Theories of Logic and Probabilities*. Macmillan, London.
- Campos, L. M. d., & Moral, S. (1995). Independence concepts for sets of probabilities. In *Proc. Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 108–115.
- Cattaneo, M. E. G. V. (2007). *Statistical decisions based directly on the likelihood function*. Ph.D. thesis, ETH.
- Chateauneuf, A., & Faro, J. (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics*, 45, 535 – 558.
- Couso, I., Moral, S., & Walley, P. (1999). Examples of independence for imprecise probabilities. In *Proc. First International Symposium on Imprecise Probabilities and Their Applications (ISIPTA '99)*.
- de Cooman, G. (2005). A behavioral model for vague probability assessments. *Fuzzy Sets and Systems*, 154(3), 305–358.
- Dubois, D., & Prade, H. (1998). Possibility measures: qualitative and quantitative aspects. In Gabbay, D. M., & Smets, P. (Eds.), *Quantified Representation of Uncertainty and*

Imprecision, Vol. 1 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, pp. 169–226. Kluwer, Dordrecht, Netherlands.

- Epstein, L., & Schneider, M. (2007). Learning under ambiguity. *Review of Economic Studies*, 74(4), 1275–1303.
- Fagin, R., Halpern, J. Y., & Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation*, 87(1/2), 78–128.
- Farkas, J. (1902). Theorie der einfachen ungleichungen. *J. Reine und Angewandte Math.*, 124, 1–27.
- Gärdenfors, P., & Sahlin, N. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, 53, 361–386.
- Gärdenfors, P., & Sahlin, N. (1983). Decision making with unreliable probabilities. *British Journal of Mathematical and Statistical Psychology*, 36, 240–251.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18, 141–153.
- Gilboa, I., & Schmeidler, D. (1993). Updating ambiguous beliefs. *Journal of Economic Theory*, 59, 33–49.
- Giles, R. (1982). Foundations for a theory of possibility. In Gupta, M. M., & Sanchez, E. (Eds.), *Fuzzy Information and Decision Processes*, pp. 183–195. North-Holland.
- Good, I. J. (1980). Some history of the hierarchical Bayesian methodology. In Bernardo, J. M., DeGroot, M. H., Lindley, D., & Smith, A. (Eds.), *Bayesian Statistic I*, pp. 489–504. University Press: Valencia.
- Halpern, J. Y. (1997). Defining relative likelihood in partially-ordered preferential structures. *Journal of A.I. Research*, 7, 1–24.
- Halpern, J. Y. (2003). *Reasoning About Uncertainty*. MIT Press, Cambridge, Mass.
- Halpern, J. Y., & Leung, S. (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. To appear, *Theory and Decision*.
- Halpern, J. Y., & Pucella, R. (2002). A logic for reasoning about upper probabilities. *Journal of A.I. Research*, 17, 57–81.
- Huber, P. J. (1976). Kapazitäten statt Wahrscheinlichkeiten? Gedanken zur Grundlegung der Statistik. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 78, 81–92.
- Huber, P. J. (1981). *Robust Statistics*. Wiley, New York.
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849–1892.

- Kyburg, Jr., H. E. (1988). Higher order probabilities and intervals. *International Journal of Approximate Reasoning*, 2, 195–209.
- Levi, I. (1985). Imprecision and uncertainty in probability judgment. *Philosophy of Science*, 52, 390–406.
- Lorentz, G. G. (1952). Multiply subadditive functions. *Canadian Journal of Mathematics*, 4(4), 455–462.
- Maccheroni, F., Marinacci, M., & Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6), 1447–1498.
- 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.
- Ostrogradsky, M. V. (1838). Extrait d'un mémoire sur la probabilité des erreurs des tribunaux. *Memoires d'Académie St. Petersbourg, Séries 6, 3*, xix–xxv.
- Pearl, J. (1987). Do we need higher-order probabilities and, if so, what do they mean?. In *Proc. Third Workshop on Uncertainty in Artificial Intelligence (UAI '87)*, pp. 47–60.
- Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46, 55–67.
- Schrijver, A. (1986). *Theory of Linear and Integer Programming*. Wiley, New York.
- Scott, D. (1964). Measurement structures and linear inequalities. *Journal of Mathematical Psychology*, 1, 233–247.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*, Vol. 42 of *Monographs on Statistics and Applied Probability*. Chapman and Hall, London.
- Walley, P. (1997). Statistical inferences based on a second-order possibility distribution. *International Journal of General Systems*, 26(4), 337–383.
- Williams, P. M. (1976). Indeterminate probabilities. In Przelecki, M., Szaniawski, K., & Wojcicki, R. (Eds.), *Formal Methods in the Methodology of Empirical Sciences*, pp. 229–246. Reidel, Dordrecht, Netherlands.
- Wolf, G. (1977). *Obere und untere Wahrscheinlichkeiten*. Ph.D. thesis, ETH, Zurich.
- Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1, 3–28.