

Time and No Lotteries: A Simple Axiomatization of Maxmin Expected Utility*

Asen Kochov[†]

February 15, 2012

Abstract

The paper provides an axiomatization of the maxmin expected utility model in a multi-period Savage-style framework of purely subjective uncertainty.

*I thank Larry Epstein and William Thomson for helpful comments.

[†]Department of Economics, University of Rochester, Rochester, NY 14627. E-mail: asen.kochov@rochester.edu.

1 Introduction

The paper provides an axiomatization of maxmin expected utility in a multi-period Savage-style framework of subjective uncertainty. Two of the conditions we impose are surprisingly simple: history independence and stationarity. The third condition is novel. It captures the idea that ‘smoothing across states’ is more important than ‘smoothing across time’. A number of reasons motivate the analysis.

The seminal paper on maxmin, Gilboa and Schmeidler [13], assumes that state-contingent outcomes are lotteries with exogenously specified probabilities. This specification of the domain of choice was introduced by Anscombe and Aumann [3], and is henceforth referred to as the AA framework. The AA framework simplifies the analysis but is rarely used in descriptive modeling. In contrast, most applications of the maxmin model in finance and macroeconomics involve intertemporal decisions. The reader is referred to the recent survey by Epstein and Schneider [8]. Characterizing the maxmin model in an intertemporal setting facilitates the understanding and interpretation of these results.

Many papers have emphasized the foundational and practical problems associated with adopting an AA domain. In response, [1, 4, 12] characterize the maxmin model in a general Savage-style framework without time or lotteries. These characterizations involve axioms that are not directly verifiable. The axiomatization provided here is arguably as simple as the original one in Gilboa and Schmeidler [13]. Thus, the intertemporal framework proves amenable to axiomatic analysis and raises the possibility of new ways to test the model.

The main idea is simple. As a model of choice under uncertainty, maxmin is a model of *intratemporal* behavior. This is why previous studies abstract from issues pertaining to time altogether. If time is introduced, however, intratemporal properties have *intertemporal* manifestations. The latter are easier to analyze if one makes the standard assumption that utility is additive across time periods.

Making the link between intratemporal and intertemporal behavior operational comes at a price. We have already mentioned time additivity. Its merits and weaknesses have been discussed extensively in both axiomatic and empirical work; we do not repeat them here. More significantly, our

approach leads to a model of intertemporal utility which does not permit a separation between risk attitudes and the degree of intertemporal substitution. Such a separation has proved useful in understanding many issues in macroeconomics. An extension of the analysis in this direction is undoubtedly important.

Even though the focus of this paper is on the maxmin model, the arguments can be adapted to characterize other models of ambiguity such as the Choquet model of Schmeidler [24], and the multiplier model introduced by Anderson et al. [2] and recently axiomatized in an AA setting by Strzalecki [26]. The derivation of subjective expected utility is also technically easier in a multi-period rather than an atemporal setting. In particular, it requires no restrictions on the cardinality of the state space. This point has been made by Gorman [15]. It is the subject of ongoing work to see if the setting can be used to study the more general models of ambiguity aversion developed by Maccheroni et al. [20] and Cerreia-Vioglio et al. [5]. It is nonetheless reasonable to conclude that the multi-period setting provides a viable alternative to the Anscombe-Aumann framework. It is tractable and more relevant to applied work.

2 Domain

Time is discrete and varies over an infinite horizon $T := \{0, 1, 2, \dots\}$. The information structure is described by a filtered space $(\Omega, \{\mathcal{F}_t\}_t)$ where Ω is an arbitrary set of states of the world and $\{\mathcal{F}_t\}_t =: \mathcal{F}$ is an increasing sequence of algebras such that $\mathcal{F}_0 = \{\Omega, \emptyset\}$. If Ω is finite, one may think of the information structure as an event tree whose nodes correspond to time-event pairs $(t, \omega) \in T \times \Omega$.

In every period, outcomes lie in a connected, separable, compact space X . An **act** is an X -valued, \mathcal{F} -adapted processes. Each such process is the multi-period counterpart of a Savage act; it has the form $h = (h_t)_t$ where $h_t : \Omega \rightarrow X$ is \mathcal{F}_t -measurable for every $t \in T$.

An act h is **simple** if there exists a finitely-generated algebra $\mathcal{A} \subset \cup_t \mathcal{F}_t$ such that h_k is \mathcal{A} -measurable for every $k \in T$. The domain of choice is the set \mathcal{H} of all simple acts. The domain excludes acts whose outcomes may depend on the realization of tail events. Such events do not take place in finite time

and are therefore unobservable.¹

An act is **deterministic** if its outcomes do not depend on the state of the world. The space of deterministic acts may be identified with the space X^∞ . Generic elements of X^∞ are denoted by d, d' . The bold-faced letter \mathbf{x} is reserved for the deterministic act which gives $x \in X$ in every period. Given some $t \in T$, another important subset of \mathcal{H} is $X^{t+1} \times \mathcal{H}$, whose elements are denoted as $(x_0, x_1, \dots, x_t, h)$. For example, $g := (x, h)$ is the act such that $g_0 = x$ and $g_t = h_{t-1}$ for all $t > 0$.

3 Axioms

The primitive of the model is a preference relation over the space of simple acts \mathcal{H} . The first two axioms impose the familiar order properties and a weak continuity requirement.

Order The preference relation \succeq is complete and transitive.

Continuity For every $h \in \mathcal{H}$, the sets $\{d : d \succeq h\}$ and $\{d : h \succeq d\}$ are closed in the product topology on X^∞ .

Notice that we are not imposing continuity in the product topology on \mathcal{H} . In fact, the latter need not hold if there are infinitely many events in $\cup_t \mathcal{F}_t$. To see this, consider a sequence of events A_n that increases monotonically to Ω . Continuity in the product topology would require that the certainty equivalent of a bet on A_n changes continuously as $A_n \nearrow \Omega$. In the context of ambiguity, a discontinuity may arise since Ω is unambiguous whereas all of the events A_n may be ambiguous.

The next axiom is a monotonicity requirement. To facilitate a comparison with its atemporal analogue, it is convenient to think of an act h as a function from Ω into X^∞ . Thus, $h(\omega)$ denotes the infinite stream of outcomes delivered by h in state ω .

¹A tail event is an event in $\sigma(\cup_t \mathcal{F}_t) \setminus \cup_t \mathcal{F}_t$. A preference on the set of simple acts can be extended to the set of all acts using the techniques developed in Epstein and Wang [11].

Monotonicity For every $t \in T$, $A \in \mathcal{F}_t$, $d, d' \in X^\infty$ and $h \in \mathcal{H}$,

$$d \succeq d' \quad \Rightarrow \quad \left[\begin{array}{ll} d & \text{if } \omega \in A \\ h(\omega) & \text{if } \omega \in A^c \end{array} \right] \succeq \left[\begin{array}{ll} d' & \text{if } \omega \in A \\ h(\omega) & \text{if } \omega \in A^c \end{array} \right] \quad (3.1)$$

whenever $d_k = d'_k = h_k$ for all $k < t$.

Two features of the axiom may require elaboration. First, the opposite implication in (3.1) is not assumed and may indeed be violated if the event A is null. Second, the restriction that d, d' and h agree up to period t is made to insure that all acts in (3.1) are \mathcal{F} -adapted. This restriction has no analogue in the atemporal setting. For further discussion of the axiom, see Lemma 3 in Section 5.

The next two assumptions are the core of the present axiomatization. The first one imposes a weak preference for ‘time diversification’ or ‘intertemporal hedging’. In particular, it says that combining some uncertain outcome, h_t , occurring in period t , with another uncertain outcome, h_{t+1} , occurring in period $t + 1$, is at least as good as combining their respective certainty equivalents x_t, y_{t+1} . The intuition is that the two uncertain outcomes, h_t and h_{t+1} , may hedge one another, eliminating any ambiguity about the overall outcome. In contrast, combining their certainty equivalents produces no such complementarity.

Intertemporal Hedging For every $t \in T$, $h \in \mathcal{H}$, $x, y \in X$ and $d \in X^\infty$,

$$(d_{-t}, h_t) \sim (d_{-t}, x) \text{ and } (d_{-(t+1)}, h_{t+1}) \sim (d_{-(t+1)}, y) \quad \Rightarrow \\ (d_{-t, t+1}, h_t, h_{t+1}) \succeq (d_{-t, t+1}, x, y).$$

It is interesting to note that, under the hypothesis of the axiom, the standard time-additive expected-utility model implies that $(d_{-t, t+1}, h_t, h_{t+1})$ and $(d_{-t, t+1}, x, y)$ are necessarily indifferent. This implication of the standard model has no immediate justification. It is a by-product of two assumptions that have been typically studied in isolation: separability across states and separability across time. In its own right, this feature of the standard model has come under criticism even in the context of objective uncertainty. See, for example, Epstein and Tanny [9]. To see what is at stake, suppose that h_t and h_{t+1} have the same certainty equivalent x . If h_t and h_{t+1} are different, combining them ‘smooths’ outcomes across states. At the same time, it

induces variability across time. In the standard model, these effects cancel out. Intertemporal Hedging permits a wider scope of behavior. It expresses the idea that ‘smoothing across states’ might be more important to the individual than ‘smoothing across time’. Arguably, the appeal of the axiom is stronger in settings in which probabilities are not given and difficult to assess.

A possible criticism of the axiom is that a strict preference for $(d_{-t,t+1}, h_t, h_{t+1})$ is not justified if $h_t = h_{t+1}$ in which case the two bets do not hedge one another. Although not immediately obvious, such a preference is ruled out by the next axiom.

Stationarity For all acts $h, h' \in \mathcal{H}$ and outcomes $x \in X$:

$$h \succeq h' \Leftrightarrow (x, h) \succeq (x, h')$$

The act (x, h) is obtained from h by advancing the timing or ‘delivery date’ of each outcome by one period and adding x in the initial period. The axiom requires that preference be invariant to such changes. In the context of deterministic choice, the axiom is introduced and studied extensively by Koopmans [17, 18].

One obvious implication of the axiom is **History Independence**, that is, $(x, h) \succeq (x, h')$ if and only if $(y, h) \succeq (y, h')$, for all x, y and h, h' . History Independence is typically violated in applications in which the acts h are interpreted as income streams and the individual cares only about the distribution over *total* wealth. Thus, the present analysis is intended primarily for choice situations involving consumption streams in which *intertemporal* considerations are important. In this case, the axiom is more compelling and widely maintained.

Invariance to the timing of outcomes may be violated even if the ranking of deterministic acts is stationary. An individual may become more likely to accept a bet when its consequences are delayed sufficiently far into the future. The intuition is that, after discounting, the ‘stakes’ become negligible and the bet ‘less uncertain’. Such behavior is accommodated by the variational preferences in [21], which are history independent, but its normative appeal seems debatable.

The final axiom is restrictive but common in modeling intertemporal behavior. Consider two deterministic acts which are identical from period $t = 2$

onwards. The axiom requires that their ranking does not depend on the common continuation.

Time Separability For every $x, x', y, y' \in X$ and $d, d' \in X^\infty$,

$$(x, y, d) \succeq (x', y', d) \Leftrightarrow (x, y, d') \succeq (x', y', d').$$

Koopmans [18] shows that Stationarity and Time Separability imply an additive representation over deterministic acts.

4 Main Result

Endow the space of finitely additive probability measures on $\cup_t \mathcal{F}_t$ with the weak* topology, i.e., the topology generated by all real-valued bounded measurable functions on Ω .

Theorem 1 *The preference \succeq satisfies Order, Continuity, Monotonicity, Intertemporal Hedging, Stationarity and Time Separability if and only if there exist a discount factor $\beta \in (0, 1)$, a continuous function $u : X \rightarrow \mathbb{R}$ and a nonempty weak*-closed convex set P of finitely additive probability measures on $\cup_t \mathcal{F}_t$ such that:*

$$h \succeq h' \Leftrightarrow \min_{p \in P} \int_{\Omega} [\sum_t \beta^t u(h_t)] dp \geq \min_{p \in P} \int_{\Omega} [\sum_t \beta^t u(h'_t)] dp. \quad (4.1)$$

Moreover, if not all h, h' are indifferent, then P and β are unique and u is unique up to positive affine transformations.

Despite the obvious similarities, the proof of Theorem 1 is not a straightforward application of the arguments in [13]. A key property of the multiple-prior functional is that it is positively homogenous.² In an AA framework, outcomes are objective lotteries and utility is linear in their probabilities. By varying the latter, one obtains homogeneity with respect to any given positive constant. In the present setting, the utility of deterministic acts is

²A functional $I : V \rightarrow \mathbb{R}$ defined on a real vector space V is positively homogenous if $I(\alpha x) = \alpha I(x)$ for every $x \in V$ and every $\alpha > 0$.

additive across time periods and one may think of the discount factor β^t as the ‘probability’ of the period- t outcome. The problem is that β is fixed and therefore homogeneity obtains in a limited sense only, that is, with respect to powers of the fixed constant β . In the appendix, this property is called β -homogeneity. The gist of the proof is to show that a β -homogeneous, monotone, translation invariant and superadditive functional is in fact positively homogeneous.

The infinite horizon plays an important role and cannot be replaced without some loss of generality. To see this, suppose there are finitely many periods, $T < \infty$, and all uncertainty resolves in the final period, that is, $\mathcal{F}_0 = \mathcal{F}_1 = \dots = \mathcal{F}_{T-1}$. It is then clear that Intertemporal Hedging loses most of its bite, since one of the functions h_t, h_{t+1} in the statement of the axiom is always constant. More generally, one needs to assume that $\mathcal{F}_{T-1} = \mathcal{F}_T$ in order to derive a maxmin representation over all acts. In addition, a positive rate of time preference, $\beta < 1$, which plays an indispensable role in the proof, would no longer be implied by Continuity, as shown in [18], and would have to be assumed explicitly.

Conclude this section with some comments on related literature. Kochov [16] uses Stationarity to derive the representation in (4.1) in a finite-horizon model. However, he did not recognize the full potential of the multi-period domain and followed the AA approach to defining ‘hedging’ or ‘uncertainty aversion’.

Strzalecki [25] studies a specialized environment in which shocks are independent and identically distributed across time. He restricts attention to a parametric class of recursive uncertainty-averse preferences, modeled after [5], and shows that maxmin is the only model which exhibits ‘indifference to the timing of uncertainty’. Although conceptually distinct, it can be verified that, within the environment and the class of preferences he studies, the proposed notion of ‘indifference to the timing of uncertainty’ is equivalent to Stationarity.

There is, on the other hand, no accepted notion of ‘indifference to the timing of uncertainty’ for the general environment studied here. Lemma 4 in the appendix shows that Order, Monotonicity and Stationarity imply a ‘reduction-style’ property that may be interpreted as such. However, the implied property is strictly weaker than Stationarity and does not restrict how ‘beliefs’ on Ω are modeled.

5 Dynamic Choice

This section outlines an extension of the analysis to a dynamic model in which choices can be made at any point in time. Assuming that behavior is dynamically consistent, the objective is to derive a recursive specification of the maxmin representation in (4.1). In an AA framework, such a specification is characterized axiomatically by Epstein and Schneider [7]. Building on the proof of Theorem 1, it becomes possible to derive the Epstein-Schneider results without the use of lotteries. Since, as they emphasize [7, p.4], the domain of Savage acts is typically the primary domain of interest, we view this extension as an important application of the theorem. As a separate contribution, we describe two ways in which the Epstein-Schneider axiomatization can be simplified. One of them is based on a general result that may be of independent interest.

Because the necessary changes to the arguments in [7] can be easily deduced from the proof of Theorem 1, we do not give a complete exposition.³ We focus on the primitives and additional axioms needed to derive the recursive formulation in [7].

As common in axiomatic work on dynamic choice, in this section we assume that every \mathcal{F}_t in the information structure \mathcal{F} is finitely generated. Dynamic choice is modeled by an \mathcal{F} -adapted process of preferences $\{\succeq_{t,\omega}\}$ where $\succeq_{t,\omega}$ is interpreted as the ranking of \mathcal{H} conditional on the information available at the node (t,ω) . To simplify the exposition, it is assumed, without further mention, that all preferences are complete, transitive and **nondegenerate**, that is, not all acts are indifferent.

The first axiom requires that when the individual evaluates an act h at a given node (t,ω) , he cares only about the outcomes of h in the continuation of that node. In particular, past outcomes and outcomes in unrealized parts of the tree are irrelevant.

Consequentialism For every $t \in T, \omega \in \Omega$, and acts $h, h' \in \mathcal{H}$, if $h_k(\omega') = h'_k(\omega')$ for all $k \geq t$ and $\omega' \in \mathcal{F}_t(\omega)$, then $h \sim_{t,\omega} h'$.

Unlike any of the axioms encountered so far, the next axiom links together preferences at different points in time. In particular, it insures that ‘plans’

³Theorem 1 is necessary is to establish an analogue of Lemma A2 in [7].

made at a given node in the tree remain optimal when reevaluated at later nodes.

Dynamic Consistency For every $t \in T, \omega \in \Omega$, and acts $h, h' \in \mathcal{H}$ such that $h_k = h'_k$ for all $k \leq t$,

$$\left[h \succeq_{t+1, \omega'} h' \text{ for all } \omega' \in \mathcal{F}_t(\omega) \right] \Rightarrow h \succeq_{t, \omega} h'.$$

Moreover, the latter ranking is strict, if, in addition, the former ranking is strict for some $\omega' \in \mathcal{F}_t(\omega)$.

The main result of this section can be stated as follows. Dynamic Consistency and Consequentialism, in conjunction with the assumption that \succeq_0 satisfies the axioms in Section 3, are necessary and sufficient for the recursive maximin representation in Epstein and Schneider [7]. In addition to using the ‘smaller’ domain of Savage acts, the result drops two of the assumptions made in [7]. The next lemmas elaborate and show how to deduce the additional assumptions in [7].

Versions of the first lemma are well-known in many different settings. It says that conditional preferences ‘inherit’ the properties of the ex ante preference \succeq_0 , whenever the process $\{\succeq_{t, \omega}\}$ is dynamically consistent and consequentialist.

Lemma 2 *Suppose $\{\succeq_{t, \omega}\}$ satisfies Dynamic Consistency and Consequentialism and the preference \succeq_0 satisfies the axioms in Section 3. With the exception of Order, the axioms in Section 3 are inherited by every conditional preference $\succeq_{t, \omega}$.*

If ‘tastes’ change with the passage of time, it is often difficult to isolate and identify changes in ‘beliefs’ that may be induced by the arrival of new information. It is therefore common in both axiomatic work and empirical studies to assume that the ranking of deterministic acts is invariant across states and time periods. Formally, this means that, for every $t \in T, \omega \in \Omega, x_0, \dots, x_{t-1} \in X$ and $d, d' \in X^\infty$:

$$(x_0, \dots, x_{t-1}, d) \succeq_0 (x_0, \dots, x_{t-1}, d') \Leftrightarrow (x_0, \dots, x_{t-1}, d) \succeq_{t, \omega} (x_0, \dots, x_{t-1}, d').$$

Like Dynamic Consistency, this axiom, which we call **Conditional State Independence**, links together preferences at different points in time. Versions of the axiom are explicitly assumed in [7, 21] among others. The next lemma shows that, under very general conditions, the axiom is implied by the monotonicity of \succeq_0 . Apart from dynamic consistency, the most substantive requirement is that the ranking of *deterministic* acts be stationary at any point in time. As the proof of the lemma makes it clear, this condition is far from necessary and can be replaced by other requirements depending on the context.

Lemma 3 *Suppose $\{\succeq_{t,\omega}\}$ satisfies Dynamic Consistency and Consequentialism and \succeq_0 satisfies Monotonicity. If, for every t and ω , the restriction of $\succeq_{t,\omega}$ to X^∞ is continuous, stationary and nondegenerate, then $\{\succeq_{t,\omega}\}$ satisfies Conditional State Independence.*

The result may be interpreted as follows. If the ex ante preference \succeq_0 is monotone, yet ‘tastes’ change with the passage of time, such changes reveal a failure of dynamic consistency. If there is no uncertainty, the result is a tautology. The proof of the lemma clarifies why an extension of the result to an uncertain environment is not immediate. In conclusion, we point out that there is an obvious converse to Lemma 3: If $\{\succeq_{t,\omega}\}$ satisfies Dynamic Consistency, Consequentialism and Conditional State Independence, then \succeq_0 satisfies Monotonicity.

6 Proofs

6.1 Theorem 1

With the exception of Continuity, necessity of the axioms is easily verified. To verify Continuity, it suffices to show that $\sum_t \beta^t u(x_t)$ is continuous in the product topology on X^∞ . Notice that, since X is compact, the continuous function $u : X \rightarrow \mathbb{R}$ is bounded. Since $\beta \in (0, 1)$, the series $\sum_t \beta^t u(x_t)$ is uniformly convergent. Continuity in the product topology follows from [22, Thm 7.10-7.11].

Turn to sufficiency. If all acts are indifferent, the result follows trivially. From now on, we assume that \succeq is nondegenerate. We begin by establishing two useful properties of \succeq .

Say that \succeq satisfies **Reduction** if, for every $h, h' \in \mathcal{H}$, $h \succeq h'$ whenever $h(\omega) \succeq h'(\omega)$ for every $\omega \in \Omega$. It is readily verified that Reduction implies Monotonicity.⁴ Assuming Order and Stationarity, the next lemma provides a converse.

Lemma 4 *Order, Monotonicity and Stationarity imply Reduction.*

Proof. Let h, h' be such that $h(\omega) \succeq h'(\omega)$ for every ω . Because h, h' are simple, there exists some t such that h_k, h'_k are \mathcal{F}_t -measurable for every k . Fix some $a = (x_0, \dots, x_{t-1}) \in X^t$ and consider the acts $(a, h), (a, h')$. By construction, $(a, h)(\omega) = (a, h(\omega))$. By Stationarity, $(a, h(\omega)) \succeq (a, h'(\omega))$ for every $\omega \in \Omega$. Moreover, $h \succeq h'$ if and only if $(a, h) \succeq (a, h')$, so it suffices to show the latter. Viewed as functions from Ω into X^∞ , both $(a, h), (a, h')$ have finite range. Let $\{A_1, A_2, \dots, A_n\}$ be the coarsest partition of Ω making both of them measurable. Replace the infinite stream of (a, h') on A_1 by the respective infinite stream of (a, h) . By Monotonicity, the new act is preferred to (a, h') . Take the new act and replace its infinite stream on A_2 by the respective infinite stream of (a, h) . Apply Monotonicity again. After n such steps, we obtain (a, h) . By transitivity, (a, h) is preferred to the initial act (a, h') , as desired. ■

Say that \succeq is **sensitive** if there exist $x, y \in X$ and $d \in X^\infty$ such that $(x, d) \succ (y, d)$.

Lemma 5 *If \succeq is nondegenerate and satisfies Order, Continuity, Monotonicity and Stationarity, then \succeq is sensitive.*

Proof. Suppose, by way of contradiction, that $(x, d) \sim (y, d)$ for all $x, y \in X$ and $d \in X^\infty$. From this and Stationarity in turn, it follows that $(z, x, d) \sim (z', x, d) \sim (z', y, d)$ for all z, z', x, y, d . Repeating the argument, conclude that any $d, d' \in X^\infty$ which differ in at most finitely many periods must be indifferent. Since the set of all pairs (d, d') that differ in at most finitely many periods is dense in $X^\infty \times X^\infty$ and preference is continuous, all $d, d' \in X^\infty$ are indifferent. Lemma 4 then implies that all $h, h' \in \mathcal{H}$ are indifferent, contradicting nondegeneracy. ■

⁴Reduction implies that the individual care only about the ‘distribution’ over outcome streams. The axiom is typically violated by the recursive preferences introduced in Kreps and Porteus [19] and further extended by Epstein and Zin [10].

Sensitivity is needed to apply the following lemma due to Koopmans [18]. It delivers an additively separable utility function for the set of deterministic acts X^∞ .

Lemma 6 *There exists a continuous function $u : X \rightarrow \mathbb{R}$ and a discount factor $\beta \in (0, 1)$, such that the restriction of \succeq to X^∞ is represented by the utility function:*

$$U(d) := (1 - \beta) \sum_t \beta^t u(x_t) \quad (6.1)$$

Moreover, β is unique and u is unique up to positive affine transformations.

Proof. There are some minor differences in the domain and axioms used by Koopmans [18] to derive (6.1). First, he assumes that X is metrizable and that \succeq is continuous in the sup metric on X^∞ . Second, he assumes a monotonicity axiom (P5) which is not made here. It is obvious that continuity in the product topology, assumed here, implies continuity in the sup metric. Moreover, because X^∞ is compact in the product topology, there exist best and worst elements in X^∞ . As Koopmans [18, p.84] notes, the existence of best and worst elements implies his monotonicity postulate P5. Finally, Koopmans' arguments are based on the results of [6, 14] for which metrizability is not needed. ■

Because X is compact and connected, the range of u is a compact interval. Because preference is sensitive, the interval has nonempty interior. Rescaling appropriately, it is w.l.o.g. to set $u(X) = [-1, 1]$. Let $x^* \in X$ be an outcome such that $u(x^*) = 0$.

The product space X^∞ inherits the topological properties of X . In particular, X^∞ is compact, separable and connected.

Let $\bar{\mathbf{x}}$ and $\underline{\mathbf{x}}$ be the best and worst elements in X^∞ , respectively. Lemma 4 implies that $\bar{\mathbf{x}} \succeq h \succeq \underline{\mathbf{x}}$. Because X^∞ is connected, a standard argument shows that for every act $h \in \mathcal{H}$ there exists a deterministic act $d_h \in X^\infty$ such that $d_h \sim h$. This allows us to extend the utility function U from X^∞ to the space of all acts \mathcal{H} . For every $h \in \mathcal{H}$, define $\tilde{U}(h) := U(d_h)$. The transitivity of \succeq implies that $\tilde{U} : \mathcal{H} \rightarrow \mathbb{R}$ is well-defined and represents the preference relation \succeq .

For every act $h \in \mathcal{H}$, define the **util-act** $U \circ h : \Omega \rightarrow \mathbb{R}$:

$$[U \circ h](\omega) := (1 - \beta) \sum_t \beta^t u(h_t(\omega)) \quad (6.2)$$

Let $\mathcal{U} = \{U \circ h : h \in \mathcal{H}\}$ be the set of all util-acts. Define the functional $I : \mathcal{U} \rightarrow \mathbb{R}$:

$$I(U \circ h) := \tilde{U}(h)$$

By Lemma 4, I is well-defined.

Let B° denote the set of all simple, real-valued, $\cup_t \mathcal{F}_t$ -measurable functions on Ω . For any t , let B_t° denote the set of simple \mathcal{F}_t -measurable functions from Ω into the closed interval $[-\beta^t, \beta^t]$.

Lemma 7 *For every $t \geq 0$, $B_t^\circ \subset \mathcal{U}$. In particular, \mathcal{U} is an absorbing subset of B° .*

Proof. Let $a = \sum_{i=1}^k \alpha_i \mathbf{1}_{A_i}$ be the canonical representation of $a \in B_t^\circ$ and note that, for every $i = 1, \dots, k$, $|\frac{\alpha_i}{\beta^t}| \leq 1$. Let x_i be such that $u(x_i) = \frac{\alpha_i}{\beta^t}$. Define $f(\omega) = x_i$ for every $\omega \in A_i$, $i \in \{1, \dots, k\}$. Define an act h by setting $h_\tau := x^*$ for all $\tau < t$ and $h_\tau = f$ for all $\tau \geq t$. It is easy to check that $U \circ h = a$.

To see that \mathcal{U} is absorbing, fix any $a \in B^\circ$. Because a is simple, it must be \mathcal{F}_t -measurable for some t . Moreover, for k large enough, $\beta^k a(\omega) \in [-\beta^t, \beta^t]$ for every $\omega \in \Omega$. Conclude from the above that $\beta^k a \in \mathcal{U}$ for k large enough.

■

Lemma 8 *For every $t, k \geq 0$ and $a \in B_t^\circ$, $I(\beta^k a) = \beta^k I(a)$.*

Proof. Fix some t and $a \in B_t^\circ$. It suffices to show that the result holds for $k = 1$. By Lemma 7, there exists an act $h \in \mathcal{H}$ such that $U \circ h = a$. Let $d_h \in X^\infty$ be such that $d_h \sim h$. Consider the act (x^*, h) and note that $U \circ (x^*, h) = \beta a$. Moreover, by Stationarity, $(x^*, h) \sim (x^*, d_h)$. Conclude that:

$$I(\beta a) = \tilde{U}(x^*, h) = U(x^*, d_h) = \beta U(d_h) = \beta \tilde{U}(h) = \beta I(U \circ h) = \beta I(a),$$

as desired. ■

Lemma 9 *The functional $I : \mathcal{U} \rightarrow \mathbb{R}$ can be uniquely extended to a functional $\tilde{I} : B^\circ \rightarrow \mathbb{R}$ such that $\tilde{I}(\beta a) = \beta a$ for every $a \in B^\circ$.*

Proof. Fix any $a \in B^\circ$ and let k, t be such that $\beta^k a \in B_t^\circ$. Define,

$$\tilde{I}(a) := \frac{1}{\beta^k} I(\beta^k a) \quad (6.3)$$

To see that \tilde{I} is well-defined, let (k', t') be another pair such that $\beta^{k'} a \in B_{t'}^\circ$. If $k' \geq k$, then $\beta^{k'} a \in B_t^\circ$, If $k > k'$, then $\beta^k a \in B_{t'}^\circ$. It is therefore without loss of generality to assume that $t = t'$, $k' > k$ and $\beta^{k'} a, \beta^k a \in B_t^\circ$. Then, Lemma 8 implies that:

$$\frac{1}{\beta^{k'}} I(\beta^{k'} a) = \frac{1}{\beta^{k'}} I(\beta^{k'-k} \beta^k a) = \frac{1}{\beta^{k'}} \beta^{k'-k} I(\beta^k a) = \frac{1}{\beta^k} I(\beta^k a),$$

and hence that \tilde{I} is well-defined. By construction, it is clear that $\tilde{I}(\beta a) = \beta \tilde{I}(a)$ for every $a \in B^\circ$, and moreover, that \tilde{I} is the unique extension of I that has the latter property. ■

The above property of \tilde{I} may be termed β -**homogeneity**. Next, we establish that the functional \tilde{I} is **translation-invariant**. Given any real number α , let $\alpha^* : \Omega \rightarrow \mathbb{R}$ be the constant function such that $\alpha^*(\omega) = \alpha$ for every state $\omega \in \Omega$.

Lemma 10 *For every $a \in B^\circ, \alpha \in \mathbb{R}$, $\tilde{I}(a + \alpha^*) = \tilde{I}(a) + \alpha$.*

Proof. Fix any $a \in B^\circ$ and $\alpha \in \mathbb{R}$. By Lemma 9, it is enough to show that $\tilde{I}(\beta^k a + \beta^k \alpha^*) = \tilde{I}(\beta^k a) + \beta^k \alpha$ for some $k > 0$. Because a is simple, it is \mathcal{F}_t -measurable for some t . Choose k large enough so that $\beta^k a(\omega) \in [-\beta^{2t}, \beta^{2t}]$ for every ω , and $\beta^k \alpha \in [\beta^t - 1, 1 - \beta^t]$. The first inclusion allows us to find an act $h \in \mathcal{H}$ such that $U \circ g = \beta^k a$ where $g_\tau = x^*$ for all $\tau < t$ and $g_\tau = h_{\tau-t}$ for all $\tau \geq t$. In more suggestive notation, we can write the act g as $(x_0^*, \dots, x_{t-1}^*, h)$. The construction is analogous to the one used in the proof of Lemma 7. The second inclusion allows us to find $(x_0, \dots, x_{t-1}) \in X^t$ such that:

$$(1 - \beta) \sum_{\tau=0}^{t-1} u(x_\tau) = \beta^k \alpha$$

Consider the act (x_0, \dots, x_{t-1}, h) and notice that $U \circ (x_0, \dots, x_{t-1}, h) = \beta^k a + \beta^k \alpha^*$. If d_h is such that $d_h \sim h$, Stationarity implies that:

$$\begin{aligned} (x_0, \dots, x_{t-1}, h) &\sim (x_0, \dots, x_{t-1}, d_h), \quad \text{and} \\ g = (x_0^*, \dots, x_{t-1}^*, h) &\sim (x_0^*, \dots, x_{t-1}^*, d_h). \end{aligned}$$

Conclude that:

$$\begin{aligned} I(\beta^k a + \beta^k \alpha^*) &= U(x_0, \dots, x_{t-1}, d_h) = \\ &= U(x_0^*, \dots, x_{t-1}^*, d_h) + \beta^k \alpha = \tilde{U}(g) + \beta^k \alpha = I(\beta^k a) + \beta^k \alpha. \end{aligned}$$

Since \tilde{I} is an extension of I , the above completes the proof. ■

The next lemma shows that the functional \tilde{I} is **superadditive**.

Lemma 11 *For every $a, b \in B^\circ$, $\tilde{I}(a + b) \geq \tilde{I}(a) + \tilde{I}(b)$.*

Proof. Fix any $a, b \in B^\circ$. By Lemma 9, it is enough to show that $\tilde{I}(\beta^k a + \beta^k b) \geq \tilde{I}(\beta^k a) + \tilde{I}(\beta^k b)$ for some k . Since a, b are simple, they are \mathcal{F}_t -measurable for some t . Choose k large enough so that $\beta^k a, \beta^k b$ take values in the interval $[-\beta^{t+1}(1 - \beta), \beta^{t+1}(1 - \beta)]$. This allows us to find functions $h_t, h_{t+1} : \Omega \rightarrow X$ such that $U \circ (\mathbf{x}_{-t}^*, h_t) = \beta^k a$ and $U \circ (\mathbf{x}_{-(t+1)}^*, h_{t+1}) = \beta^k b$. Let $z_a, z_b \in X$ be such that $(\mathbf{x}_{-t}^*, h_t) \sim (\mathbf{x}_{-t}^*, z_a)$ and $(\mathbf{x}_{-(t+1)}^*, h_{t+1}) \sim (\mathbf{x}_{-(t+1)}^*, z_b)$. Their existence follows easily from Continuity and Lemma 4. Recall that \mathbf{x}^* is the deterministic act which gives x^* in every period. Intertemporal Hedging implies that:

$$\begin{aligned} I(\beta^k a + \beta^k b) &= U(\mathbf{x}_{-t, (t+1)}^*, h_t, h_{t+1}) \geq U(\mathbf{x}_{-t, (t+1)}^*, z_a, z_b) = \\ &= U(\mathbf{x}_{-t}^*, z_a) + U(\mathbf{x}_{-(t+1)}^*, z_b) = I(\beta^k a) + I(\beta^k b). \end{aligned}$$

Since \tilde{I} is an extension of I , the above completes the proof. ■

The next lemma is the main step of the argument.

Lemma 12 *The functional \tilde{I} is positively homogenous, that is, $\tilde{I}(\alpha a) = \alpha \tilde{I}(a)$ for every $a \in B^\circ$ and $\alpha \in \mathbb{R}_+$.*

Proof. By translation-invariance, it suffices to show that $\tilde{I}(\alpha a) = 0$ for all $\alpha > 0$ and all a such that $\tilde{I}(a) = 0$. Suppose by way of contradiction that $\lambda := \tilde{I}(\alpha a) \neq 0$. By translation-invariance, again, it is without loss of generality to assume that $\lambda < 0$. Else, set $b := \alpha a - \lambda^*$ and note that $\tilde{I}(b) = 0$ whereas $\tilde{I}(\frac{1}{\alpha}b) = -\frac{\lambda}{\alpha} < 0$. Moreover, by β -homogeneity, it is also without loss of generality to assume that $\alpha < 1$.

So let $\alpha \in (0, 1)$ and let a be such that $\tilde{I}(a) = 0$ and $\lambda := \tilde{I}(\alpha a) < 0$. Define $b := \alpha a - \lambda^*$ and note that $\tilde{I}(b) = 0$. Let A be the positive ray through a :

$$A := \{\gamma a : \gamma \geq 0\} \quad (6.4)$$

Let V be the span of a, b and let $B := \{b' \in V : b \gg b'\}$ where $b \gg b'$ if $b(\omega) > b'(\omega)$ for every ω . Because a , and hence b , are simple, the set B is open in V . Moreover, for every $b' \in B$, there exists a $\gamma > 0$, sufficiently small, such that $b \gg b' + \gamma^* \gg b'$. By translation-invariance and monotonicity, conclude that:

$$0 = \tilde{I}(b) > \tilde{I}(b') \quad (6.5)$$

for all $b' \in B$. By construction, αa belongs to $A \cap B$. Since A is a ray and B is convex and open, $A \cap B$ is an open line segment. Let (γ_1, γ_2) be the open interval such that $\gamma a \in A \cap B$ for every $\gamma \in (\gamma_1, \gamma_2)$. If $\gamma_1 = 0$, then $\beta^k a \in A \cap B$, for k sufficiently large. By β -homogeneity, $\tilde{I}(\beta^k a) = \beta^k \tilde{I}(a) = 0 = \tilde{I}(b)$, for every k , which contradicts (6.5). So suppose that $\gamma_1 > 0$ and fix any k such that $\gamma_1 > \beta^k > 0$. By superadditivity and β -homogeneity in turn,

$$\tilde{I}(\beta^k a + \beta^t b) \geq \tilde{I}(\beta^k a) + \tilde{I}(\beta^t b) = 0, \quad (6.6)$$

for every t . Applying superadditivity inductively, conclude that:

$$\tilde{I}(\beta^k a + n\beta^t b) \geq 0 \quad (6.7)$$

for every $n, t \in \mathbb{N}$. Let

$$A' := \{\beta^k a + \gamma b : \gamma \geq 0\} \quad (6.8)$$

be the ray through $\beta^k a$ parallel to b . By construction $\alpha > \gamma_1 > \beta^k$ which implies that $\frac{\beta^k}{\alpha} \in (0, 1)$. Recall that $b = \alpha a - \lambda^* \gg \alpha a$, since $\lambda < 0$. Conclude that

$$b \gg \frac{\beta^k}{\alpha}(\alpha a) + (1 - \frac{\beta^k}{\alpha})b \quad (6.9)$$

and so $A' \cap B$ is a nonempty open subset of A' . It is not difficult to see that

$$\{\beta^k a + n\beta^t b : n \in \mathbb{N}, t \in \mathbb{N}\} \quad (6.10)$$

is dense in the A' . Conclude that $b \gg \beta^k a + n\beta^t b$ for some n, t . But then

$$0 = \tilde{I}(b) > \tilde{I}(\beta^k a + n\beta^t b) \quad (6.11)$$

contradicting (6.7). ■

We have verified that the functional \tilde{I} satisfies all the properties necessary to apply Theorem 1 in Gilboa and Schmeidler [13], which delivers the desired representation.

6.2 Proof of Results in Section 5

This proof that conditional preferences inherit the properties of \succeq_0 is long but elementary. It is omitted.

Proof of Lemma 3. Let $t > 0$. Dynamic Consistency, Consequentialism and the monotonicity of \succeq_0 imply that:

$$\begin{aligned} (x_0, \dots, x_{t-1}, d) \succeq_0 (x_0, \dots, x_{t-1}, d') &\Rightarrow \\ (x_0, \dots, x_{t-1}, d) \succeq_{t,\omega} (x_0, \dots, x_{t-1}, d'), & \end{aligned}$$

for every $\omega, x_0, \dots, x_{t-1}$ and d, d' . In words, all conditional preferences $\succeq_{t,\omega}$ ‘agree weekly’ with the ex ante preference \succeq_0 . In particular, there are no distinct ω, ω' and $d, d' \in X^\infty$ such that $d \succ_{t,\omega} d'$ and $d' \succ_{t,\omega'} d$. If we can show the stronger result that the preferences $\succeq_{t,\omega}, \succeq_{t,\omega'}$ are identical on X^∞ for all ω, ω' , the proof of the lemma will be complete. To see this, note that Dynamic Consistency and Consequentialism would then imply the converse of the above implication:

$$\begin{aligned} (x_0, \dots, x_{t-1}, d) \succeq_{t,\omega} (x_0, \dots, x_{t-1}, d') &\Rightarrow \\ (x_0, \dots, x_{t-1}, d) \succeq_0 (x_0, \dots, x_{t-1}, d'). & \end{aligned}$$

for every $\omega, x_0, \dots, x_{t-1}$ and d, d' .

In general, however, the fact that there are no d, d' such that $d \succ_{t,\omega} d'$ and $d' \succ_{t,\omega'} d$ is not enough to imply that $\succeq_{t,\omega}, \succeq_{t,\omega'}$ are identical on X^∞ . To see

this, suppose that $\succeq_{t,\omega}$ has ‘thick’ indifference curves. Then, $\succeq_{t,\omega'}$ may ‘split’ those indifference curves without the indifference maps of the two preferences ever being transversal. Such cases are arguably nongeneric and, as the next lemma shows, they can be ruled out whenever $\succeq_{t,\omega}, \succeq_{t,\omega'}$ are both stationary, nondegenerate and continuous on the subdomain of deterministic acts X^∞ . [The result is somewhat surprising, since we are not assuming any form of local nonsatiation.]

Lemma 13 *Let \succeq, \succ' be two continuous, stationary, nontrivial preferences on X^∞ . If the preferences are distinct, there exist alternatives $d, d' \in X^\infty$ such that $d \succ d'$ and $d' \succ' d$.*

Proof. Say that $d \in X^\infty$ is **repeating** if there exists an $n > 0$ and a vector $a \in X^n$ such that $d = (a, a, a, \dots)$. For example, (x, x', x, x', \dots) is repeating with $a := (x, x') \in X^2$. The subset X^{rep} of all repeating d 's is dense in the product topology on X^∞ . Since X^∞ is connected and \succeq, \succ' are continuous, the two preferences are distinct if and only if their restrictions to X^{rep} are distinct. Moreover, the preferences are nontrivial if and only if the restrictions are nontrivial. By the first observation, we can find $d, d' \in X^{\text{rep}}$ such that $d \succ d'$ and $d' \succeq' d$. If $d' \succ' d$, the proof is complete, so suppose that $d' \sim' d$. Because the restriction of \succ' to X^{rep} is nontrivial, there exists some $d'' \in X^{\text{rep}}$ such that $d'' \succ' d' \sim' d$. Suppose that $d'' \succ' d' \sim' d$, the opposite case follows analogously. If $d \succ d''$, the proof is complete, so suppose that $d'' \succeq d \succ d'$. Because d' is repeating, then $d' = (a, a, \dots)$ for some $a \in X^n$, and some n . Define $d^1 := (a, d'')$, $d^2 := (a, a, d'')$, and so on. By construction, the sequence $\{d^k\}$ converges to d' . Since $d \succ d'$, it follows that, for k large enough,

$$d \succ d^k. \tag{6.12}$$

On the other hand, it follows from $d'' \succ' d'$ and the stationarity of \succ' that

$$d^k := (a, \dots, a, d'') \succ' (a, \dots, a, d') = d' \sim' d \tag{6.13}$$

for every k . Conclude that, for k large enough, $d \succ d^k$ and $d^k \succ' d$, as desired. ■

References

- [1] S. Alon and D. Schmeidler. Purely subjective maxmin expected utility. Technical report, 2008.
- [2] E. Anderson, L. Hansen, and T. Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123, 2003.
- [3] F. Anscombe and R. Aumann. A definition of subjective probability. *The Annals of Mathematical Statistics*, 34:199–205, 1963.
- [4] R. Casadesus-Masanell, P. Klibanoff, and E. Ozdenoren. Maxmin expected utility over Savage acts with a set of priors. *Journal of Economic Theory*, 92(1):35–65, 2000.
- [5] S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and L. Montrucchio. Uncertainty averse preferences. *Journal of Economic Theory*, 146:1275–1330, 2011.
- [6] G. Debreu. Topological methods in cardinal utility theory. In *Mathematical Methods in the Social Sciences*. Stanford University Press, 1960.
- [7] L. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113:1–31, 2003.
- [8] L. Epstein and M. Schneider. Ambiguity and asset markets. *Annual Review of Financial Economics*, 2:315–346, 2010.
- [9] L. Epstein and S. Tanny. Increasing generalized correlation: a definition and some economic consequences. *Canadian Journal of Economics*, pages 16–34, 1980.
- [10] L. Epstein and S. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.
- [11] L. G. Epstein and T. Wang. Beliefs about beliefs without probabilities. *Econometrica*, 64(6):1343–1373, 1996.
- [12] P. Ghirardato, F. Maccheroni, M. Marinacci, and M. Siniscalchi. A subjective spin on roulette wheels. *Econometrica*, 71:1897–1908, 2003.

- [13] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.
- [14] W. M. Gorman. The structure of utility functions. *Review of Economic Studies*, 35:367–390, 1968.
- [15] W. M. Gorman. Facing an uncertain future. In *Separability and Aggregation*. Clarendon Press, 1995.
- [16] A. Kochov. *Two Essays in Decision Theory*. PhD thesis, University of Rochester, 2010.
- [17] T. Koopmans. Stationary ordinal utility and impatience. *Econometrica*, 28(2):287–309, 1960.
- [18] T. Koopmans. Representation of preference orderings over time. In C. McGuire and R. Radner, editors, *Decision and organization*, pages 79–100. Amsterdam: North Holland, 1972.
- [19] D. Kreps and E. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978.
- [20] F. Maccheroni, M. Marinacci, and A. Rustichini. Ambiguity aversion, robustness and the variational representation of preferences. *Econometrica*, 74:1447–1498, 2006.
- [21] F. Maccheroni, M. Marinacci, and A. Rustichini. Dynamic variational preferences. *Journal of Economic Theory*, 128(1):4–44, 2006.
- [22] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.
- [23] L. Savage. *The Foundations of Statistics*. Wiley, New York, NY, 1954.
- [24] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57:571–587, 1989.
- [25] T. Strzalecki. Temporal resolution of uncertainty and recursive models of ambiguity aversion. 2011. Working paper.
- [26] T. Strzalecki. Axiomatic foundations of multiplier preferences. *Econometrica*, 79(1):47–73, 2011.