

ON DYNAMIC CONSISTENCY IN AMBIGUOUS GAMES

ANDREW ELLIS

ABSTRACT. I consider static, incomplete information games where players may not be ambiguity neutral. Every player is one of a finite set of types, and each knows her own type but not that of the other players. Ex ante, players differ only in their taste for outcomes. If every player is dynamically consistent with respect to her own information structure and respects Consequentialism, then players act as if expected utility for uncertainty about types.

1. OBJECTIVES

Recently, the theory of incomplete information games (Harsanyi, 1967-8) has been extended to allow players to exhibit uncertainty averse behavior and applied to economic settings such as auctions, mechanism design and voting. This paper formalizes a modeling trade-off for such games, showing that uncertainty aversion poses difficulties for some of expected utility's particularly appealing properties, including dynamic consistency (DC). The following three properties imply no player is ambiguity averse with respect to the uncertainty about her own and other players' types: (i) DC, i.e. a strategy optimal conditional on each of a given player's types is also ex ante optimal for that player; (ii) Consequentialism, i.e. each player updates her preferences in a manner independent of other players' strategy choices and ignores counter-factual signal realizations; and (iii) common ex ante behavior, i.e. each player has the same preferences (up to tastes) before observing her type, a generalization of common priors.

Date: April, 2018.

Key words and phrases. Ambiguity, incomplete information, dynamic consistency, strategic interaction.

Department of Economics, London School of Economics. I would like to thank the referees, Larry Epstein, Sujoy Mukerji, and participants at SAET 2016 and the 2nd "Ambiguity in Games and Mechanisms" workshop for helpful comments, as well as Faruk Gul and a referee of another paper for comments that led to the results presented herein.

The main result is general, in that it applies to all standard models incorporating uncertainty, and has several immediate implications for mechanism design under ambiguity. While many of these games do not explicitly model an ex ante stage, their interpretation typically admits an ex ante stage with (at least) a participation decision. The result suggests, for instance, that satisfaction of the ex post participation constraints need not imply ex ante participation constraint. Furthermore, renegotiation may be part of the optimal contract, even when contracts are complete. In games whose interpretation does not admit any ex ante decisions, DC is particularly important for welfare comparisons. Violations of DC imply ex post welfare may differ substantially from ex ante.

The result, presented formally in Section 2, implies that much of the literature that focuses on the strategic interaction of ambiguity averse agents with incomplete information studies agents that violate one of the three properties. For instance, Theorem 1 shows that at least one of these properties fails in (discretized versions of) nearly all of the literature on auctions and multi-agent mechanism design with ambiguity aversion, including Salo and Weber (1995), Lo (1998), Bose et al. (2006), Chen et al. (2007), Bose and Daripa (2009), Bodoh-Creed (2012), Bose and Renou (2014), Ellis (2016), and Renou (2015).¹ In other contexts, various authors have explored agents that explicitly violate either DC or Consequentialism; for instance, Baliga et al. (2013) study polarization and ambiguity by maintaining DC but relaxing Consequentialism. The results also provide restrictions on the ex ante behavior of agents in non-strategic interactions with uncertainty, such as the general equilibrium model of Condie and Ganguli (2011).

While it is well-known that requiring DC for all information structures poses difficulties for ambiguity sensitive models, the approach pioneered by Sarin and Wakker (1998) and Epstein and Schneider (2003b) requires DC to hold only for a particular information structure.² In a single player decision problem with fixed information,

¹While non-discretized versions of the above papers are not covered by the result, it is easy to see that DC fails for the common sets of priors implicitly or explicitly assumed for the agents in those papers as well. The sets of priors in these papers have a product structure that fails rectangularity, as noted by Epstein and Schneider (2003a).

²See e.g. Epstein and Le Breton (1993), Ghirardato (2002), and the debate between Al-Najjar and Weinstein (2009) and Siniscalchi (2009).

they provide conditions that permit DC as well as a non-trivial role for ambiguity. The fixed information assumption also applies in the strategic settings considered here, but the results herein show that the conditions are much more restrictive when combined with the analog of common priors.

Section 3 discusses the assumptions and alternative approaches. The trade-off is tight, in that there are examples that satisfy any two of the three key assumptions. Thus, given common ex ante behavior, the modeler faces the familiar trade-off between Consequentialism and DC. On the one hand, DC can be maintained by allowing the update rule to depend on the equilibrium strategy profile, e.g. Hanany and Klibanoff (2007). Hanany et al. (2016) develop an equilibrium concept based on this idea for games when players have smooth ambiguity preferences (Klibanoff et al., 2005), for which there is no issue with ex ante versus ex post for choice, equilibrium, or welfare. In the games considered by the present paper, their solution concept reduces to solving the (ex ante) normal form.³

On the other hand, while DC has very strong normative appeal, violations thereof are well documented. The experimental literature has provided some insights into how conflicts between Consequentialism and DC are resolved. For instance, Maher and Kashima (1997), Cohen et al. (2000), and Dominiak et al. (2012) study dynamic behavior in an Ellsberg-style experiment, albeit in non-strategic settings.⁴ All find support for non-DC update rules. The last of the three explicitly tests DC against Consequentialism in an experiment where the two make conflicting predictions. They find that more subjects satisfy Consequentialism than DC.

Common priors itself is an ubiquitous, if sometimes controversial, assumption in incomplete information games. It requires that differences in beliefs are entirely due to information. A number of classic results in information economics depend critically on common priors, so understanding the corresponding assumption's implications is of interest. A drawback of this paper's results is that its analog, common ex ante behavior, implies that both perception of and attitude towards uncertainty

³Their setting is richer in that they allow sequential moves, and in these cases, the reduction to normal form need not obtain. In particular, they consider the appropriate analog to sequential equilibrium for this setting.

⁴Recently, Li et al. (2017) find evidence for non-expected utility behavior in strategic settings, though in a setting with complete information.

are identical. Justifications for common priors typically apply to the former but not the latter. A full analysis of the best way to relax common ex ante behavior as above is beyond the scope of this paper, but in Appendix B an extension of the main result to the α -maxmin expected utility model is provided.

2. MODEL AND RESULTS

Following Harsanyi (1967-8), I model incomplete information games as follows; for related formulations with ambiguity aversion, see e.g. Lo (1999), Kajii and Ui (2005), or Azrieli and Teper (2011). Let $I = \{1, \dots, n\}$ with $n \geq 2$ be the set of players, T_0 be a finite set of states of nature, T_i be a finite set of types of player i where $\#T_i \geq 2$ for all $i > 0$, and $T = T_0 \times T_1 \times \dots \times T_n$ be the set of states of the world. As is standard, denote $T_{-i} = T_0 \times \dots \times T_{i-1} \times T_{i+1} \times \dots \times T_n$ and (t_i, t'_{-i}) for the state where Player i 's type is t_i and the type of Player $j \neq i$ is t'_j . The space of outcomes of the game is a convex set X , such as the set of lotteries over action profiles. It will sometimes be convenient to write ΔZ for the set of all finite support probability measures on a set Z with an appropriate σ -algebra. As standard, Player i learns her type before choosing a strategy.⁵

I abstract away from the formal details of the game and equilibrium. Instead, I consider each player's preference over acts that map T to X . Every pair of a strategy profile and a mechanism or game with the above information structure maps to one such an act. Denote player i 's ex ante preference over acts by \succeq_0^i and her preference conditional on learning that her type is $t_i \in T_i$ by $\succeq_{t_i}^i$. Each $\succeq_{t_i}^i$ and \succeq_0^i are complete and transitive, and \succeq_0^i has a utility representation $U_0^i(\cdot)$.

2.1. Assumptions. I impose three assumptions on each player's family of preference relations, and one assumption that links preferences across players. Section 3 discusses their importance for the result, their appeal in a game theoretic setting, and some reasons why one might relax each.

Assumption 1 (Consequentialism). For every $i \in I$ and every $t_i \in T_i$, if $f(t_i, t_{-i}) = g(t_i, t_{-i})$ for all $t_{-i} \in T_{-i}$, then $f \sim_{t_i}^i g$.

⁵Formally, the information of Player i is modeled as a filtration $\mathcal{F}^i = \{\mathcal{F}_0, \mathcal{F}_1^i\}$ where $\mathcal{F}_0 = \{T\}$ and $\mathcal{F}_1^i = \{T_0 \times T_1 \times \dots \times T_{i-1} \times \{t\} \times T_{i+1} \times \dots \times T_n : t \in T_i\}$.

Assumption 2 (Dynamic Consistency). For every $i \in I$, if $f \succeq_{t_i}^i g$ for all $t_i \in T_i$, then $f \succeq_0^i g$. If in addition there exists $t'_i \in T_i$ such that $f \succ_{t'_i}^i g$, then $f \succ_0^i g$.

Assumption 3 (Full Support). For all $i \in I$, any $f \in \mathcal{F}$ and any $t \in T$, there exist outcomes $x, y \in X$ such that $f_x \succ_0^i f_y$ where $f_w(t')$ equals w when $t' = t$ and equals $f(t')$ otherwise for $w = x, y$.

The first two assumptions are adapted from Epstein and Schneider (2003b), and the third is ubiquitous in the work on updating to avoid conditioning on a null event. Consequentialism implies that each player cares only about the outcomes in states that agree with her information.⁶ DC says that if player i knows she will prefer f to g regardless of the signal she receives, then she prefers f to g ex ante as well. In contrast to Ghirardato (2002) or Epstein and Le Breton (1993), this version of DC applies only to the particular information structure in the game considered. By themselves, these assumptions do not restrict the perception of ambiguity perceived by player i about her own type. In fact, the results of Epstein and Schneider (2003b) imply that for any set of marginal probability distributions over T_i , there exists a DM satisfying the above assumptions for set of priors with that set of marginals.

Our next assumption explicitly uses the structure imposed by the game.

Assumption 4 (Common Ex Ante Behavior). There exists an interval $B \subseteq \mathbb{R}$, a continuous $U_0 : B^T \rightarrow \mathbb{R}$, and a family of continuous, onto functions $\{u_i : X \rightarrow B\}_{i=1}^n$ so that $U_0^i(f) = U_0(u_i \circ f)$ for all i and all f .

This assumption generalizes the common prior assumption typical of expected utility. In fact, whenever $U_0(\cdot)$ is an expected utility function, Common Ex Ante Behavior holds if and only if the priors are identical and the range of each u_i is the same. The normalization that the range of every u_i is the same is harmless if U_0 is maxmin expected utility, but may entail some loss of generality with other models of ambiguity aversion. Note also that every paper in mechanism design or auction theory cited above has quasi-linear utility, implying that the range of each $u_i(\cdot)$ equals \mathbb{R} .

⁶Because of this assumption, one could have defined ex post preference on acts that depend only on the states consistent with received information. I opt not to do so for notational ease.

It is easy to verify that Common Ex Ante Behavior allows for players whose ex ante preference accommodates all standard models of ambiguity aversion, including:

- $U_0(\psi) = \int \psi dp$ (Savage, 1954)
- $U_0(\psi) = \min_{p \in C} \int \psi dp$ (Gilboa and Schmeidler, 1989, henceforth, MEU)
- $U_0(\psi) = \alpha \min_{p \in C} \int \psi dp + (1 - \alpha) \max_{p \in C} \int \psi dp$ (henceforth, (α, C) -MEU)
- $U_0(\psi) = \int \phi(\int \psi dp) \mu(dp)$, $\mu \in \Delta(\Delta T)$ (Klibanoff et al., 2005).

In any of these models, one can easily state Common Ex Ante Behavior solely in terms of the preference relation.

2.2. Results. The main result shows that under the assumptions, players do not exhibit Ellsberg behavior – that is, violations of Savage (1954)’s sure thing principle – when the actions considered depend only on the type profile. To state it, say that an act f does not depend on Nature’s type if f is $f(\tau_0, t_{-i}) = f(\tau'_0, t_{-i})$ for all $\tau_0, \tau'_0 \in T_0$ and $t_{-i} \in T_{-i}$.⁷ Then, the type of every player suffices to determine its outcome. The result shows that the DM acts as if expected utility over these acts.

Theorem 1. *Under Assumptions 1-4, for uncertainty about players’ types, each player is additive. That is, for any $i \in I$ there are continuous functions $\{v_t\}_{t \in T}$ such that for any acts f, g that do not depend on Nature’s type:*

$$f \succeq_0^i g \iff \sum_{t \in T} v_t(f(t)) \geq \sum_{t \in T} v_t(g(t)).$$

Additivity immediately rules out violations of the sure thing principle and other anomalies for expected utility. Hence, DC and common ex ante behavior greatly limit the scope for ambiguity in games. No ambiguity is perceived about the types of players, only about Nature’s type.

To provide an intuition, the following specializes the result to the case of MEU with a common set of priors.

Corollary 1. *If players satisfy Assumptions 1-4 and U_0 is MEU with set of priors Π , then for any $t_{-0} \in T_{-0}$ and any $\pi, \pi' \in \Pi$, $\pi(T_0 \times \{t_{-0}\}) = \pi'(T_0 \times \{t_{-0}\})$.*

As is well known, an MEU player satisfies DC and Consequentialism only if her set of priors is rectangular (Epstein and Schneider, 2003b) with respect to her own signal.

⁷Equivalently, f is measurable with respect to the algebra $\{\emptyset, T_0\} \times \prod_{i=1}^n 2^{T_i}$.

Since all players have the same set of priors, the set must in fact be rectangular with respect to every player's signal. This is a demanding condition that forces the set of priors to collapse to a singleton set over the players' types. In fact, the priors have a very particular structure: there exists a probability measure $q \in \Delta T_{-0}$ and a closed, convex $\Pi_{t_{-0}} \subset \Delta(T_0 \times \{t_{-0}\})$ for every $t_{-0} \in T_{-0}$ so that Π equals

$$\left\{ \sum_{t_{-0} \in T_{-0}} q(t_{-0}) \pi_{t_{-0}}(\cdot) : \pi_{t_{-0}} \in \Pi_{t_{-0}} \forall t_{-0} \in T_{-0} \right\}.$$

Players have a common prior, q , about the distribution of the other players' types. Indeed, if Nature has no types, i.e. T_0 is a singleton, then each is actually (state independent) subjective expected utility.⁸

The proof relies on tools developed by Gorman (1968) that have had numerous applications in economics.⁹ As noted in Epstein and Seo (2011), DC with respect to a partition and Consequentialism implies that each cell has a property that Gorman calls Separability.¹⁰ Epstein and Seo show that exchangeable models have separable singletons, implying an additive representation. In the present paper, no such exchangeability exists, and singletons need not be separable when T_0 is not a singleton. Instead, a second result of Gorman is applied to show that the representation must be additive on "components" of the state space which include all sets of the form $T_0 \times \{t_{-0}\}$.

Proof of Theorem 1. Suppose Assumptions 1-4 are satisfied. By construction, U_0 satisfies Gorman (1968)'s P1 and P2. Full support implies Gorman's P4 (which implies his P3). For any $E \subset T$ and $f, g \in B^T$, define fEg by the element of B^T so that $fEg(t) = f(t)$ if $t \in E$ and $fEg(t) = g(t)$ if $t \notin E$. A set $E \subset T$ is *separable* if

⁸A similar result holds when U_0 is Choquet Expected Utility (CEU) (Schmeidler, 1989), i.e. there exist a probability measure q over T_{-0} and capacities $\nu_{t_{-0}}$ over T_0 for each $t_{-0} \in T_{-0}$ so that the capacity the capacity $\nu = \sum_{t_{-0}} q(t_{-0}) \nu_{t_{-0}}(\cdot)$. This is a further corollary not only of Theorem 1 but also of Sarin and Wakker (1998), who show that a CEU DM satisfies DC only if it is additive at the first stage, i.e. over each player's type.

⁹See e.g. Wakker (1989), Epstein and Seo (2011) and Mongin and Pivato (2015).

¹⁰To keep the paper self-contained, the appendix includes statements of Gorman's theorems used in the proof.

for any $x, y, z, z' \in B^T$,

$$U_0(xEz) \geq U_0(yEz) \iff U_0(xEz') \geq U_0(yEz').$$

Consequentialism and DC imply that every set of the form $\{t : t_i = \tau\}$ for any $\tau \in T_i$ is separable.

Say that $A, E \subseteq T$ *overlap* if they intersect and neither contains the other. Gorman's Theorem 1 states that if $A, E \subseteq T$ overlap and are separable, $A \cup E$ is separable. For each $t \in T_{-0}$, define

$$Q_t = \{\tau \in T : t_i \neq \tau_i \text{ for all } i > 0\}.$$

For any $\hat{t} \in T_{-0}$, I claim $Q_{\hat{t}}$ is separable. Fix arbitrary $\hat{t} \in T_{-0}$ and set $N = \prod_{i=1}^n (\#T_i - 1)$ and order the player-type pairs via $n : \{1, \dots, N\} \rightarrow I$ and $\tau : \{1, \dots, N\} \rightarrow \cup_{i \in I} T_i$ with the following properties: (i) for $i \in I$, $n(i) = i$; (ii) for all k , $\tau(k) \in T_{n(k)}$; and (iii) for every i , $\{\tau(k) : n(k) = i\} = T_i \setminus \hat{t}_i$. That is: (i) the first I players listed by n are all distinct, (ii) $\tau(k)$ is a possible type for player $n(k)$, and (iii) for player i , all types of i except \hat{t}_i occur exactly once.

Let $A_0 = \emptyset$ and for every $k \in \{1, \dots, N\}$, define

$$E_k = \{t : t_{n(k)} = \tau(k)\}$$

and

$$A_k = \cup_{k'=1}^k E_k = A_{k-1} \cup E_k.$$

Every A_k is separable. Clearly $A_1 = E_1$ is separable by the above. Suppose that A_{k-1} is separable for $k > 1$. If A_{k-1} and E_k overlap, then Gorman's Theorem 1 implies that A_k is separable. $A_{k-1} \cap E_k \neq \emptyset$ since it contains t^1 where $t_{n(k)}^1 = \tau(k)$ and $t_i^1 = \tau(i)$ for all $i \in I \setminus n(k)$, $t^1 \in A_{k-1} \cap E_k$ for $t_0 \in T_0$. $E_k \not\subseteq A_{k-1}$ because t^2 with $t_{n(k)}^2 = \tau(k)$ and $t_i^2 = \hat{t}_i$ for all $i \in I \setminus n(k)$ belongs to E_k but not A_{k-1} . Finally, $A_{k-1} \not\subseteq E_k$ since t^3 with $t_{n(k)}^3 = \hat{t}_{n(k)}$ and $t_i^3 = \tau(\bar{i})$ for $\bar{i} = \min I \setminus \{n(k)\}$ belongs to A_{k-1} but not E_k . Thus the sets overlap, and successive application yields that $A_N = Q_{\hat{t}}$ is separable.

Say that $A \subset T$ is a *top element* if $A \neq T$, A is separable, and the only separable $E \supset A$ without equality is T itself. Let $\mathcal{B} = \{C_1, \dots, C_m\}$ be the set of

top elements. Fix distinct $t, t' \in T_{-0}$ and let $C^t, C^{t'} \in \mathcal{B}$ contain $Q_t, Q_{t'}$ respectively. Note $Q_t \cup Q_{t'} = T$ so if $Q_{t'} \subset C^t$, then $C^t = T$, a contradiction. Hence $C^t \neq C^{t'}$. Now, $C^t \cap C^{t'} \neq \emptyset$ so, combining with $\#T_{-0} \geq 4$, $\#\mathcal{B} \geq 4$, conclude \mathcal{B} satisfies the hypothesis of Gorman's Theorem 2. Applying the result and noting $T_0 \times \{t\} = \cup\{T \setminus C_i : Q_t \subseteq C_i, C_i \in \mathcal{B}\}$ obtains the desired representation. \square

3. DISCUSSION

This section discusses the content of the assumptions within a game theoretic context and argues that all of them are important for the result to hold.

Full-support ensures all type profiles are viewed as possible, and rules out that one can reduce a multi-agent problem to a single agent one.¹¹ For instance, suppose that $T_1 = T_2 = \{R, B\}$ and that U_0 is MEU. If the set of priors equals

$$\Pi = \{\pi(RR) = p \text{ and } \pi(BB) = 1 - p : p \in [\underline{p}, \bar{p}]\}$$

then all players satisfy Assumptions 1, 2, and 4. However, each violates full support, and at the interim stage, each is certain of the type of the other.

Consequentialism relates closely to sequential reasoning in games, requiring that each player's preference is determined entirely by the histories possible given her current information. In this paper, different information sets are reached by virtue of Nature's choice as opposed to that of other players, so a player's realized type reveals nothing about the strategy played by others. As is explicit in Hanany et al. (2016) and implicit herein, players know (believe with probability 1) the strategy profile of others. Thus, each player need only update her beliefs about types of other players and the parameters of the game, not their opponents' strategies.

In dynamic incomplete information games, Consequentialism may be less appealing, especially off the equilibrium path. In these information sets, knowledge of the strategy profile is no longer assured. For instance, if a player engages in forward-induction reasoning (e.g. Kohlberg and Mertens, 1986), then she violates Consequentialism. Specifically, she uses their incentives in unrealized parts of the game tree to update her beliefs about the actions taken by her opponents after finding herself in

¹¹Of course, one can make assumptions weaker than full support that rule such behavior out.

an unexpected information set; see e.g. Battigalli and Siniscalchi (2002). Consequentialism requires that only incentives and strategies within the realized information set matter for her choice.

Hanany et al. (2016) provide a framework compatible with appropriate versions of Assumptions 2, 3, 4 and the Smooth Ambiguity Model by relaxing Consequentialism. At the interim stage, Player i maximizes a preference that depends both on her information and the equilibrium strategy profile. Specifically, the second-order belief over types is updated via the “smooth rule,” which ensures that DC holds but depends on the entire strategy profile, including the player’s ex ante expectation of her own choice in the current information set. In contrast, Consequentialism does not allow one’s previous expectation of her own choice to determine her updated second-order beliefs. If DC is critical for a particular application, then adopting their framework provides a way out of the impossibility result presented here.

DC requires that no player has an incentive to deviate from her ex ante optimal strategy upon learning her type. This is the property that permits reduction of the strategic form to the normal form. Modeling players who fail to satisfy DC requires one to make assumptions regarding their sophistication (see e.g. Siniscalchi, 2011). Moreover, the economic interpretation of many games, especially in mechanism design problems, admits at least two stages at which decisions are made. In such contexts, failure of DC has consequences for the optimal mechanism, including the aforementioned examples of the renegotiation and participation decisions. Of course, if players are completely naive or if the ex ante stage does not actually occur, then DC may not be a necessary property.

As noted, Common Ex Ante Behavior is the analog of the common priors assumption under SEU. Common priors captures the idea that in the absence of differential information, i.e. at the ex ante stage, there is no reason for the players to have different beliefs. It is a routinely invoked in incomplete information games dating to Harsanyi (1967-8) that has methodological and normative appeal. For instance, it provides modeling discipline that isolates the effects of asymmetric information, and evaluating expected welfare from a particular game without it is difficult because one must take a position on which prior to use. Nonetheless it has generated some

controversy – see Morris (1995), Aumann (1998), and Gul (1998) for arguments both for and against it.¹²

Adapted to non-expected utility, the arguments in favor imply that the perception of uncertainty should be the same for all players, but Common Ex Ante Behavior implies that *both* perception of and attitude towards uncertainty are identical. The classic arguments for common priors do not apply to the latter. While a full discussion of how best to relax it is left for future work, I analyze incomplete information games where players have (α, C) -MEU preferences in Appendix B. The (α, C) -MEU model makes a parametric separation between perception of uncertainty (the set of priors C) and attitude towards ambiguity (the index α).¹³ I relax Common Ex Ante Behavior by requiring that players have the same set of priors but allow them to have different α parameters. Under the other assumptions of Theorem 1, this implies that the common set of priors must have the same structure as in Corollary 1. That is, they have a common, single prior over the distribution of player’s types. This supports, but by no means confirms, the intuition that the trade-off between DC and Consequentialism applies when players have a common perception of uncertainty rather than only those with Common Ex Ante Behavior.

APPENDIX A. STATEMENTS OF GORMAN’S THEOREMS

P1, P2 and P4 of Gorman (1968), adapted to a preference \succsim on X^T , are that:¹⁴

P1: \succsim is complete, transitive, and continuous.

P2: X is topologically separable and arc connected.

P4: For any $t \in T$ and $x \in X^T$, there exist $y, z \in X$ s.t. $(y, x_{-t}) \succ (z, x_{-t})$.

¹²For instance, Gul argues that either the ex ante stage either is a real point in time and thus common priors has strong bite, or else is a fictitious modeling device and that what matters for analyzing the game is each player’s actions conditional on her realized type. In the latter case, what matters is where in the universal type space the player lies, and the possible types need not be consistent with any common prior.

¹³The extent to which the separation obtains has generated controversy in the literature, as there may exist distinct (α, C) and (β, D) representing the same preference. See Ghirardato et al. (2004) (henceforth, GMM) and the corrections by Eichberger et al. (2011) for discussion and interpretation, and Theorem 4.1 of Klibanoff et al. (2014) for a setting in which there is a unique set of “relevant” priors. Remark 1 discusses other models that obtain the separation.

¹⁴P3 is implied by P4 and hence omitted.

Under these conditions, there exists a continuous utility function U that represents \succsim . An event E is *separable* if for any $x, y, x', y' \in X^T$ such that $x(t) = y(t)$ and $x'(t) = y'(t)$ for all $t \notin E$, $x \succsim y \iff x' \succsim y'$. The subsets $A, A' \subset T$ *overlap* if $A \cap A' \neq \emptyset$, $A \not\subseteq A'$, and $A' \not\subseteq A$.

The clauses of the following theorems not used by the proofs are omitted.

Theorem (Gorman, 1968, Theorem 1). *Under P1, P2 and P4, if $A, A' \subseteq T$ overlap and are separable, then $A \cup A'$ and $A \cap A'$ are separable.*

For a vector $\psi \in X^T$ and $E \subset T$, let ψ_E be its projection onto X^E . Recall $A \subset T$ is a *top element* if $A \neq T$, A is separable, and the only separable $E \supset A$ without equality is T itself.

Theorem (Gorman, 1968, Theorem 2). *For $m \geq 3$, let $\mathcal{B} = \{C_1, \dots, C_m\}$ be the collection of top elements and $E_j = T \setminus C_j$ for each j . Under P1, P2 and P4, if two distinct elements of \mathcal{B} intersect, then $\{E_1, \dots, E_m\}$ is a partition of T and $U(\psi) = \sum_{j=1}^m U_j(\psi_{E_j})$, perhaps after a normalization.*

See Gorman (1968) for proofs.

APPENDIX B. RELAXING COMMON EX ANTE BEHAVIOR WITH α -MEU

In this section of the Appendix, I relax Common Ex Ante Behavior in the (α, C) -MEU model and argue that the result extends. Recall that a player has a (α_i, C) -MEU representation for $\alpha_i \in [0, 1]$ and a closed, convex $C \subseteq \Delta(T)$ if her preference is represented by

$$U_0^i(u_i \circ f) = \alpha_i \min_{p \in C} \int u_i \circ f dp + (1 - \alpha_i) \max_{p \in C} \int u_i \circ f dp.$$

Intuitively, C represents perception of ambiguity and α represents the agents attitude towards that ambiguity, with a lower α reflecting more uncertainty averse behavior.¹⁵

Formally, I relax Common Ex Ante Behavior as follows.

Assumption 5. There exist sets $C, C_t \subset \Delta T$ and an $\alpha_i \in [0, 1] \setminus \{\frac{1}{2}\}$ for every $i \in I$ such that for all $i \in I$:

¹⁵This correspondence is not exact because the representation is not unique. Lemma 1 shows that if (α, C) and (β, D) represent the same preference and $\alpha = \beta$, then $C = D$.

- (i) \succ_0^i has a (α_i, C) -MEU representation, and
- (ii) $\succ_{t_i}^i$ has a (α_i, C_{t_i}) -MEU representation for each $t_i \in T_i$.

Because C is constant across players, each perceives the same ambiguity ex ante. The posterior beliefs, C_{t_i} , must relate to C through DC. A necessary, but not sufficient, condition for DC and Consequentialism is that C_{t_i} results from prior-by-prior Bayesian updating of C (Lemma 2). However, the attitude towards ambiguity, α_i , may vary from player to player. Thus players have the same perception of ambiguity but differ in their attitude towards it.

I show that replacing Common Ex Ante Behavior with the above assumption implies existence of a common prior over players' types.

Theorem 2. *Under Assumptions 1, 2, 3 and 5, for any $t \in T_{-0}$ and any $\pi, \pi' \in C$, $\pi(T_0 \times \{t\}) = \pi'(T_0 \times \{t\})$.*

The result suggests that common perception of uncertainty, as opposed to common attitude towards ambiguity, drives the impossibility result. To prove the result, I show that if an event is separable, in the sense of Gorman, for an (α_i, C) -MEU DM, then it also is separable for a $(1, C)$ -MEU DM. The result then follows from applying Corollary 1 to the hypothetical game where all players are $(1, C)$ -MEU and so Common Ex Ante Behavior is satisfied.

Remark 1. A number of other models also feature a parametric separation between perception of and attitude towards uncertainty. Of them, Klibanoff et al. (2009) has a well-understood dynamic extension. However, its functional form depends on the information structure of the agent and cannot be generated by simply altering the index of ambiguity aversion or the second order prior. The general case of Gajdos et al. (2008), wherein a DM maps a primitive “probability-possibility” sets of priors to a subjective set of priors used to evaluate acts as in MEU, has enough degrees of freedom to allow the same probability-possibility set of priors and DC. For instance, one can take it to be the entire simplex, and then have each player map it to an appropriate rectangular set of priors. In neither case are the differences between individuals naturally interpretable solely in terms of their attitudes towards uncertainty. Other models that provide the separation either do not have well-understood dynamic extensions, such as Cerreia-Vioglio et al. (2011), or rely on a

particular state space structure, such as Klibanoff et al. (2014). As noted in the introduction, a full analysis of the best way to relax common ex ante behavior is beyond the scope of this paper.

Proof of Theorem 2. As in Theorem 1, identify acts with real vectors. To save space, for any measure p over all subsets of T and vector $f \in \mathbb{R}^T$, write $p(f)$ for $\int f dp$. For any measure p and event E define another measure p_E by $p_E(A) = p(A \cap E)$.¹⁶

I first prove two Lemmas. The first shows that fixing an $\alpha_i \neq \frac{1}{2}$ (and tastes), the resulting set of priors C is unique.

Lemma 1. *If a preference \succeq has both (α, \mathcal{C}) -MEU and (α, \mathcal{D}) -MEU representations for closed, convex $\mathcal{C}, \mathcal{D} \subseteq \Delta T$ and $\alpha \neq \frac{1}{2}$, then $\mathcal{C} = \mathcal{D}$.*

Proof. The following adapts part of the proof of Proposition 20 in GMM. It is without loss to renormalize so that both preferences have the same utility index u . If $\mathcal{C} \neq \mathcal{D}$, then one can use a separating hyperplane theorem to pick $\varphi \in \mathbb{R}^T$ such that $\bar{c} = \max_{p \in \mathcal{C}} p(\varphi) \neq \max_{p \in \mathcal{D}} p(\varphi) = \bar{d}$. Define $\min_{p \in \mathcal{C}} p(\varphi) = \underline{c}$ and $\min_{p \in \mathcal{D}} p(\varphi) = \underline{d}$. Since the two representations assign the same certainty equivalent to any act f with $u \circ f = \varphi$,

$$\alpha \underline{c} + (1 - \alpha) \bar{c} = \alpha \underline{d} + (1 - \alpha) \bar{d},$$

and also for any act f' with $u \circ f' = -\varphi$, implying that

$$-\alpha \bar{c} - (1 - \alpha) \underline{c} = -\alpha \bar{d} - (1 - \alpha) \underline{d}.$$

Combining yields that $\bar{d} + \underline{d} = \bar{c} + \underline{c}$ and $(1 - 2\alpha)[\bar{d} - \underline{d}] = (1 - 2\alpha)[\bar{c} - \underline{c}]$. Since $\alpha \neq \frac{1}{2}$, these can hold only if $\bar{d} = \bar{c}$, contradicting that $\max_{p \in \mathcal{C}} p(\varphi) \neq \max_{p \in \mathcal{D}} p(\varphi)$. \square

Now, I show that C is rectangular with respect to each player's partition. One implication of Lemma 2 is that rectangular priors is a necessary but not sufficient condition for DC and Consequentialism in the (α, C) -MEU model.

Lemma 2. *Assumptions 1-3 and 5 imply that for any $i \in I$ and $\tau_i \in T_i$, the event $\{t \in T : t_i = \tau_i\}$ is separable for \succeq' , when \succeq' has a $(1, C)$ -MEU representation.*

Proof of Lemma 2. The result is trivial if $\alpha_i \in \{0, 1\}$, so consider only $\alpha \equiv \alpha_i \neq 0, 1, \frac{1}{2}$. Fix arbitrary $i \in I$ and $\tau_i \in T_i$, and let V represent \succeq_0^i , V_{i, τ_i} represent $\succeq_{\tau_i}^i$,

¹⁶Note p_E may not be a probability measure even if p is.

and $E = \{t \in T : t_i = \tau_i\}$. Consider any $f \in \mathbb{R}_+^T$. Note $fE0 \sim_0^i V_{i,\tau_i}(f)E0$ and $fE0 \succeq_0^i xE0$ if and only if $V_{i,\tau_i}(f) \geq x$. Since $V_{i,\tau_i}(f) \geq 0$

$$\begin{aligned} V(fE0) &= \alpha \min_{p \in C} p_E(f) + (1 - \alpha) \max_{p \in C} p_E(f) \\ &= V(V_{i,\tau_i}(f)E0) = [\alpha \min_{p \in C} p(E) + (1 - \alpha) \max_{p \in C} p(E)]V_{i,\tau_i}(f) \end{aligned}$$

so setting $p^* = \alpha \min_{p \in C} p(E) + (1 - \alpha) \max_{p \in C} p(E)$ and defining

$$C_{\tau_i}^* = \left\{ \frac{p_E(\cdot)}{p^*} : p \in C \right\}$$

gives an $(\alpha, C_{\tau_i}^*)$ -MEU representation of $\succeq_{\tau_i}^i$.¹⁷ Applying Lemma 1 gives that $C_{\tau_i}^* = C_{\tau_i}$ (the argument above goes through even if there exists $\mu \in C_{\tau_i}^*$ with $\mu(T) < 1$) for each τ_i . This requires that $p(E) = p^*$ for all $p \in C$ and thus $C_{\tau_i} = \{p(\cdot|E) : p \in C\}$.

Now for an arbitrary g ,

$$\begin{aligned} V(g) &= V(V_{i,\tau_i}(g)Eg) = \alpha \min_{p \in C} [p^*V_{i,\tau_i}(g) + (1 - p^*)p(g|E^c)] \\ &\quad + (1 - \alpha) \max_{p \in C} [p^*V_{i,\tau_i}(g) + (1 - p^*)p(g|E^c)] \\ &= p^*[\alpha \min_{q \in C} q(g|E) + (1 - \alpha) \max_{q \in C} q(g|E)] + \\ &\quad + (1 - p^*)[\alpha \min_{p \in C} p(g|E^c) + (1 - \alpha) \max_{p \in C} p(g|E^c)] \\ &= \alpha \min_{p \in D} p(g) + (1 - \alpha) \max_{p \in D} p(g) \end{aligned}$$

where

$$D = \{p^*q(\cdot|E) + (1 - p^*)q'(\cdot|E^c) : q, q' \in C\}.$$

Thus \succeq_0^i has both (α, C) -MEU and (α, D) -MEU representations. Lemma 1 implies $D = C$. Applying Epstein and Schneider (2003b) yields that E is \succeq' -separable. \square

Using Lemma 2, C must be rectangular with respect to each player's filtration. One can then apply Corollary 1 to obtain the desired conclusion. \square

¹⁷This extends to any vector f , since if $V_{i,\tau_i}(f) < 0$, we can find $k \in \mathbb{R}$ so that $V_{i,\tau_i}(f + k) \geq 0$ and use the identity $V_{i,\tau_i}(f) = V_{i,\tau_i}(f + k) - k$.

REFERENCES

- Al-Najjar, Nabil I. and Jonathan Weinstein (2009), “The ambiguity aversion literature: A critical assessment.” *Economics and Philosophy*, 25, 249–284.
- Aumann, Robert J. (1998), “Common priors: A reply to gul.” *Econometrica*, 66, 929–938.
- Azrieli, Yaron and Roe Teper (2011), “Uncertainty aversion and equilibrium existence in games with incomplete information.” *Games and Economic Behavior*, 73, 310–317.
- Baliga, Sandeep, Eran Hanany, and Peter Klibanoff (2013), “Polarization and ambiguity.” *American Economic Review*, 103, 3071–83.
- Battigalli, Pierpaolo and Marciano Siniscalchi (2002), “Strong belief and forward induction reasoning.” *Journal of Economic Theory*, 106, 356 – 391.
- Bodoh-Creed, Aaron (2012), “Ambiguous beliefs and mechanism design.” *Games and Economic Behavior*, 75, 518–537.
- Bose, Subir and Arup Daripa (2009), “A dynamic mechanism and surplus extraction under ambiguity.” *Journal of Economic Theory*, 144, 2084–2114.
- Bose, Subir, Emre Ozdenoren, and Andreas Pape (2006), “Optimal auctions with ambiguity.” *Theoretical Economics*, 1, 411–438.
- Bose, Subir and Ludovic Renou (2014), “Mechanism design with ambiguous communication devices.” *Econometrica*, 82, 1853–1872.
- Cerreia-Vioglio, Simone, Paolo Ghirardato, Fabio Maccheroni, Massimo Marinacci, and Marciano Siniscalchi (2011), “Rational preferences under ambiguity.” *Economic Theory*, 48, 341–375.
- Chen, Yan, Peter Katuscak, and Emre Ozdenoren (2007), “Sealed bid auctions with ambiguity: Theory and experiments.” *Journal of Economic Theory*, 136, 513 – 535.
- Cohen, M., I. Gilboa, J. Y. Jaffray, and D. Schmeidler (2000), “An experimental study of updating ambiguous beliefs.” *Risk Decision and Policy*, 5, 123–133.
- Condie, Scott and Jayant Ganguli (2011), “Ambiguity and rational expectations equilibria.” *Review of Economic Studies*, 78, 821–845.
- Dominiak, Adam, Peter Duersch, and Jean-Philippe Lefort (2012), “A dynamic ellberg urn experiment.” *Games and Economic Behavior*, 75, 625–638.

- Eichberger, Jurgen, Simon Grant, David Kelsey, and Gleb A. Koshevoy (2011), “The α -MEU model: A comment.” *Journal of Economic Theory*, 146, 1684 – 1698.
- Ellis, Andrew (2016), “Condorcet meets Ellsberg.” *Theoretical Economics*, 11, 865–895.
- Epstein, Larry and Michel Le Breton (1993), “Dynamically consistent beliefs must be bayesian.” *Journal of Economic Theory*, 61, 1–22.
- Epstein, Larry G. and Martin Schneider (2003a), “IID: independently and indistinguishably distributed.” *Journal of Economic Theory*, 113, 32–50.
- Epstein, Larry G. and Martin Schneider (2003b), “Recursive multiple priors.” *Journal of Economic Theory*, 113, 1–31.
- Epstein, Larry G. and Kyoungwon Seo (2011), “Symmetry or dynamic consistency?” *The B.E. Journal of Theoretical Economics*, 11.
- Gajdos, T., T. Hayashi, J.-M. Tallon, and J.-C. Vergnaud (2008), “Attitude toward imprecise information.” *Journal of Economic Theory*, 140, 27–65.
- Ghirardato, Paolo (2002), “Revisiting savage in a conditional world.” *Economic Theory*, 20, 83–92.
- Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci (2004), “Differentiating ambiguity and ambiguity attitude.” *Journal of Economic Theory*, 118, 133–173.
- Gilboa, Itzhak and David Schmeidler (1989), “Maxmin expected utility with a unique set of priors.” *Journal of Mathematical Economics*, 18.
- Gorman, W. M. (1968), “The structure of utility functions.” *The Review of Economic Studies*, 35, 367–390.
- Gul, Faruk (1998), “A Comment on Aumann’s Bayesian View.” *Econometrica*, 66, 923–928.
- Hanany, Eran and Peter Klibanoff (2007), “Updating preferences with multiple priors.” *Theoretical Economics*, 2, 261–298.
- Hanany, Eran, Peter Klibanoff, and Sujoy Mukerji (2016), “Incomplete information games with smooth ambiguity preferences.” *mimeo*.
- Harsanyi, John C (1967-8), “Games with incomplete information played by Bayesian players.” *Management Science*, 14.

- Kajii, Atsushi and Takahashi Ui (2005), “Incomplete information games with multiple priors.” *Japanese Economic Review*, 56, 332–351.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2005), “A smooth model of decision making under ambiguity.” *Econometrica*, 73, 1849–1892.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji (2009), “Recursive smooth ambiguity preferences.” *Journal of Economic Theory*, 144, 930–976.
- Klibanoff, Peter, Sujoy Mukerji, and Kyoungwon Seo (2014), “Perceived ambiguity and relevant measures.” *Econometrica*, 82, 1945–1978.
- Kohlberg, Elon and Jean-Francois Mertens (1986), “On the strategic stability of equilibria.” *Econometrica*, 54, 1003–1037.
- Li, Zhihua, Graham Loomes, and Ganna Pogrebna (2017), “Attitudes to uncertainty in a strategic setting.” *The Economic Journal*, 127, 809–826.
- Lo, Kin Chung (1998), “Sealed bid auctions with uncertainty averse bidders.” *Economic Theory*, 12, 1–20.
- Lo, Kin Chung (1999), “Extensive form games with uncertainty averse players.” *Games and Economic Behavior*, 28, 256 – 270.
- Maher, Patrick and Yoshihisa Kashima (1997), “Preference reversal in ellisberg problems.” *Philosophical Studies*, 88, 187–207.
- Mongin, Philippe and Marcus Pivato (2015), “Ranking multidimensional alternatives and uncertain prospects.” *Journal of Economic Theory*, 157, 146–171.
- Morris, Stephen (1995), “The common prior assumption in economic theory.” *Economics and Philosophy*, 11, 227–253.
- Renou, Ludovic (2015), “Rent extraction and uncertainty.”
- Salo, Ahti A. and Martin Weber (1995), “Ambiguity aversion in first price sealed bid auctions.” *Journal of Risk and Uncertainty*, 11, 123–137.
- Sarin, Rakesh and Peter P. Wakker (1998), “Dynamic choice and unexpected utility.” *Journal of Risk and Uncertainty*, 17, 87–120.
- Savage, Leonard J. (1954), *The Foundations of Statistics*. Dover.
- Schmeidler, David (1989), “Subjective probability and expected utility without additivity.” *Econometrica*, 57, 571–87.
- Siniscalchi, Marciano (2009), “Two out of three ain’t bad: A comment on “the ambiguity aversion literature: A critical assessment”.” *Economics and Philosophy*,

25, 335–356.

Siniscalchi, Marciano (2011), “Dynamic choice under ambiguity.” *Theoretical Economics*, 6, 379–421.

Wakker, Peter P. (1989), *Additive Representations of Preferences: A New Foundation of Decision Analysis*. Kluwer Academic Publishers, Dordrecht.