DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

PROPER SCORING RULES FOR GENERAL DECISION MODELS

Christopher P. Chambers

SOCIAL SCIENCE WORKING PAPER 1231

August 2005

Proper Scoring Rules for General Decision Models

Christopher P. Chambers

Abstract

On the domain of Choquet expected utility preferences with risk neutral lottery evaluation and totally monotone capacities, we demonstrate that proper scoring rules do not exist. This implies the non-existence of proper scoring rules for any larger class of preferences (CEU with convex capacities, multiple priors). We also show that if an agent whose behavior conforms to the multiple priors model is faced with a scoring rule for a subjective expected utility agent, she will always announce a probability belonging to her set of priors; moreover, for any prior in the set, there exists such a scoring rule inducing the agent to announce that prior.

JEL classification numbers: D81,C49

Key words: scoring rule, proper scoring rule, subjective expected utility, Choquet expected utility, multiple priors, implementation, probability elicitation

Proper Scoring Rules for General Decision Models *

Christopher P. Chambers

1 Introduction

The primary purpose of this note is to determine whether a simple mechanism can be designed whereby an agent facing subjective uncertainty is asked to reveal her "beliefs" about the uncertainty, and it is strictly incentive compatible for her to do so. Our primary finding is that if the possible "beliefs" that the agent may possess belong to a certain class, the answer is no.

Consider an environment with an exogenous set of states of the world. In classical models of subjective expected utility, agents make choices between *acts*, or state-contingent outcomes. Savage [16] establishes that an agent whose behavior conforms with several intuitive axioms acts as if she behaves in an expected utility fashion. For such a decision maker, there exists a unique probability measure over the states of the world, and a utility function over money. The decision maker always makes choices over acts in order to maximize her expected utility. In theory, this unique probability measure can be recovered by observing all possible choices between pairs of acts. While Savage's theory requires an infinite set of states of the world, other theories, most notably that of Anscombe and Aumann [2], do not.

Suppose that the agent in question is risk-neutral, so that her utility index over monetary payoffs is affine (we will see that this is without loss of generality). For such an agent, it is not necessary to observe all possible choices between all pairs of acts in order to elicit the unique probability measure representing beliefs. In fact, it is enough to offer such a decision maker *one* choice over a *menu* of acts, an insight originally due to Brier [4]. Optimizing behavior of the decision maker reveals her probability measure. Such a menu of acts is referred to as a *scoring rule*. A scoring rule is *proper* if the unique optimizing choice is to reveal her probability measure. The theory of scoring rules can easily be viewed as a subset of the implementation literature (surveyed, for example, by Jackson [13])–a scoring rule is a single-agent mechanism whereby it is always a strictly dominant strategy for an agent to reveal her true preference.

*I would like to thank Federico Echenique, PJ Healy, and John Ledyard for useful comments and suggestions. All errors are my own.

Importantly, the theory of scoring rules meshes well with the "as if" approach of classical decision theory. The probability measure corresponding to a subjective expected utility agent is merely part of the representation of preference; however, it is uniquely defined. This is no problem for the theory of scoring rules–an agent is only required to make a choice from among a menu of *acts*; there need not be any mention of the word "probability" by whoever offers this menu to the agent.

While Brier provides the first example of a proper scoring rule, McCarthy [14] (see also Savage [17]) fully characterizes the proper scoring rules. Scoring rules are commonly used in the experimental literature to elicit probabilities, starting with the original work of McKelvey and Page [15]. Camerer [5] (p. 592-593) discusses the use of scoring rules in experimental economics. The most important point from the perspective of this note is that proper scoring rules *exist* in subjective expected utility environments.

The preceding results all operate under the presumption that a given agent's behavior conforms to the subjective expected utility axioms. Of course, the subjective expected utility paradigm is not universal. Ellsberg [7] demonstrates this, and it is at this point taken as given in the decision theory literature that the behavior of many agents does *not* conform to either the Savage or Anscombe and Aumann axioms. Informally speaking, there may be uncertainty about probabilities of certain events. Such uncertainty is referred to in the literature as "ambiguity."[1] A well-known model due to Schmeidler [18], the Choquet expected utility model, is able to accommodate this type of behavior. An agent whose behavior conforms to the Choquet expected utility model also has a unique set function representing beliefs; however, this function is not necessarily additive (we often refer to this unique set function as the "beliefs" of an agent). Thus, the model is more general than the subjective expected utility model. Another important model that can accommodate Ellsberg-type behavior is the multiple priors model, axiomatized by Gilboa and Schmeidler [12]. This model features an agent who can be viewed as possessing a set of priors. Such an agent evaluates the utility of an act by taking the minimal expected utility of the act across all priors in her set.

Ultimately, one would like to design a scoring rule allowing an agent to express beliefs reflecting subjectively ambiguous situations. Unfortunately, our first primary result is that there exists no analogue of a proper scoring rule for these more general decision models. We demonstrate this impossibility on the smallest well-known extension of the subjective expected utility paradigm. The particular extension under consideration is compatible with the Choquet expected utility model; but it additionally requires that the set function representing beliefs is *totally monotone*. This gives a broad-ranging impossibility result for non-expected utility models. Indeed, the result establishes that the existence of proper scoring rules for subjective expected utility models is knife-edge.

Given the impossibility, we attempt to understand how an agent whose behavior does not conform to the subjective expected utility paradigm might act if she faces a proper

[1] General theories of ambiguity are found in the works of Epstein [8], Epstein and Zhang [9], and Ghirardato and Marinacci [10].

scoring rule that is designed to elicit the beliefs of a subjective expected utility agent. For this environment, we consider an agent whose behavior conforms with the multiple priors model.[2] Suppose that she is offered a proper scoring rule that could be used to elicit the probability measure of a subjective expected utility agent. Our second main result shows that she will choose an act corresponding to some probability measure in her set of priors. The main result shows exactly how to find which of these acts she will choose. Moreover, for every probability in her set of priors, there exists some proper scoring rule for which it is optimal for her to choose the act corresponding to the probability. To establish these results, we determine a dual optimization problem that the agent may equivalently perform.

Section 2 introduces the model. Section 3 discusses our primary results. Section 4 concludes.

2 The model

Let Ω be a finite set of states of the world. An **act** is a function $x : \Omega \to \mathbb{R}$. The set of acts is denoted $\mathcal{F}$. A **capacity** is a function $\nu : 2^{\Omega} \to \mathbb{R}$ which is monotonic (i.e. for all $E, F \subset \Omega$, if $E \subset F$, then $\nu(E) \leq \nu(F)$), and is normalized so that $\nu(\varnothing) = 0$, and $\nu(\Omega) = 1$. A capacity is **totally monotone** if for all $\{E_1, ..., E_n\} \subset 2^{\Omega}$,

$$\nu\left(\bigcup_{j=1}^{n} E_j\right) \geq \sum_{J \subset \{1,...,n\}} (-1)^{|J|+1} \nu\left(\bigcap_{j \in J} E_j\right).$$

A capacity is **convex** if for all $A, B \in 2^{\Omega}$,

$$\nu(A) + \nu(B) \leq \nu(A \cap B) + \nu(A \cup B).$$

For all $E \subset \Omega$, $E \neq \varnothing$, define $\nu_E(F) = \begin{cases} 1 \text{ if } E \subset F \\ 0 \text{ otherwise} \end{cases}$. A classical representation theorem, due to Dempster, Shafer, and Shapley (see Shapley [19], for example), states that ν is a capacity if and only if for all $E \subset \Omega$, there exists $\alpha(\nu, E) \in \mathbb{R}$ for which $\sum_{E \neq \varnothing} \alpha(\nu, E) = 1$ such that $\nu = \sum_E \alpha(\nu, E)\, \nu_E$. Further, ν is totally monotone if and only if $\alpha(\nu, E) \geq 0$ for all E (Choquet [6], Chapter VII, has a more general theorem which predates the work of Shapley). Totally monotone capacities form the convex hull of $\{\nu_E\}_{E \subset 2^{\Omega} \setminus \varnothing}$. Denote the set of totally monotone capacities on Ω by $\mathcal{TM}(\Omega)$ and denote the set of probability measures on Ω by $\Delta(\Omega)$.

[2]While the multiple priors model does not generalize the Choquet expected utility model, it is a generalization of the Choquet expected utility model under the additional assumption that decision makers are uncertainty averse, as defined by Schmeidler [18], (see p. 582-583, especially parts (ix) and (x) of the Proposition).

All probability measures are totally monotone capacities and all totally monotone capacities are convex. A capacity is a probability measure if and only if the corresponding α assigns positive value only to singletons.

A risk-neutral Choquet expected utility maximizer evaluates acts $x : \Omega \rightarrow \mathbb{R}$ through the use of the Choquet integral [3]:

$$E_{\nu}[x] = \int_{\Omega} x(\omega)\, d\nu(\omega).$$

The Choquet integral is defined as follows: for all ν, let $\{\alpha(\nu, E)\}_E$ be its associated list of weights. Then

$$E_{\nu}[x] = \sum_{E \in 2^{\Omega} \setminus \varnothing} \min_{\omega \in E} \{x(\omega)\} \alpha(\nu, E).$$

Preferences conforming to the Choquet expected utility model were introduced and axiomatized by Schmeidler [18].

Let $\mathcal{C}$ be some set of capacities. A **scoring rule on** $\mathcal{C}$ is a function $f : \mathcal{C} \rightarrow \mathcal{F}$ for which for all $\nu, \nu' \in \mathcal{C}$,

$$E_{\nu}[f(\nu)] \geq E_{\nu}[f(\nu')].$$

Obviously, scoring rules exist; simply fix $x \in \mathcal{F}$, and let $f(\nu) \equiv x$ for all $\nu \in \mathcal{C}$. A scoring rule is **proper**[4] if for all $\nu, \nu' \in \mathcal{C}$ for which $\nu \neq \nu'$, $E_{\nu}[f(\nu)] > E_{\nu}[f(\nu')]$. An agent facing a proper scoring rule acts in her best interest (in an ex-ante sense) by choosing the act corresponding to her capacity.

3 Results

3.1 The nonexistence of proper scoring rules

The following result is due to McCarthy [14], Theorem 1.

Theorem 1 (McCarthy): The function $f : \Delta(\Omega) \rightarrow \mathcal{F}$ is a proper scoring rule on $\Delta(\Omega)$ if and only if there exists some strictly convex function $g : \Delta(\Omega) \rightarrow \mathbb{R}$ for which for all $\nu \in \Delta(\Omega)$, the function $h_{\nu} : \Delta(\Omega) \rightarrow \mathbb{R}$ defined by $h_{\nu}(\nu') \equiv E_{\nu'}[f(\nu)]$ is contained in the subdifferential of g at ν.[5]

[3] Note that we require our decision-maker to be risk-neutral. This is also a feature of the pioneering work of Brier [4], McCarthy [14], and Savage [17]. However, the key feature of risk-neutrality is that the decision maker has a von Neumann-Morgenstern utility index which is linear. If one accepts the theory of Anscombe and Aumann [2], then decision makers need not be risk-neutral, and one can use lotteries as payoffs. Instead of monetary compensation, state-contingent payoffs would be in the probability of winning some alternative. This idea is first introduced in Allen [1]. The only difference is the requirement that probabilities must be bounded. However, any monetary proper scoring rule will also necessarily be bounded, at least in the subjective expected utility model.

[4] Sometimes, such a scoring rule is referred to as **strictly proper**.

[5] Formally speaking, $g \geq h_{\nu}$, and $g(\nu) = h_{\nu}(\nu)$.

Proof (sketch): It is simple to verify that such a function g induces a proper scoring rule on $\Delta(\Omega)$. Conversely, suppose that f is a proper scoring rule on $\Delta(\Omega)$. Define $g : \Delta(\Omega) \to \mathbb{R}$ as $g(\nu) = \sup_{\nu'} h_{\nu'}(\nu)$. As the supremum of convex functions, g is also convex. Moreover, as f is a proper scoring rule, we know that $g(\nu) = \max_{\nu'} h_{\nu'}(\nu)$, and moreover, the maximum is achieved uniquely at $\nu' = \nu$. Thus, g is in fact strictly convex. In other words, the linear functional h_ν is included in the subdifferential of g at ν, but is included in no other subdifferential. ■

The preceding characterization illustrates that the set of proper scoring rules on $\Delta(\Omega)$ is quite large. When considering more general decision models; however, the situation is dramatically different. We will show that there does not exist a proper scoring rule on $\mathcal{TM}(\Omega)$. This implies a corresponding result for any superset of $\mathcal{TM}(\Omega)$, as well as other related models. Note that this impossibility result *does not* imply that the function representing beliefs cannot be elicited. It merely implies that it cannot be elicited from a single choice. An interesting question is to understand which families of beliefs can be elicited from a finite (possibly sequential) number of choices.

The intuition for the result is simple. By the Dempster, Shafer, Shapley representation, the set $\mathcal{TM}(\Omega)$ is isomorphic to $\Delta\left(2^\Omega \backslash \varnothing\right)$. Hence the set $\mathcal{TM}(\Omega)$ is of an exponentially higher dimension that the set $\Delta(\Omega)$. The set of acts has a higher dimension than $\Delta(\Omega)$ (as $\Delta(\Omega)$ is simply a simplex in $\mathbb{R}^\Omega$, and $\mathbb{R}^\Omega$ is the same as $\mathcal{F}$). However, as soon as there is more than one state of the world, the set $\mathcal{TM}(\Omega)$ has a higher dimension than $\mathcal{F}$. It therefore becomes more difficult to use acts in $\mathcal{F}$ to distinguish between decision makers with differing capacities.

Theorem 2: If $|\Omega| > 1$, there does not exist a proper scoring rule on the domain $\mathcal{TM}(\Omega)$.

Proof. Without loss of generality, we may work with the Dempster, Shafer, and Shapley representation of totally monotone capacities. A scoring rule then maps from $\Delta\left(2^\Omega \backslash \varnothing\right)$ into $\mathcal{F}$. If $\alpha, \beta \in \Delta\left(2^\Omega \backslash \varnothing\right)$ and $\alpha \neq \beta$, by properness,

$$\sum_{E \in 2^\Omega \backslash \varnothing} \min_{\omega \in E} \{f(\alpha)(\omega)\} \alpha(E) > \sum_{E \in 2^\Omega \backslash \varnothing} \min_{\omega \in E} \{f(\beta)(\omega)\} \alpha(E).$$

As the payoff from telling the truth is $U(\alpha) = \sum_{E \in 2^\Omega \backslash \varnothing} \min_{\omega \in E} \{f(\alpha)(\omega)\} \alpha(E)$, the function

$$U(\alpha) = \sup_{\beta \in \Delta(2^\Omega \backslash \varnothing)} \sum_{E \in 2^\Omega \backslash \varnothing} \min_{\omega \in E} \{f(\beta)(\omega)\} \alpha(E)$$

is a supremum of linear functionals, and must be strictly convex.

We will show that it is impossible for U to be strictly convex. Let $W(\Omega)$ be the set of weak orders on Ω.[6] We will say an act $x : \Omega \to \mathbb{R}$ **is monotonic with respect to** $\preceq \in W(\Omega)$ if $x(\omega) \geq x(\omega') \Leftrightarrow \omega \succeq \omega'$. Note that $W(\Omega)$ is a finite set.

[6]An order $\preceq$ is a weak order if it is complete and transitive.

For all acts $x \in \mathcal{F}$, there exists an order $\preceq$ with respect to which x is monotonic. Note that if x is monotonic with respect to $\preceq$, then

$$\sum_{E \in 2^{\Omega} \backslash \varnothing} \min_{\omega \in E} \{x(\omega)\} \alpha(E) = \sum_{E \in 2^{\Omega} \backslash \varnothing} x\left(\arg\min_{\omega \in E} \preceq\right) \alpha(E).$$

Denote by $x(\preceq)$ the set of acts that are monotonic with respect to $\preceq$. For all $\preceq$ for which $x(\preceq) \cap f(\mathcal{TM}(\Omega))$ is nonempty, define $U^{\preceq}(\alpha) = \sup_{x \in x(\preceq) \cap f(\mathcal{TM}(\Omega))} \sum_{E \in 2^{\Omega} \backslash \varnothing} x(\arg\min_{\omega \in E} \preceq) \alpha(E)$. Note that each such $U^{\preceq}$ is convex and subdifferentiable on the boundary of $\Delta\left(2^{\Omega} \backslash \varnothing\right)$.

We claim that for all V open in $\Delta\left(2^{\Omega} \backslash \varnothing\right)$ and convex, there exists $\alpha, \beta \in V$, $\alpha \neq \beta$, and $\lambda \in (0,1)$ such that $U^{\preceq}(\alpha) = U^{\preceq}(\beta) = U^{\preceq}(\lambda\alpha + (1-\lambda)\beta)$, so that $U^{\preceq}$ is not strictly convex. Let $\omega^{*} \equiv \arg\min_{\omega \in \Omega} \preceq$, and consider any $E \neq \{\omega^{*}\}$ which contains ω^{*}. Then in particular, $\arg\min_{\omega \in E} \preceq = \omega^{*}$ and $\arg\min_{\omega \in \{\omega^{*}\}} \preceq = \omega^{*}$. As V is open, there exists $\alpha \in V$ for which $\alpha(\{\omega^{*}\}) > 0$. Let $\varepsilon < \alpha(\{\omega^{*}\})$ be small enough so that

$$\beta(F) = \left\{ \begin{array}{c} \alpha(F) - \varepsilon \text{ if } F = \{\omega^{*}\} \\ \alpha(F) + \varepsilon \text{ if } F = E \\ \alpha(F) \text{ otherwise} \end{array} \right\}$$

is contained in V (as V is open, such an ε exists). As for all $F \notin \{\{\omega^{*}\}, E\}$, $\beta(F) = \alpha(F)$, it is clear that

$$\begin{aligned} & \sup_{x \in x(\preceq) \cap f(\mathcal{TM}(\Omega))} \sum_{E \in 2^{\Omega} \backslash \varnothing} x\left(\arg\min_{\omega \in E} \preceq\right) \alpha(E) \\ = & \sup_{x \in x(\preceq) \cap f(\mathcal{TM}(\Omega))} \sum_{E \in 2^{\Omega} \backslash \varnothing} x\left(\arg\min_{\omega \in E} \preceq\right) \beta(E). \end{aligned}$$

Moreover, it is also clear that for any $\lambda \in (0,1)$,

$$\begin{aligned} & \sup_{x \in x(\preceq) \cap f(\mathcal{TM}(\Omega))} \sum_{E \in 2^{\Omega} \backslash \varnothing} x\left(\arg\min_{\omega \in E} \preceq\right) \alpha(E) \\ = & \sup_{x \in x(\preceq) \cap f(\mathcal{TM}(\Omega))} \sum_{E \in 2^{\Omega} \backslash \varnothing} x\left(\arg\min_{\omega \in E} \preceq\right) (\lambda\alpha + (1-\lambda)\beta)(E). \end{aligned}$$

Therefore, $U^{\preceq}$ is not strictly convex on *any* open neighborhood V.

Clearly, $U = \sup_{\preceq} U^{\preceq}$. We claim that there exists some open V and some $\preceq$ for which $U|_V = U^{\preceq}|_V$. The theorem will then be complete, as U is not strictly convex on V. Enumerate the functions $\left\{U^{\succeq}\right\}$ as $\left\{U^{1}, ..., U^{K}\right\}$. Clearly, $\Delta\left(2^{\Omega} \backslash \varnothing\right)$ is relatively open. Set $V^{1} = \Delta\left(2^{\Omega} \backslash \varnothing\right)$. For all $i = 2, ..., K$, set $V^{i} = \{\alpha \in V^{i-1} : \text{there exists } j \text{ such that } U^{j}(\alpha) > U^{i-1}(\alpha)\}$. Set $V^{K+1} = \varnothing$. Clearly, V^{i} is open for all i, and for all $i = 1, ...K$, $V^{i} \subset V^{i-1}$. Let i^{*} satisfy $V^{i^{*}} \neq \varnothing$ and

$V^{i^*+1} = \varnothing$ (clearly such an i^* must exist). As $V^{i^*+1} = \varnothing$, it must be that $U^{i^*}(\alpha) \geq U^i(\alpha)$ for all i on V^{i^*}. We therefore establish that there exists some open set V for which there exists $\preceq^*$ for which $U = U^{\preceq^*}$. This is an immediate contradiction to the strict convexity of U. Therefore, there exists no proper scoring rule on $\mathcal{TM}(\Omega)$. ■

Corollary 1: For all $\mathcal{C}$ for which $\mathcal{TM}(\Omega) \subset \mathcal{C}$, there does not exist a proper scoring rule on $\mathcal{C}$.

The preceding corollary applies especially to the case of those individuals who evaluate acts with respect to convex capacities (as all totally monotone capacities are convex). It also applies to the general model of biseparable preferences of Ghirardato and Marinacci [11]. We will now discuss a related model for which an impossibility is obtained.

Say a preference ordering is a **multiple priors** ordering if there exists a nonempty, closed, convex set $P \subset \Delta(\Omega)$ such that the decision maker evaluates actions x according to $\min_{p \in P} E_p[x]$. Multiple priors preferences were first axiomatized by Gilboa and Schmeidler [12].

We may define another notion of a scoring rule; one in which agents announce sets of priors. Denote by $\mathcal{K}(\Delta(\Omega))$ the nonempty, compact, convex subsets of $\Delta(\Omega)$.

Corollary 2: There does not exist a function $f : \mathcal{K}(\Delta(\Omega)) \to \mathcal{F}$ for which for all $P, P' \in \mathcal{K}(\Delta(\Omega))$ with $P \neq P'$

$$\min_{p \in P} E_p[f(P)] > \min_{p \in P} E_p[f(P')].$$

Proof. For all $\nu \in \mathcal{TM}(\Omega)$, let $C(\nu) = \{p \in \Delta(\Omega) : p(E) \geq \nu(E) \text{ for all } E \subset \Omega\}$. A classic result of Schmeidler (for example, see Schmeidler [18], p. 582-583) states that for all $E \subset \Omega$, $\nu(E) = \min\{p(E) : p \in C(\nu)\}$, and moreover, $E_\nu[x] = \min_{p \in C(\nu)} E_p[x]$ for all $x \in \mathcal{F}$. The corollary now follows trivially. ■

3.2 Choices made by agents facing scoring rules on $\Delta(\Omega)$

The question addressed here is the following. Let f be a scoring rule on $\Delta(\Omega)$, and let $P \subset \Delta(\Omega)$ be closed and convex. For which p^* is it true that

$$\min_{p \in P} E_p[f(p^*)] \geq \min_{p \in P} E_p[f(p')]$$

for all $p' \in \Delta(\Omega)$? When a multiple priors decision maker is faced with a scoring rule on $\Delta(\Omega)$, which probability measure can we expect her to reveal?

One important question is the general existence of a solution to the optimization problem. The following example demonstrates that existence is not guaranteed.

Example: Given a proper scoring rule on $\Delta(\Omega)$, say, f, and a set of priors $P \subset \Delta(\Omega)$, there need not exist a solution to the optimization problem

$$\max_{p^* \in \Delta(\Omega)} \min_{p \in P} E_p \left[f \left(p^* \right) \right].$$

To see this, let $\Omega = \{0,1\}$ and let $P = \Delta(\Omega)$. Without loss of generality, we may identify $\Delta(\Omega)$ with $[0,1]$, so that a decision maker is required to announce the probability of $\{1\}$. By Theorem 1, a scoring rule on $\Delta(\Omega)$ is identified with a subdifferential mapping of a strictly convex function $g : [0,1] \to \mathbb{R}$. Consider the function g defined by:

$$g(x) \equiv \begin{cases} (2x-1)^2 & \text{if } x < 1/2 \\ 2(2x+1)(2x-1) & \text{if } x \geq 1/2 \end{cases}.$$

Note that g is smooth everywhere except for its boundaries, and at the point $1/2$. Therefore, when choosing a scoring rule on $\Delta(\Omega)$, we are free to pick any element of the subdifferential at $1/2$. We may therefore choose the following scoring rule: for $p < 1/2$,

$$\begin{aligned} f(p)(0) &= -4p^2 + 1 \\ f(p)(1) &= -4p^2 + 8p - 3, \end{aligned}$$

and for $p \geq 1/2$,

$$\begin{aligned} f(p)(0) &= -2 - 8p^2 \\ f(p)(1) &= 16p - 2 - 8p^2. \end{aligned}$$

Therefore, for $p < 1/2$, $\min_{p' \in \Delta(\Omega)} E_{p'}[f(p)] = -4p^2 + 8p - 3$, and for $p \geq 1/2$, $\min_{p' \in \Delta(\Omega)} E_{p'}[f(p)] = -2 - 8p^2$. It is easy to see that a maximum does not exist in this case. Note, however, that $\sup_{p \in \Delta(\Omega)} \min_{p' \in \Delta(\Omega)} E_{p'}[f(p)] = 0$, and can be achieved by taking a sequence $\{p_n\}$ such that $p_n = 1/2 - \varepsilon_n$, where $\varepsilon_n > 0$ and $\varepsilon_n \to 0$. Thus, the supremum is achieved in a neighborhood of $1/2$, which is the minimizer of g. We will see that the minimizer of g over a compact set of priors will be the typical choice for a multiple priors decision maker facing a scoring rule on $\Delta(\Omega)$.

Note the feature of the preceding example: the scoring rule f is discontinuous at the point $p = 1/2$. This is because there is a "kink" in the function g from which f is defined. To this end, we are concerned primarily with *continuous* scoring rules on $\Delta(\Omega)$; those are the scoring rules for which the corresponding function from Theorem 1 $g : \Delta(\Omega) \to \mathbb{R}$ is everywhere differentiable.

Theorem 3: Let f be a continuous proper scoring rule on $\Delta(\Omega)$ and let P be convex and compact. Then $\arg\max_{p^* \in \Delta(\Omega)} \min_{p \in P} E_p[f(p^*)]$ exists, is a singleton, and is equal to $\arg\min_{p^* \in P} \max_{p \in \Delta(\Omega)} E_{p^*}[f(p)]$. In particular, the unique solution to the optimization problem is an element of P. Furthermore, for all $p \in P$, there exists a continuous probabilistic scoring rule whose corresponding solution is p.

Proof. Existence of a solution follows trivially from continuity and compactness of $\Delta(\Omega)$. The remainder of the proof is an application of the minimax theorem. Define X to be the convex hull of $\{f(p)\}_{p\in\Delta(\Omega)}$. As x is continuous in p and as $\Delta(\Omega)$ is compact, the set $\{f(p)\}_{p\in\Delta(\Omega)}$ is compact, thus X is compact (as Ω is finite). Consider the function $G:\Delta(\Omega)\times X\to\mathbb{R}$ defined by

$$G(p,x)=E_p[x].$$

This function is clearly bilinear. Moreover, by the Sion minimax Theorem (Berge [3], p. 210), there exist $p^*\in P$ and $x^*\in X$ such that for all $(p,x)\in P\times X$,

$$E_{p^*}[x]\leq E_{p^*}[x^*]\leq E_p[x^*].$$

In particular, for given p^*, the unique maximizer of $E_{p^*}[x]$ over X is $f(p^*)$, so that the inequality reads:

$$E_{p^*}[x]\leq E_{p^*}[f(p^*)]\leq E_p[f(p^*)].$$

Moreover, we therefore may write

$$E_{p^*}[f(p)]\leq E_{p^*}[f(p^*)]\leq E_p[f(p^*)],$$

for all $p\neq p^*$. Hence,

$$\min_{p\in P}\max_{p'\in\Delta(\Omega)}E_p[f(p')]=\max_{p'\in\Delta(\Omega)}\min_{p\in P}E_p[f(p')],$$

and is achieved at $p=p^*$, $p'=p^*$. We claim that p^* is the unique element of $\arg\min_{p\in P}\max_{p'\in\Delta(\Omega)}E_p[f(p')]$; this follows trivially by the strict convexity of $g(p')=\max_{p'\in\Delta(\Omega)}E_p[f(p')]$ and the fact that P is convex and compact. Moreover, p^* is also the unique element of $\arg\max_{p'\in\Delta(\Omega)}\min_{p\in P}E_p[f(p')]$. To see this, let $p'\in\Delta(\Omega)$, $p'\neq p^*$. Then

$$\begin{aligned}&\min_{p\in P}E_p[f(p')]\\ \leq\ & E_{p^*}[f(p')]\\ <\ & E_{p^*}[f(p^*)]\\ =\ & \max_{p''\in\Delta(\Omega)}\min_{p\in P}E_p[f(p'')].\end{aligned}$$

Here, the first inequality follows as $p^*\in P$ and the second follows as f is a proper scoring rule on $\Delta(\Omega)$. Hence, $\max_{p'\in\Delta(\Omega)}\min_{p\in P}E_p[f(p')]$ is achieved (uniquely) at $p^*\in P$, which is the unique minimizer of the strictly convex function $g(p')=\max_{p'\in\Delta(\Omega)}E_p[f(p')]$ over P.

To see that for all $p^*\in P$, there exists a continuous probabilistic scoring rule for which $\arg\max_{p\in\Delta(\Omega)}\min_{p^*\in P}E_{p^*}[f(p)]=\{p^*\}$, simply let g be a strictly convex and smooth function $g:\Delta(\Omega)\to\mathbb{R}$ whose global minimum is achieved at p^* (for example, let $g(p)\equiv\|p-p^*\|^2$, where $\|\cdot\|$ denotes the Euclidean norm). Let $x:\Delta(\Omega)\to\mathbb{R}$ be the continuous probabilistic scoring rule which consists of the subdifferentials of g. Then $g(p') = \max_{p \in \Delta(\Omega)} E_{p'}[f(p)]$, from which we utilize our preceding result. ■

Suppose there are only two states of the world, say $\Omega = \{1, 2\}$. In such an environment, a multiple priors agent is characterized by the maximal and minimal probabilities she attributes to one of the states. These maximal and minimal probabilities can be completely recovered by using two distinct scoring rules on $\Delta(\Omega)$. The first such scoring rule, say f_1, has a corresponding function g_1 from Theorem 1 which is minimized uniquely at the probability measure placing probability one on state 1. The second scoring rule, f_2, has a corresponding function g_2 from Theorem 1 which is minimized uniquely at the probability measure placing probability one on state 2.

4 Conclusion

We have shown that proper scoring rules do not exist for agents whose behavior is not consistent with the subjective expected utility paradigm. However; this does not preclude the existence of simple methods of determining a multiple priors agent's set of priors, for example. It only demonstrates the impossibility of eliciting this set of priors using a single decision. For example, we have shown how, in a two-state environment, one may elicit the set of priors using only two menus of acts. An interesting question, therefore, is to find a minimal collection of menus that an agent must face in order to completely determine her set of priors in a more general enviornment.

References

[1] F. Allen, Discovering personal probabilities when utility functions are unknown, *Management Science* 33 (1987), 542-544.

[2] F.J. Anscombe and R. Aumann, A definition of subjective probability, *Annals of Mathematical Statistics* 34 (1963), 199-205.

[3] C. Berge, Topological Spaces, Dover Publications, Minneola NY, (1997).

[4] G. Brier, Verification of forecasts expressed in terms of probability, *Monthly Weather Review* 78 (1950), 1-3.

[5] C. Camerer, Individual decision making *in* Handbook of Experimental Economics, J. Kagel and A. Roth (eds.), Princeton University Press, Princeton NJ, 1995.

[6] G. Choquet, Theory of capacities, *Annales de l'institut Fourier* 5 (1954), 131-295.

[7] D. Ellsberg, Risk, ambiguity, and the Savage axioms, *Quarterly Journal of Economics* 75 (1961), 643-669.

[8] L. Epstein, A definition of uncertainty aversion, *Review of Economic Studies* 66 (1999), 579-608.

[9] L. Epstein and J. Zhang, Subjective probabilities on subjectively unambiguous events, *Econometrica* 69 (2001), 265-305.

[10] P. Ghirardato and M. Marinacci, Ambiguity made precise: A comparative foundation, *Journal of Economic Theory* 102 (2002), 251-289.

[11] P. Ghirardato and M. Marinacci, Risk, ambiguity, and the separation of utility and beliefs, *Mathematics of Operations Research* 26 (2001), 864-890.

[12] I. Gilboa and D. Schmeidler, Maxmin expected utility with nonunique prior, *Journal of Mathematical Economics* 18 (1989), 141-153.

[13] M. Jackson, A crash course in implementation theory, *Social Choice and Welfare* 18 (2001), 655-708.

[14] J. McCarthy, Measures of the value of information, *Proceedings of the National Academy of Sciences* 42 (1956), 654-655.

[15] R.D. McKelvey and T. Page, Public and private information: An experimental study of information pooling, *Econometrica* 58 (1990), 1321-1339.

[16] L.J. Savage, The Foundations of Statistics, Dover Publications, Minneola NY, (1972).

[17] L.J. Savage, Elicitation of personal probabilities and expectations, *Journal of the American Statistical Association* 66 (1971), 783-801.

[18] D. Schmeidler, Expected utility and probability without additivity, *Econometrica* 57 (1989), 571-587.

[19] L.S. Shapley, A value for n-person games, *Annals of Mathematical Studies* 28 (1953), 307-317.