SYMMETRY OR DYNAMIC CONSISTENCY?*

Larry G. Epstein Kyoungwon Seo

July 5, 2010

Abstract

In a setting with repeated experiments, where evidence about the experiments is symmetric, a decision-maker ranks bets (or acts) over their outcomes. We describe a stark modeling trade-off between symmetry of preference (indifference to permutations), dynamic consistency and ambiguity. Then, assuming that experiments are ordered in time, we outline an axiomatic model of preference that exhibits dynamic consistency and yet models learning under ambiguity.

1. INTRODUCTION

We consider a setting with many experiments and an individual who ranks bets (or acts) on their outcomes. Prior evidence about the experiments is symmetric, and thus there is no reason for her to distinguish between them. This translates intuitively into a restriction on preference that we call Symmetry. However, there may be little information available about any of them and thus experiments may be seen to be ambiguous.

We make two contributions. First, we prove a theorem that establishes a stark trade-off between Symmetry and the possibility of dynamically consistent updating. One or the other must be relaxed if ambiguity is to be accommodated. In [7], we describe a model that accommodates ambiguity, satisfies Symmetry,

*Boston University, lepstein@bu.edu and Department of Managerial Economics and Decision Sciences, Northwestern University, k-seo@kellogg.northwestern.edu. This research was supported by the National Science Foundation (awards SES-0917740 and SES-0918248). First version: May 2009.

and also admits dynamically consistent updating in a limited, but still interesting class of environments, namely, where an individual first samples and observes the outcomes of some experiments, and then chooses how to bet on the outcomes of remaining experiments. Thus each experiment serves *either* as a signal *or* is payoff relevant, but *not both*. Statistical decision-making often fits into this framework.

Our second contribution in this paper is to describe an alternative response to the above trade-off for the case where experiments are ordered in time, and where the time interval between experiments is significant. Then, nonindifference to the way that uncertainty is resolved over time (in the spirit of Kreps and Porteus [17]) is plausible and can justify relaxing Symmetry, even where evidence about experiments is symmetric. Essentially, the fact that experiments are conducted at significantly different times introduces an asymmetry between them. This permits recursive preferences and hence (unqualified) dynamic consistency.

The bottom line is a model that generalizes de Finetti's [11] exchangeable Bayesian model, the canonical model of learning about a "parameter", to incorporate ambiguity while retaining simple and dynamically consistent updating. Applications in macroeconomics and finance come to mind; an example is portfolio choice where there is learning about the unknown mean return to the uncertain security, while returns are thought to be affected also by poorly understood idiosyncratic factors.[1]

2. A MODELING TRADE-OFF

2.1. A Theorem

Let S be finite and $\Omega = S^N = S_1 \times ... \times S_N$, where $S_i = S$ for all i and $N < \infty$. Think of a series of experiments (coin-tosses, for example), where the outcome of each lies in S. An act f is a mapping from Ω into $[0, 1]$. The set of all acts is $\mathcal{F}$. An individual has a (complete and transitive) preference order $\succeq_0$ on $\mathcal{F}$. The real-valued outcomes of an act can be thought of as denominated in utils, and derived in the familiar way from a more primitive ranking of Anscombe-Aumann acts.

Suppose that evidence about the experiments is symmetric, so that there is no reason to distinguish between them. The natural implication for preference is as follows. Let Π be the set of permutations on $\{1, ..., N\}$. For $\pi \in \Pi$ and

[1] The model of learning in Epstein and Schneider [4] has similar applications and is closely related. It and other related literature are discussed below in the concluding remarks.

$\omega = (s_1, ..., s_N) \in \Omega$, let $\pi\omega = \left(s_{\pi(1)}, ..., s_{\pi(N)}\right)$. Given an act f, define the permuted act πf by $(\pi f)(s_1, ..., s_N) = f\left(s_{\pi(1)}, ..., s_{\pi(N)}\right)$.

Axiom 1 (Symmetry). *$f \sim_0 \pi f$ for all acts f and permutations π.*

At issue is whether this ex ante preference can be updated so as to deliver dynamic consistency. Let $\succeq_{n,s_1^n}$ denote preference on $\mathcal{F}$ conditional on the sample $s_1^n = (s_1, ..., s_n)$, $1 \leq n < N$. Dynamic consistency imposes the following restriction across conditional and ex ante preferences.

Axiom 2 (Dynamic Consistency (DC)). *For all $1 \leq n < N$, samples s_1^{n-1}, and acts $f', f \in \mathcal{F}$,*

$$f' \succeq_{n,(s_1^{n-1},s_n)} f \text{ for all } s_n \implies f' \succeq_{n-1,s_1^{n-1}} f,$$

and the latter preference is strict if, in addition, $f' \succ_{n,(s_1^{n-1},s_n)} f$ for some s_n.

One way to ensure dynamic consistency is to update in the trivial way whereby every $\succeq_{n,s_1^n}$ is set equal to $\succeq_0$. We rule this out by imposing Consequentialism - the conditional ranking given the sample s_1^n does not take into account what the acts might have delivered had a different sample been realized.

Axiom 3 (Consequentialism). *For all $1 \leq n < N$, samples s_1^n, and acts $f', f \in \mathcal{F}$,*

$$f' \sim_{n,s_1^n} f \text{ if } f'\left(s_1^n, \cdot\right) = f\left(s_1^n, \cdot\right).$$

The theorem below shows that the preceding (plus some mild regularity conditions) imply that the ex ante preference $\succeq_0$ is *additive*: $\succeq_0$ has a utility function V of the form

$$V(f) = \Sigma_{r \in S_1 \times ... \times S_N} \, v_r\left(f\left(s_1, ..., s_N\right)\right), \quad f \in \mathcal{F}.$$

The two regularity conditions are continuity (in the standard sense, identifying $\mathcal{F}$ with a subset of Euclidean space) and *sensitivity to every sample path*, which means: For every act f with $f\left(S^N\right) \subset (0,1]$, and for every $(s'_1, ..., s'_N) \in S_1 \times ... \times S_N$, then $f' \not\sim_0 f$ for *some* act f' such that $f'(s_1, ..., s_N) = f(s_1, ..., s_N)$ for all $(s_1, ..., s_N) \neq (s'_1, ..., s'_N)$. The obvious notion of strict monotonicity of preference is sufficient, but is too strong for present purposes since it is often violated by

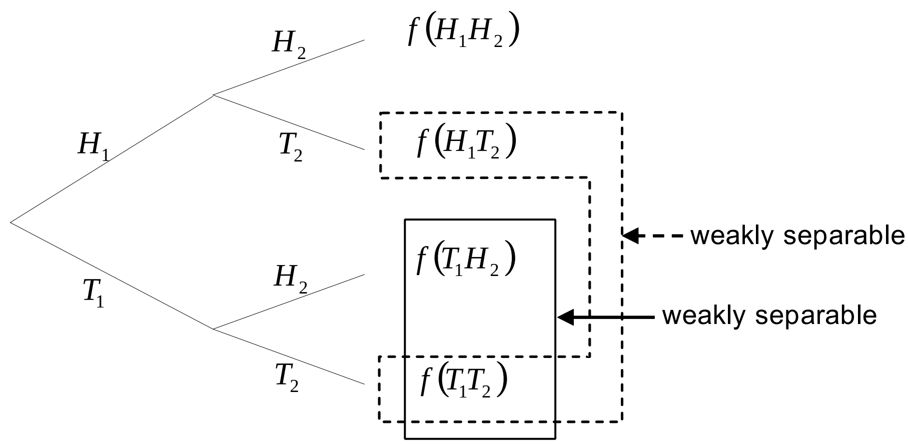

Figure 2.1: Overlapping Weak Separability

ambiguity averse preferences.[2] However, sample sensitivity is an extremely weak condition.

Theorem 2.1. *Let $\succeq_0$ satisfy Symmetry, and suppose that it is continuous and sensitive to every sample path. Suppose also that $\{\succeq_{n,s_1^n}: 1 \leq n < N\}$ satisfy Consequentialism and, in conjunction with $\succeq_0$, Dynamic Consistency. Then $\succeq_0$ and all conditional preferences are additive.*

The theorem is a corollary of Gorman's "overlapping theorem" [13, Theorem 1]. The intuition is conveyed by Figure 2.1 for two coin tosses ($N = 2$).

Each act f over $\{H_1, T_1\} \times \{H_2, T_2\}$ can be identified with the vector $(f(H_1H_2), f(H_1T_2), f(T_1H_2), f(T_1T_2))$. Say that $\{H_1H_2, H_1T_2\}$ is *weakly separable* for the preference $\succeq_0$ if for any two acts f' and f that agree on $\{T_1H_2, T_1T_2\}$, their ranking according to $\succeq_0$ is invariant to any change in those common outcomes; define weak separability similarly for other sets of terminal nodes. Conditional preferences at the intermediate nodes in the tree satisfy Consequentialism, and they jointly agree with ex ante preference $\succeq_0$ in the sense of DC. It follows that $N_1 = \{H_1H_2, H_1T_2\}$ is weakly separable for $\succeq_0$. By Symmetry, therefore, applying the permutation that switches the two experiments, deduce that

[2] Think of the Leontief utility function $W(x) = \min\{x_i : i = 1, .., n\}$, which is not strictly increasing on $(0, 1]^n$, but is sensitive to every co-ordinate, since $W(x_{-i}, .)$ is not constant, at least if $x_{-i} >> 0$.

$N_2 = \{H_1H_2, T_1H_2\}$ is also weakly separable. Since these sets overlap (they intersect but neither is contained in the other), Gorman's theorem implies that the following sets are also weakly separable:

$$N_1 \cap N_2 = \{H_1H_2\},\ N_1 \backslash N_2 = \{H_1T_2\},\ N_2 \backslash N_1 = \{T_1H_2\},$$

$$N_3 = N_1 \cup N_2 = \{H_1H_2, H_1T_2, T_1H_2\}, \text{ and } N_1 \triangle N_2 = \{H_1T_2, T_1H_2\}.$$

Similarly, reversing the roles of Heads and Tails in N_3, it follows that $N_4 = \{T_1T_2, H_1T_2, T_1H_2\}$ is weakly separable. Because N_3 and N_4 are overlapping weakly separable sets, Gorman's theorem implies the weak separability of $N_3 \triangle N_4 = \{H_1H_2, T_1T_2\}$. Separability of N_1 and N_4 implies that of $N_1 \triangle N_4 = \{T_1T_2, H_1H_2, T_1H_2\}$. Proceed in this way to show that all subsets of terminal nodes are weakly separable. Additivity follows from Debreu [2, Theorem 3].

2.2. Discussion

The significance of additivity that concerns us is that it rules out ambiguity, which, as illustrated by the Ellsberg Paradox, is inherently about complementarities across states.[3] This is illustrated by the fact that a multiple-priors preference (Gilboa and Schmeidler [12]) on $\mathcal{F}$ is additively separable if and only if it conforms to subjective expected utility (SEU) theory; and similarly for every other axiomatic model of ambiguity averse preferences in the literature. However, the conflict revealed by the theorem is deeper since it is not limited to any particular axiomatic model or functional form for utility, nor is it tied to a particular updating rule (except via Consequentialism).

There are many instances where there is no reason for distinguishing between experiments, particularly in a cross-sectional context. Repeated tosses of a single coin is the canonical example. The common assumption that realizations of experiments are drawn independently from a fixed but possibly unknown distribution implies Symmetry - this is one direction in the de Finetti Theorem [11] on exchangeable probability measures.[4] Dynamic consistency is a compelling normative requirement. In particular, both Symmetry and DC have normative appeal for statistical decision-making. Therefore, one might interpret the theorem as an

[3] See Epstein and Zhang [9] for a formal behavioral definition of ambiguity in terms of such complementarities.

[4] Though the de Finetti theorem can be viewed as a result in probability theory alone, it is typically understood in economics as describing the prior in the subjective expected utility model of choice. That is how we view it.

argument against introducing ambiguity into normative models given the setting of repeated experiments.

Our view, expressed in [7], is that ambiguity is to be expected. For the case of repeated tosses of a single coin, for example, if prior information about the coin's bias is poor, then the bias might be treated like the draw of a ball from an ambiguous Ellsberg urn, at least for the first few tosses before learning can have much effect. A different reason has to do specifically with the multiplicity of experiments. For example, the single coin may be tossed by different people - their tossing techniques may be thought to matter, but be poorly understood and seen as idiosyncratic. More generally, if experiments are thought to be affected by many poorly understood idiosyncratic factors, then even a statistician may not be certain that experiments are identical, or conducted under identical conditions. A normative model should respect the decision-maker's limitations and help her to make choices given her lack of confidence. We argue in the cited paper that ambiguity averse preferences can, but SEU preferences cannot, qualify as such guides to choice. Thus, we interpret the theorem as posing a modeling trade-off between Symmetry and DC.

As noted in the introduction, in [7] we put forth a model that relaxes Dynamic Consistency. Next we describe a model that relaxes Symmetry.

3. "SYMMETRY" IN A TEMPORAL CONTEXT

When experiments are ordered in time and the time interval between experiments is significant, a case can be made for relaxing Symmetry even where evidence about experiments is symmetric, which then permits recursive preferences and hence Dynamic Consistency.

Consider the bet H_1T_2 that Heads will be followed by Tails in the first two tosses of a coin, where the prizes are 1 and 0 denominated in utils; define T_1H_2 similarly. If both tosses are carried out (almost) simultaneously, then presumably the two bets would be indifferent - there is no reason for distinguishing between the two tosses. Suppose, however, that the coins are tossed one week apart. Suppose further that prior beliefs are that Heads is very likely and Tails very unlikely on either toss. Then for T_1H_2 most of the uncertainty is resolved after the outcome of the first toss is realized, while for H_1T_2, the individual would likely have to wait an extra week to learn how she fared. If she does not enjoy living with the uncertainty, she might prefer T_1H_2, while if she enjoys consuming the hope of winning for as long as possible, then she would prefer H_1T_2. In general, the fact that experiments are conducted at significantly different times introduces an asymmetry between them if the individual cares about when a given uncertainty is resolved (in the spirit of Kreps and Porteus [17]). The resulting violation of Symmetry opens the door to Dynamic Consistency, though the normative status of such a model is not clear.[5]

We outline a formal model. Assume that there are infinitely many experiments ($N = \infty$).[6] Since the timing of resolution must matter, the evaluation of an act (that is, preference) depends on the information structure. First we consider (ex ante) preference $\succeq_0^*$ on $\mathcal{F}$ when it is expected that all uncertainty will be resolved at a single time - the precise time is not important as long as it occurs after the choices dictated by $\succeq_0^*$ are made. The significance of one-shot resolution is that the information structure treats all experiments symmetrically, which eliminates any argument against Symmetry. Thus assume that $\succeq_0^*$ satisfies Symmetry. One can think of it as being the preference relevant when experiments are cross-sectional, which is one interpretation of the model in [7].

Similarly, $\succeq_{n,s_1^n}^*$ denotes preference on $\mathcal{F}$ after observing a sample s_1^n, and assuming that all future experiments will be resolved at a single time. The intention is not of a dynamic process where the ex ante expectation of one-shot resolution is later violated. Rather, $\succeq_0^*$ and $\succeq_{n,s_1^n}^*$ are two separate snap-shots, and both can in principle be elicited.

View $\{\succeq_0^*, \succeq_{n,s_1^n}^*\}$ as auxiliary primitives. Our main focus is preference given the natural filtration corresponding to the given ordering of experiments. Thus we consider also the preference process $\{\succeq_0, \succeq_{n,s_1^n}\}$, with the obvious interpretation. As illustrated by the coin-tossing example above, Symmetry is not compelling for these preferences. We assume, however, that they satisfy Dynamic Consistency. All conditional preferences are assumed to satisfy Consequentialism.

The following basic properties are also assumed for every (conditional or unconditional) preference.

Axiom 4 (Basic). *Each preference is complete, transitive, continuous, and admits a unique certainty equivalent, that is, for every act f, there exists a unique constant act p in $[0,1]$ that is indifferent to f.*

[5] Epstein and Le Breton [3] describe another functional-form-free modeling trade-off that can be interpreted as saying that, (under suitable conditions that are not comparable to those considered here), ambiguity and dynamic consistency can co-exist only if one permits nonindifference to the way in which a given uncertainty is resolved over time.

[6] The trade-off described in the theorem still applies since the theorem implies additivity on the set of all acts that depend on any finite number of experiments.

The main content of the model lies in the connection between the two collections of preferences. Both $\succeq^*_{n,s_1^n}$ and $\succeq_{n,s_1^n}$ are conditioned on the same sample. On the other hand, they are based on different perspectives about the future since only the latter expects gradual resolution of the remaining uncertainty. However, for acts f that depend only on the next experiment, all relevant uncertainty is resolved in one-shot also for the natural filtration. Therefore, the two preferences should agree on one-step-ahead acts. Similarly for ex ante preferences $\succeq^*_0$ and $\succeq_0$. Thus we adopt:

Axiom 5 (One-Step-Ahead). *For all $n \geq 0$, p in $[0,1]$ and acts f over S_{n+1}, if $f \sim^*_{n,s_1^n} p$, then $f \sim_{n,s_1^n} p$.*

By Basic, we can define "certainty equivalent" utility functions for each preference. Thus define U_0^* by

$$U_0^*(f) = p \text{ iff } f \sim^*_0 p,$$

and similarly for U_n^*, U_0 and U_n. The model's content is summarized by the implied relation between $\{U_n\}$ and $\{U_n^*\}$.

The upshot of the axioms is clear. Dynamic Consistency implies that the U_n's satisfy a recursive relation, which at each stage involves only the evaluation of the uncertain utility payoff next period, and thus only one-step-ahead acts. The axiom One-Step-Ahead implies that such acts are evaluated as if using U_n^*. Therefore, the entire process $\{U_n\}$ is pinned down once one takes a stand on the utility functions that apply under one-shot resolution of uncertainty.

Instead of restating all this formally, which the reader can easily do for himself, we describe (without proof, which is straightforward) the result when one-shot resolution preferences $\{\succeq^*_0, \succeq^*_{n,s_1^n}\}$ are modeled as in [7]. In fact, so as to minimize new notation and for illustrative simplicity, we specialize that model further and adopt the (axiomatically based) specification in [8]. In that model, ambiguity is modeled by using *belief functions* instead of additive probability measures. (See Shafer [18], for example. When $S = \{H, T\}$, a belief function on S can be thought of simply as a probability *interval* for Heads.) Denote by $Bel\,(S)$ the set of belief functions on S.[7]

[7] With the topology induced by Choquet integration of continuous real-valued functions, $Bel\,(S)$ is compact metric. Below integrals with respect to a belief function are intended in the Choquet sense.

If one adds the assumptions Consequentialism, Dynamic Consistency, Basic, and One-Step-Ahead for $\{\succeq_0, \succeq_{n,s_1^n}\}$, then one obtains the following representation for the corresponding utility functions $\{U_n\}$: For each n (and suppressed history s_1^n), there exists a (Borel) probability measure μ_n on $Bel\,(S)$ such that $\{U_n\}$ satisfies the recursive relation, (for $n \geq 0$ and in obvious notation),

$$U_n\,(f) = \int_{Bel(S)} \left(\int_{S_{n+1}} U_{n+1}\,(f \mid s_{n+1})\, d\theta\,(s_{n+1}) \right) d\mu_n\,(\theta)\,. \tag{3.1}$$

The interpretation is that each experiment is seen to be characterized by a common belief function θ, which is unknown. Given θ, then $\int_{S_{n+1}} U_{n+1}\,(f \mid s_{n+1})\, d\theta\,(s_{n+1})$ gives the "expected," in the Choquet sense, continuation value of f. Since uncertainty about θ at stage n is described by μ_n, the outer expectation, (this time a standard integral, because μ_n is a probability measure) gives the value of f today. The exchangeable Bayesian model due to de Finetti is the special case where each μ_n has support on $\Delta\,(S) \subset Bel\,(S)$, the set of probability measures on S.

The other important component of the Bayesian model is Bayes' Rule. In our model, the updating process generating $\{\mu_n\}$ is identical to the one described in [7] and has the following form: There exists a likelihood function $L : Bel\,(S) \rightarrow \Delta\,(\Omega)$ such that, for all $n \geq 1$,

$$d\mu_n\,(\theta) = \frac{L_n\,(s_n \mid \theta)}{\overline{L}_n\,(s_n)} d\mu_{n-1}\,(\theta)\,, \tag{3.2}$$

where $L_n\,(s_n \mid \theta)$ is the one-step-ahead conditional of $L\,(\cdot \mid \theta)$ at stage n, $L_n : S^{n-1} \times Bel\,(S) \rightarrow \Delta\,(S)$, and

$$\overline{L}_n\,(\cdot) = \textstyle\int L_n\,(\cdot \mid \theta)\, d\mu_{n-1}\,(\theta)\,,$$

is a probability measure on S having full support.

In other words, the process of posteriors $\{\mu_n\}$ is identical to that in an expected utility model where the Bayesian prior $\overline{L}\,(\cdot) \in \Delta\,(\Omega)$ is given by

$$\overline{L}\,(\cdot) = \textstyle\int L\,(\cdot \mid \theta)\, d\mu_0\,(\theta)\,.$$

The existence of a shadow Bayesian model promotes tractability, since it permits application of results from the Bayesian literature about the dynamics of posteriors, though interpretations differ since our model of choice is not Bayesian. Some such applications are provided in [7].

4. CONCLUDING REMARKS

We conclude with some perspective on the model.

The preference process, and in particular, the ex ante preference $\succeq_0$, depends on the given ordering of experiments. Thus, for any other ordering and the natural filtration that it induces, the ex ante preference constructed as above will in general be different, reflecting again that the timing of resolution matters. Such dependence is a feature of all recursive models, since the result of backward induction depends, outside the expected utility framework, on the filtration being assumed. We view such dependence as plausible and as an unavoidable by-product, rather than as an objective, of our modeling approach. A parallel in the literature may be helpful. Epstein and Zin [10] show that allowing temporal resolution to matter permits a disentangling of risk aversion from intertemporal substitution. In the same spirit, we allow temporal resolution to matter in order to achieve a distinction between risk and ambiguity in a setting with repeated experiments.

The temporal setting, and the focus on learning when the environment is not completely understood, suggests connections to literature in macroeconomics/finance (Epstein and Schneider [4, 5], Hansen [14], Chen, Ju and Miao [1] and Ju and Miao [15], for example). Though our model is formulated in terms of terminal consumption or payoffs, it is straightforward to extend it to have consumption streams as the source of utility. Thus we view our temporal model as adding to the tool-kit of dynamic models of learning under ambiguity. One difference from the other models cited is that only the present one is a fully axiomatic generalization of de Finetti's canonical Bayesian model of learning.

Our model has two other noteworthy features. First, though preferences are ambiguity averse, and the individual is concerned that experiments may differ in some unspecified way, updating is identical to that for a shadow Bayesian agent with suitable prior. The simplicity that this affords is attractive - there is no need to deal with the issue of how to update sets of priors, for example, and one can import results from Bayesian learning theory. The models in [14, 15, 1, 16] share this simplicity - in all cases, updating proceeds exactly as in a Bayesian model and ambiguity aversion enters only in the way that posterior beliefs are used to define preference. Updating is more complicated in [4, 5]. On the other hand, simplicity is not the only relevant criterion. Epstein and Schneider [5] show, by means of an Ellsberg-style thought experiment, that their model of updating *can* capture an individual's concern with "signal or information quality," but that

Bayesian updating, *cannot*; see also a similar example in [6]. Thus the modeler faces another trade-off, between simplicity and a desire to capture information quality.

Finally, the Epstein-Schneider model has the following "model closure" property: not only do all conditional preferences over one-step-ahead acts conform to the multiple-priors model, but ex ante preference on the whole tree also belongs to the multiple-priors class. Our model is not "closed" in this sense.[8] For example, for the belief function based specification (3.1), it is *not* the case that there exists a belief function ν on S^∞ such that $U_0(\cdot)$ is the Choquet expected utility function using ν. We find "model closure" appealing because it reduces the degrees of freedom available to the modeler. On the other hand, if it comes at a cost - and here the cost is that simple Bayesian updating is excluded thereby - then the more important question becomes "is model closure necessary for a coherent model?" Our answer is that it is not. For example, multiple-priors utility is a two-period model, where choice is made in period 1 and all uncertainty is resolved in period 2; and the intuition for Gilboa and Schmeidler's axioms relies implicitly on this one-shot resolution of uncertainty. There is no room in their model for a different temporal resolution and for the associated preference consequences. Thus it makes sense to impose the axioms on all the one-shot-resolution preferences $\succeq_0^*$ and $\succeq_{n,s_1^n}^*$, and on the evaluation of one-step-ahead acts by $\succeq_0$ and $\succeq_{n,s_1^n}$, but there is no compelling reason for imposing them more broadly. A similar remark applies for the axioms corresponding to belief function utility.

A. Appendix: Proof of Theorem

The proof amounts to translating the obvious event tree into the formalism of demand theory so that the Gorman theorem is seen to apply. (In fact, we use a corollary of his theorem that is easier to apply in the general (nonbinary) case.) It is enough to show that $\succeq_0$ is additive.

Each tuple $t = (s_1, ..., s_N) \in S_1 \times ... \times S_N \equiv T$ defines a terminal node in the N-stage tree, of which there are $M = K^N$, where K is the cardinality of S. The set of acts $\mathcal{F}$ may be identified with the Cartesian product $Y_T = \Pi_{t \in T} Y_t$, where $Y_t = [0,1]$ for all t. (In particular, the domain of $\succeq_0$ is a connected and topologically separable product space.) For any subset of terminal nodes $T' \subset T$, let $Y_{T'} = \Pi_{t \in T'} Y_t$, and write $Y_T = Y_{T'} \times Y_{T \setminus T'}$. Say that T' is *weakly separable*

[8]Neither are the models in [14, 15, 1, 16].

if the ranking $\left(y'_{T'}, z_{T\backslash T'}\right) \succeq_0 \left(y_{T'}, z_{T\backslash T'}\right)$ is invariant to the common vector $z_{T\backslash T'}$ describing outcomes at terminal nodes in $T\backslash T'$. Sensitivity to sample paths implies Gorman's assumption P4 - every sector is strictly essential. (This is true only if $Y_t = (0,1]$ for all t, but this difference is not important given continuity of preference.)

An intermediate node in the tree corresponds to the set of all terminal nodes with fixed outcomes for experiments $1, ...n$, for some $n < N$, that is, to sets T' of the form

$$T' = \{(\overline{s}_1, ..., \overline{s}_n, s_{n+1}, ..., s_N) : (s_{n+1}, ..., s_N) \in S_{n+1} \times ... \times S_N\}. \tag{A.1}$$

DC and the fact that each conditional preference satisfies Consequentialism imply that each subset T' of the form in (A.1) is weakly separable. But by Symmetry, the order of experiments does not matter and thus every set of the form

$$T' = \{(s_1, ..., s_n, ..., s_N) \in S_1 \times ... \times S_N : s_{1_j} = \overline{s_{i_j}},\ j = 1, ..., n\},\ n < N,$$

is also weakly separable. In other words, every nonsingleton (because $n < N$) cylinder is weakly separable. By the overlapping theorem, any singleton is weakly separable because it can be expressed as an intersection of nonsingleton cylinders. Apply Gorman's Corollary on p. 382. Let $\mathcal{B}$ denote the set of all cylinders. Any nonsingleton cylinder overlaps some other cylinder. (T' and T'' overlap if they intersect and neither contains the other.) This implies the Corollary's hypothesis that singletons are the only proper components. The conclusion is that utility is additively separable. ■

References

[1] H. Chen, N. Ju and J. Miao, Dynamic asset allocation with ambiguous return predictability, revised 2009.

[2] G. Debreu, Topological methods in cardinal utility theory, in *Mathematical Methods in the Social Sciences*, K. Arrow, G. Karlin and P. Suppes eds., Stanford U. Press, 1960.

[3] L.G. Epstein and M. Le Breton, Dynamically consistent beliefs must be Bayesian, *J. Econ. Theory* 61 (1993), 1-22.

[4] L.G. Epstein and M. Schneider, Learning under ambiguity, *Rev. Ec. Studies* 74 (2007), 1275-1303.

[5] L.G. Epstein and M. Schneider, Ambiguity, information quality and asset pricing, *J. Finance* 63 (2008), 197-228.

[6] L.G. Epstein and M. Schneider, Ambiguity and asset markets, *Annual Rev. Finan. Econ.*, forthcoming.

[7] L.G. Epstein and K. Seo, Symmetry of evidence without evidence of symmetry, *Theor. Econ.*, forthcoming.

[8] L.G. Epstein and K. Seo, A de Finetti theorem for capacities: ambiguity about correlation, 2009.

[9] L.G. Epstein and J. Zhang, Subjective probabilities on subjectively unambiguous events, *Econometrica* 69 (2001), 265-306.

[10] L.G. Epstein and S. Zin, Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework, *Econometrica* 57 (1989), 937-969.

[11] B. de Finetti, La prevision: ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincare* 7 (1937), 1-68. English translation in *Studies in Subjective Probability*, 2nd edition, H.E. Kyburg and H.E. Smokler eds., Krieger Publishing, Huntington NY, 1980, pp. 53-118.

[12] I. Gilboa and D. Schmeidler, Maxmin expected utility with non-unique priors, *J. Math. Econ.* 18 (1989), 141-153.

[13] T. Gorman, The structure of utility functions, *Rev. Econ. Stud.* 35 (1968), 367-390.

[14] L.P. Hansen, Beliefs, doubts and learning: valuing macroeconomic risk, *Amer. Ec. Rev.* 97 (2007), 1-30.

[15] N. Ju and J. Miao, Ambiguity, learning and asset returns, 2009.

[16] P. Klibanoff, M. Marinacci and S. Mukerji, Recursive smooth ambiguity preferences, *J. Econ. Theory* 144 (2009), 930-976.

[17] D. Kreps and E.L. Porteus, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46 (1978), 185-200.

[18] G. Shafer, *A Mathematical Theory of Evidence*, Princeton U. Press, Princeton, 1976.