Trembles in Extensive Games with Ambiguity Averse Players*

Gaurab Aryal and Ronald Stauber†

February 17, 2014

Abstract

We introduce and analyze three definitions of equilibrium for finite extensive games with imperfect information and ambiguity averse players. In a setting where players' preferences are represented by maxmin expected utility, as characterized in Gilboa and Schmeidler (1989), our definitions capture the intuition that players may consider the possibility of slight arbitrary mistakes. This generalizes the idea leading to trembling-hand perfect equilibrium as introduced in Selten (1975), by allowing for ambiguous trembles characterized by sets of distributions. We prove existence for two of our equilibrium notions, and relate our definitions to standard equilibrium concepts with expected utility maximizing players. Our analysis shows that ambiguity aversion can lead to behavioral implications that are distinct from those attained under expected utility maximization, even if ambiguous beliefs only arise from the possibility of slight mistakes in the implementation of unambiguous strategies.

Keywords: Extensive games; Ambiguity; Maxmin

JEL Classification: C72; D81

*Acknowledgments to be added.

†G. Aryal: University of Chicago and Australian National University, Email: `aryalg@uchicago.edu`; R. Stauber: Research School of Economics, Australian National University, ACT 0200, Australia, Email: `ronald.stauber@anu.edu.au`

1 Introduction

The main difficulty in defining equilibrium concepts for extensive games with imperfect information is how to characterize players' equilibrium beliefs about past actions at information sets that do not lie on the equilibrium path. For games where players are expected utility maximizers, the two most common ways to solve this difficulty give rise to the definitions of Weak Perfect Bayesian Equilibrium (WPBE) and Sequential Equilibrium (SE). For a WPBE, any beliefs are defined to be consistent off the equilibrium path; for SE, consistency is defined by requiring equilibrium beliefs at all information sets to be the limit of beliefs derived using Bayesian updating from a sequence of completely mixed behavioral strategies that converges to the SE strategy.

In this paper, we propose an alternative approach in an environment with ambiguity/uncertainty averse players, whose preferences are represented by maxmin expected utility, as axiomatized by Gilboa and Schmeidler (1989).[1] The intuition for our equilibrium definition is analogous to the suggestion of Selten (1975) that equilibrium play should take into account the possibility of slight mistakes. We allow for such mistakes by assuming that players believe that with a small ε probability their opponents may not be rational and may deviate from their equilibrium strategies. However, in contrast to Selten (1975), we do not view such deviations as occurring according to exogenously given probability distributions, assuming instead that if a mistake occurs, it can follow any distribution over the relevant action space. As a result, players are faced with a set of possible distributions over the actions of their opponents, which we model as an ε-contamination of the actual equilibrium strategies. Given any probability measure p defined over a set Θ and any $\varepsilon \in (0,1)$, an ε-contamination of p is defined by the set

$$p^{\varepsilon} := \{r \in \Delta\Theta \,|\, r = (1-\varepsilon)p + \varepsilon q, \text{ with } q \in \Delta\Theta\},$$

where $\Delta\Theta$ denotes the set of all probability measures with support in Θ. Thus, if say p is an equilibrium strategy of a player at some information set, his opponents believe that this player plays p with a high probability $1-\varepsilon$, and makes a mistake with probability ε. Moreover, if he makes a mistake, he may do so according to any arbitrary distribution over his action set. The players' beliefs about past actions at any informa-

[1] See also Gajdos et al. (2008) for a maxmin expected utility representation for agents who possess objective but imprecise information.

tion set are then defined using full Bayesian updating across all opponents' strategies in the ε-contamination under which the particular information set is reached with positive probability. Similarly, beliefs about future actions are defined by the ε-contamination of opponents' strategies at future information sets. If the resulting set-valued beliefs are assumed to generate the sets of "priors" corresponding to a maxmin expected utility representation of the players' preferences, the associated decision rule can be used to derive best responses at each information set, and define corresponding equilibrium notions. Hence, our equilibrium definitions implicitly require that the relevant set of probability distributions in the maxmin expected utility representation of a player's preferences is derived from an ε-contamination of his opponents' equilibrium strategies.[2] While the main part of the paper assumes that players update their beliefs using the full Bayesian updating method, many of our results also apply to more general updating rules. We discuss such generalizations after presenting our main results.

Our modeling approach has a number of advantages: The use of ε-contaminations generates ambiguous beliefs modeled by sets of distributions from an equilibrium strategy profile described by precise unambiguous distributions over action sets. Thus, equilibria are defined using standard behavioral strategy profiles, and can be characterized without having to specify an associated belief system as for WPBE or SE. Furthermore, the resulting beliefs are in a certain sense "more ambiguous" at information sets that do not lie on the equilibrium path. This has a very intuitive interpretation—if players observe a counterfactual, i.e., if they find themselves at an information set that should not have been reached according to the equilibrium strategies, their beliefs about what kind of play by their opponents led to this information set are more ambiguous than beliefs at information sets that lie on the equilibrium path. In addition, our definition relies on the standard notion of ε-contamination, which has been used extensively in Bayesian statistics to analyze robustness questions,[3] and has recently been applied in a variety economic contexts, such as in Nishimura and Ozaki (2004), Bose, Ozdenoren, and Pape (2006) and Bose and Daripa (2009).

To our knowledge, the only paper that allows for ambiguity aversion in extensive games with imperfect information is Lo (1999), who defines a notion of "multiple prior

[2] A justification for such beliefs could be inferred from the axiomatic decision-theoretic models of Nishimura and Ozaki (2006) and Kopylov (2008).

[3] See, for example, Berger (1985) or Wasserman and Kadane (1990).

Nash equilibrium" for such games. As in our model, the only uncertainty that players face in the games he considers is regarding the strategies of their opponents. In his equilibrium notion, players beliefs about their opponents are modeled using sets of probability measures over opponents' strategies, which are only updated at information sets that are in the support of a player's beliefs. Given that multiple prior Nash equilibrium is an extension of the "equilibrium in beliefs" for strategic games with uncertainty averse players (Lo, 1996), it defines a consistency notion for the players' beliefs, which does not require consideration of information sets that cannot arise according to such beliefs. This allows a definition of a Nash equilibrium for the corresponding extensive games that does not make any restrictions on "off-the-equilibrium-path" behavior, and therefore does not require any additional assumptions about counterfactual beliefs. To define consistency of beliefs corresponding to a multiple prior Nash equilibrium, the support of a player's beliefs is only allowed to include strategies of his opponents that are optimal given their own beliefs. The interpretation given is that each player knows his opponents' beliefs, and that they are rational. As a consequence, each strategy in the support of any of the players' set-valued beliefs must be optimal at every information set that is consistent with the belief sets that define the equilibrium. This requirement also yields the implication that players are ambiguous about their opponents' strategies, but have complete knowledge of their opponents' beliefs. In contrast, our equilibrium notion is defined in terms of a consistency condition on strategies and not on beliefs, and relies on players allowing for opponents who make mistakes and play sub-optimal actions. This assumption enables a careful analysis of beliefs even at information sets that lie off the equilibrium path, and a definition of optimality of strategies at such information sets.

Note that in our model players do not have the option of choosing "ambiguous strategies," i.e., they cannot use subjective randomization devices such as Ellsberg urns to determine their action choices. Bade (2011) and Riedel and Sass (2013) analyze normal form games where players can choose such ambiguous strategies. Riedel and Sass (2013) base their definition of "Ellsberg equilibrium" on maxmin expected utility, and require an ambiguous equilibrium strategy to only contain distributions over a player's actions that attain the maxmin of his expected utility, where the minimum is taken over the elements of his opponents' (ambiguous) equilibrium strategies. In a similar way, the "beliefs equilibrium" of Lo (1996) requires the ambiguous beliefs (over opponents' strategies) that define the equilibrium notion to only contain distributions that are optimal in the sense of attaining a maxmin given the opponents' own beliefs.[4] In contrast to these papers, the ambiguous beliefs in our paper are generated by trembles, which by definition are not required to be optimal. There is also a literature on incomplete information games with ambiguous beliefs about states of the world. Such Bayesian games with ambiguous beliefs are analyzed for example in Kajii and Ui (2005), Bose et al. (2006), Bose and Daripa (2009), Azrieli and Teper (2011) and Stauber (2011). In the present paper, the environment of the game is common knowledge and thus there is never any ambiguity about any states of the world—the only ambiguity present is in regards to the actions chosen by a player's opponents. A natural extension of our analysis could be achieved by considering games that include both ambiguity about states of the world and opponents' actions.

We introduce three equilibrium notions for our environment. The first, referred to as ε-Perfect Maxmin Equilibrium (ε-PME), requires equilibrium strategies β to be optimal for beliefs derived from an ε-contamination of β. The second notion defines a strategy profile to be a Perfect Maxmin Equilibrium (PME) if it is the limit of a sequence of ε-PME strategies as $\varepsilon \to 0$. Thus, the definition of a PME is analogous to Selten's definition of (trembling-hand) perfect equilibrium (Selten, 1975), which is defined as the limit of a sequence of equilibria of perturbed games. Our third equilibrium definition considers the set limit of beliefs induced by an ε-contamination of equilibrium strategies as $\varepsilon \to 0$, and requires the equilibrium strategies to be optimal given these limiting beliefs. We call such an equilibrium a Strong Perfect Maxmin Equilibrium (SPME). Our main result is that an SPME may not always exist, but that every extensive game with perfect recall always has an ε-PME and thus a PME. We also show that SPME and PME are Nash Equilibria (NE) and Subgame Perfect Equilibria (SPE) of the respective games. However, both SPME and PME are distinct from WPBE and SE when viewed as refinements of NE—strategies corresponding to WPBE or SE are not necessarily SPME or PME, and conversely, there are SPME and PME that are not part of a WPBE or SE. Given that our model considers players with ambiguity averse preferences, and the standard definitions of NE and its refinements assume expected utility maximizing players, our equilibrium

[4]The equilibrium notion of Lo (1996) also allows for disagreement and correlation in the players' beliefs about other players' strategies.

notions cannot be interpreted directly as refinements. Our analysis therefore implicitly views expected utility preferences as a subset of maxmin preferences, where the relevant beliefs are given by a singleton distribution, and only compares the respective equilibrium strategies. Based on this interpretation, our results show that ambiguity aversion may yield behavioral implications that are distinct from those derived under expected utility maximization, even when ambiguous beliefs only arise from small probability errors in the implementation of unambiguous strategies, and we consider the limiting case where the probability of such errors converges to zero. We illustrate our results using various examples.

The paper is structured as follows: A model of extensive games with imperfect information is presented in Section 2. Section 3 introduces our equilibrium definitions. Existence of equilibria is explored in Section 4, and the relation of our equilibrium definitions to standard equilibrium concepts is analyzed in Section 5. Section 6 discusses various generalizations and extensions, and Section 7 concludes.

2 Extensive games with imperfect information

We use the model and notation of Osborne and Rubinstein (1994, p. 200), with the additional assumptions that all games satisfy perfect recall, and that preferences can be represented by maxmin expected utility. A finite extensive game with imperfect information is then defined as follows:

- A finite set N representing the players of the game;
- A finite set H of sequences of actions $(a_k)_{k=1,\dots,K}$, representing the histories of the game (including the empty history $\varnothing$);
- A set $Z \subset H$ of terminal histories; for each $h \in H \setminus Z$, $A(h) := \{a \,|\, (h, a) \in H\}$ defines the set of actions available after history h;
- A player function $P : H \setminus Z \to N$;
- For each player $i \in N$ an information partition $\mathcal{I}_i$ of the set $\{h \in H \setminus Z \,|\, P(h) = i\}$, having the property that if $h, h' \in I_i \in \mathcal{I}_i$, then $A(h) = A(h')$; the player choosing an action at the information set I_i is denoted by $P(I_i)$, and the actions available at I_i by $A(I_i)$;

- All players have perfect recall, i.e., if two histories h and h' belong to the same information set I_i, then h is not a sub-history of h' and h' not a sub-history of h, and both h and h' pass through the same sequence of information sets of player i, and contain identical actions at all such information sets;

- The preferences of each player i satisfy the axioms of Gilboa and Schmeidler (1989) and can be represented by a maxmin expected utility function $u_i : Z \to \mathbb{R}$, i.e., if player i's ambiguous beliefs over terminal histories are given by a set of distributions $M \subset \Delta Z$, then his expected utility is defined by $\min_{\mu \in M} \mathbb{E}_\mu[u_i(z)]$, where $\mathbb{E}_\mu$ denotes the expectation operator associated to the distribution μ.

The set of pure strategies of player i is defined by $S_i := \times_{I_i \in \mathcal{I}_i} A(I_i)$, and the set of mixed strategies by $\mathcal{M}_i := \Delta S_i$. A behavioral strategy of player i, $\beta_i = (\beta_i(I_i))_{I_i \in \mathcal{I}_i}$, with $\beta_i(I_i) \in \Delta A(I_i)$, is a function that assigns to each information set I_i of player i a probability distribution over the set of actions available at I_i. We let β denote a profile of behavioral strategies β_i, and $\mathcal{B}_i$ denote the set of all behavioral strategies of player i, with $\mathcal{B} := \times_{i=1}^N \mathcal{B}_i$. Given a profile of strategies $\beta \in \mathcal{B}$, $\mathcal{O}(\beta)$ denotes the probability measure over terminal histories Z induced by β, and given a probability measure $\mu(I_i) \in \Delta I_i$, $\mathcal{O}(\beta, \mu(I_i))$ denotes the probability measure over Z conditional on I_i and $\mu(I_i)$. Finally, $u_i(\mathcal{O}(\beta))$ and $u_i(\mathcal{O}(\beta, \mu(I_i)))$ denote the expected utilities of the respective measures.

3 Strategies and equilibria with ambiguity averse players

As noted in the introduction, we define equilibria in terms of unambiguous strategy profiles, but endow players at each information set with ambiguous beliefs derived from their opponents' strategies. We assume that players choose behavioral strategies.[5] Given a profile of behavioral strategies, each player's ambiguous beliefs are derived from a behavioral ε-contamination of every opponent's strategy, defined as follows:

Definition 1. Let $\beta_j \in \mathcal{B}_j$ and $\varepsilon \in (0, 1)$. A *behavioral ε-contamination of β_j* is the set $\beta_j^\varepsilon \subset \mathcal{B}_j$ defined by

$$\beta_j^\varepsilon := \{\beta_j' \in \mathcal{B}_j \,|\, \beta_j'(I_j) = (1-\varepsilon)\beta_j(I_j) + \varepsilon\delta_j(I_j),\, \forall \delta_j \in \mathcal{B}_j,\, \forall I_j \in \mathcal{I}_j\}.$$

[5]We will discuss the appropriateness of this modeling approach shortly.

In deriving beliefs from behavioral ε-contaminations, we assume that from the point of view of the opponents of player j, this player is expected to make a mistake with probability ε at each of his information sets, in which case he may choose his actions according to any distribution over the actions that are feasible at the respective information sets. Since mistakes result by definition from unintentional trembles, all feasible distributions are thus viewed as equally-likely candidates for a mistake. Furthermore, the distributions over actions in case a mistake occurs are assumed to be independent across information sets.

Hence, when defining our equilibrium notions, we consider a fixed behavioral strategy profile β, but think of each player i's beliefs about his opponents' strategies as being derived from an ε-contamination β^{ε}_{-i} induced by β_{-i}. Then, given any information set I_i of player i, there always exist strategies in β^{ε}_{-i} under which this information set is reached with positive probability, as long as player i's actions are consistent with I_i. We can thus define i's beliefs at I_i over past actions as the set of all distributions over histories in I_i that are derived using full Bayesian updating across all those strategies in β^{ε}_{-i} under which I_i is reached with positive probability, assuming that player i's actions prior to I_i are consistent with I_i. Furthermore, beliefs at I_i over future actions are given by the ε-contaminations of all his opponents' relevant future strategies.

It is well-known that the maxmin expected utility model need not be dynamically consistent with full Bayesian updating.[6] Dynamic consistency is also not guaranteed to hold in the restricted environment we consider, when beliefs are derived from ε-contaminated strategies as described above. This is illustrated by the following example:

Example 1. Consider the game structure described in Figure 1, where player 2's payoffs are arbitrary, and fix a strategy for player 2 such that $\beta_2(R) = 1$. Assuming player 1 chooses a strategy at the start of the game, his available pure strategy choices are CA and CB, and his expected utility from playing CA with probability p is given by

$$\begin{aligned} u_1(p, \beta_2) = {} & (1-\varepsilon)[101p + 100(1-p)] \\ & + \varepsilon \min_{(q_O, q_L, q_R) \in \Delta\{O,L,R\}} \{-q_O + q_L(1-p)101 + q_R[101p + 100(1-p)]\}. \end{aligned}$$

[6]For a discussion and analysis of this issue in a decision-theoretic setting, see for example Epstein and Schneider (2003) and Hanany and Klibanoff (2007). See also Sass (2013) for an analysis of dynamic consistency in a game-theoretic setting where players can choose ambiguous strategies.

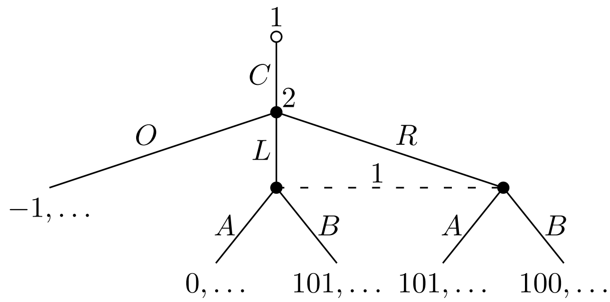

Figure 1: Example 1.

Clearly, the minimum in the second term is attained when $q_O = 1$, which implies that the optimal choice at the start of the game is to set $p = 1$, and hence to play CA.

Now consider the optimal choice of player 1 at this second information set. The histories constituting this information set are given by CL and CR, and player 1's conditional beliefs over these histories are derived from an ε-contamination of player 2's strategy $\beta_2(R) = 1$. Letting μ denote the probability assigned to CR, the set of all such conditional beliefs is given by $\mu \in [1-\varepsilon, 1]$. Then 1's expected utility from assigning probability p to action A is given by

$$\begin{aligned} u_1(p, \beta_2) &= \min_{\mu \in [1-\varepsilon,1]} \{\mu[101p + 100(1-p)] + (1-\mu)(1-p)101\} \\ &= \min_{\mu \in [1-\varepsilon,1]} \{\mu(102p - 1) + 101(1-p)\} \\ &= \begin{cases} 100 + p, & \text{if } p \leq \frac{1}{102}, \\ 100 + \varepsilon + (1 - 102\varepsilon)p & \text{if } p > \frac{1}{102}. \end{cases} \end{aligned}$$

It follows that as long as $1 - 102\varepsilon < 0$, or equivalently, $\varepsilon > \frac{1}{102}$, the optimal strategy of player 1 is to set $p = \frac{1}{102}$, so player 1's preferences are not dynamically consistent.[7] $\lhd$

Since dynamic consistency will in general not hold in our environment, in defining our equilibrium notions, we follow the approach to dynamic choice introduced by Strotz (1955-1956), usually referred to as consistent planning.[8] Assuming consistent planning,

[7]Note that dynamic inconsistency can also arise with arbitrarily small values of ε. For example, if player 2 is assumed to play O with probability 1, then for any $\varepsilon > 0$, every $p \in [0, 1]$ is optimal at the start of the game, but only $p = \frac{1}{102}$ is optimal at player 1's second information set.

[8]See also Siniscalchi (2011) for a recent axiomatic justification for consistent planning, which shows

each information set of every player is viewed as defining a separate individual who decides independently what strategy to choose at that information set. When making a choice at an information set, each corresponding individual will correctly predict the strategies chosen at future information sets by different instances of the same player. Furthermore, given that we assume perfect recall, each player always knows exactly what actions he has chosen at past information sets.

Consider now, as an alternative to introducing trembles information set-by-information set through behavioral ε-contaminations, the possibility of introducing trembles to mixed strategies in the strategic form of an extensive game. Such an approach would essentially assume that players can commit to a mixed strategy ex-ante, and that trembles occur at this commitment stage. If dynamic inconsistency arises in the game under consideration, this would contradict the dynamic structure of the game, since it would deprive the players of the option of choosing actions at individual information sets, which, in the case of ambiguity averse players, would not be without loss of generality. In contrast, when using the consistent planning approach, each information set is treated as an individual player who chooses an action given the information available at that information set. As a consequence, the fact that actions are chosen individually at each information set suggests that if mistakes can occur, they should be modeled as individual trembles that are made when actions are chosen, i.e., at individual information sets, as in the definition of a behavioral ε-contamination.[9]

Epstein and Schneider (2003) and Hanany and Klibanoff (2007) show that dynamic consistency can be guaranteed to hold in dynamic choice settings with maxmin expected utility, by either restricting the class of feasible initial priors, or carefully choosing updating rules that differ from full Bayesian updating. We discuss in Section 6.2 how these approaches might be used in our game-theoretic framework to restore dynamic consistency, but also show that even for our initial model, in a certain sense, any dynamic inconsistency associated with our equilibrium strategies disappears at the limit, when $\varepsilon \to 0$.

Before introducing our equilibrium definitions, we now characterize the ambiguous

that knowledge of ex-ante preferences over decision trees is sufficient to characterize the behavior of a consistent planner.

[9]Note that this discussion applies even to games such as the one discussed in Example 1, where each player only has one information set at which he must choose an action.

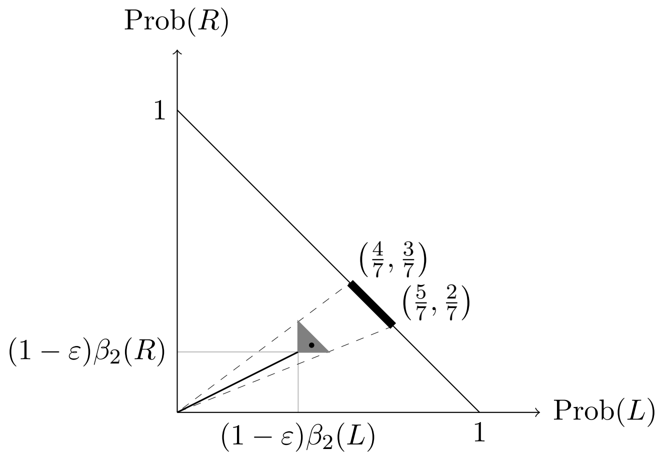

Figure 2: The derivation of $\mu_{I_i}(\beta^\varepsilon)$.

beliefs that can arise through the previously described updating procedure. For any information set I_i such that player i moves at I_i, there is a unique, but possibly empty, sequence of actions of player i that leads to I_i and is contained in every history $h \in I_i$. Given a strategy profile β, let $\beta_{-i}[I_i]$ denote the probability of reaching I_i under β_{-i}, assuming that i plays the unique action sequence that leads to I_i. If β^ε is an ε-contamination of β, let $\mu_{I_i}(\beta^\varepsilon) \subset \Delta I_i$ denote the set of probability distributions over I_i that are derived from β^ε_{-i} using full Bayesian updating across all $\beta'_{-i} \in \beta^\varepsilon_{-i}$ with $\beta'_{-i}[I_i] > 0$. Thus, every element in $\mu_{I_i}(\beta^\varepsilon)$ yields a distribution over histories in I_i, conditional on player i playing the sequence of actions that leads to I_i. Define

$$\mu_{I_i}(\beta) := \limsup_{\varepsilon \to 0} \mu_{I_i}(\beta^\varepsilon),$$

where the lim sup is with respect to set limits.[10] Note that $\limsup_{\varepsilon\to 0} \mu_{I_i}(\beta^\varepsilon)$ is always a closed set.

The following example illustrates the derivation of $\mu_{I_i}(\beta^\varepsilon)$, and its properties. These properties generalize to some extent to any finite extensive game, as shown in the subsequent two lemmas.

[10]If E_n is a sequence of sets, $\limsup_n E_n$ is the set of cluster points of sequences $y_n \in E_n$, and $\liminf_n E_n$ is the set of limit points of sequences $y_n \in E_n$. If $\limsup_n E_n = \liminf_n E_n$, the resulting set is equal to the Hausdorff limit of E_n. See Aubin and Frankowska (1990, pp. 16-23) and Nadler (1978, pp. 4-6) for a discussion of set limits.

Example 2. Consider again the game from Example 1, and let $\varepsilon = \frac{1}{10}$, $\beta_2(L) = \frac{4}{9}$, and $\beta_2(R) = \frac{2}{9}$. The unconditional probabilities at player 1's second information set corresponding to 2's strategy assign probability $\beta_2(L)$ to history L, and $\beta_2(R)$ to history R. The unconditional probabilities derived from the ε-contamination of 2's strategy are given by the sum of the vector $(1-\varepsilon)(\beta_2(L), \beta_2(R))$ and the projection of $\varepsilon\Delta\{O, L, R\}$ to the coordinates corresponding to L and R. The resulting set of all such unconditional probabilities are represented by the gray shaded triangle in Figure 2. Bayesian updating of all unconditional probabilities in this set results in a set of conditional probability distributions over $\{L, R\}$ where the probability assigned to L ranges from $\frac{4}{7}$ to $\frac{5}{7}$. In Figure 2, this set is represented by the thick black line connecting the points $\left(\frac{4}{7}, \frac{3}{7}\right)$ and $\left(\frac{5}{7}, \frac{2}{7}\right)$. If we let $\varepsilon \to 0$, the gray shaded triangle shrinks towards the black dot located at the point $(\beta_2(L), \beta_2(R))$, and the set of conditional probability distributions derived from the ε-contamination converges to the point $\left(\frac{2}{3}, \frac{1}{3}\right)$, which is just the conditional distribution derived from β_2. Note also that for a constant value of ε, the set of conditional distributions derived from the corresponding ε-contamination varies continuously in β_2. (The size of the gray shaded triangle stays constant, but the line connecting the origin of the coordinate system with the bottom left corner of the triangle changes continuously with β_2.)

If however we had started with $(\beta_2(L), \beta_2(R)) = (0, 0)$, the bottom left corner of the gray shaded triangle in Figure 2 would coincide with the origin, and the resulting conditional distributions over L and R would be represented by the line connecting the points $(1, 0)$ and $(0, 1)$. In this case, the set of conditional distributions derived from such an ε-contamination would not vary with ε, and its limit as $\varepsilon \to 0$ would therefore also be given by all distributions over $\{L, R\}$. $\triangleleft$

Note that the example illustrates two appealing properties of our modeling framework: If we interpret the "size" of the set of conditional beliefs as a measure for the degree of ambiguity, then ambiguity is clearly increasing in the probability ε of making a mistake. Moreover, given a fixed value of ε, the smaller the probability that a particular information set is reached according to the equilibrium strategies of a player's opponents, the larger the weight of the opponent's potential mistakes will be in the calculation of the player's conditional beliefs at this information set. The resulting effect is particularly stark at information sets that are not reached with positive probability—in the example above, this results in a set of conditional beliefs that is given by all distributions over histories in the corresponding information set. The greatest difference in the degree of ambiguity on and off the equilibrium path in the example is attained at the limit, when $\varepsilon \to 0$. In this case, beliefs on the equilibrium path are singletons, and beliefs off the equilibrium path have "full" ambiguity.

The following lemma shows that at information sets for which $\beta_{-i}[I_i] > 0$, $\mu_{I_i}(\beta)$ is always a singleton set that is equal to the Hausdorff limit of $\mu_{I_i}(\beta^\varepsilon)$.

Lemma 1. *If $\beta_{-i}[I_i] > 0$, then for any $\varepsilon \in (0,1)$, $\mu_{I_i}(\beta^\varepsilon)$ is a compact set, and $\mu_{I_i}(\beta)$ is the singleton set defined by the conditional distribution induced by β_{-i}, which is equal to the Hausdorff limit of $\mu_{I_i}(\beta^\varepsilon)$ as $\varepsilon \to 0$.*

Proof. Given β_{-i}, consider an information set I_i such that $\beta_{-i}[I_i] > 0$. Let I_j, with $j \neq i$, be an information set for which there exists an action $a_j \in A(I_j)$ such that a_j is contained in some history $h \in I_i$, i.e., I_j is an information set that "leads to I_i." For any such I_j, $\beta_j^\varepsilon(I_j)$ is a compact subset of a Euclidean space, since it is defined as the sum of $\varepsilon(\Delta A(I_j))$, which is compact, and the vector $(1-\varepsilon)\beta_j(I_j)$. Thus, the projection of $\beta_j^\varepsilon(I_j)$ onto coordinates that correspond to actions belonging to some history in I_i is also compact. Taking the Cartesian product of all such projections across all information sets I_j that lead to I_i will therefore also yield a compact set. The set $\mu_{I_i}(\beta^\varepsilon) \subset \Delta I_i$ is the image of a continuous function defined on this Cartesian product, and is therefore compact—the corresponding function is continuous since its domain is a compact subset of $\mathbb{R}^l_+$ for some l-dimensional Euclidean space, which does not include any vector that assigns an unconditional probability of zero to the information set I_i.

As $\varepsilon \to 0$, each $\beta_j^\varepsilon(I_j)$ converges in the corresponding Hausdorff metric to $\{\beta_j(I_j)\}$, and hence the Cartesian product described above converges to a singleton set. It follows that $\mu_{I_i}(\beta^\varepsilon)$ converges in the Hausdorff metric to the singleton set defined by the conditional distribution induced by β_{-i}. □

Note that when $\beta_{-i}[I_i] = 0$ at some information set I_i, the set $\mu_{I_i}(\beta)$ will in general be a non-singleton, which is sometimes equal to the set of all distributions over I_i, as in Example 2, but which can also be a strict subset of this set. As an example where this is the case, consider the game structure displayed in Figure 3, and assume that $\beta_1(L) = p \in (0,1)$ and $\beta_2(C) = 1$, so $\beta_{-3}[I_3] = 0$. Then every distribution in $\mu_{I_3}(\beta)$

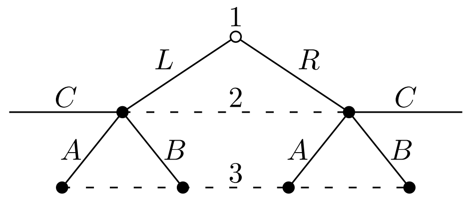

Figure 3: $\beta_{-3}[I_3] = 0$ if $\beta_2(C) = 1$.

at player 3's information set must assign probability p to the set of histories $\{LA, LB\}$, and must satisfy the property that the ratios of the probabilities assigned to LA vs. RA, and LB vs. RB, if defined, are equal to $\frac{p}{1-p}$. Hence, $\mu_{I_3}(\beta)$ would be a strict subset of the set of all distributions over I_3. More generally, the following lemma shows that the sets $\mu_{I_i}(\beta^\varepsilon)$ are ordered by set inclusion, which allows a simplification of the definition of $\mu_{I_i}(\beta)$ as the Hausdorff limit of $\mu_{I_i}(\beta^\varepsilon)$.[11]

Lemma 2. *For every strategy profile β and information set I_i, the sets $\mu_{I_i}(\beta^\varepsilon)$ are monotonic in ε relative to set inclusion, and thus $\mu_{I_i}(\beta) = \lim_{\varepsilon\to 0}\mu_{I_i}(\beta^\varepsilon)$.*

Proof. As in the proof of Lemma 1, consider an information set I_j leading to I_i. Then the ε-contamination of $\beta_j(I_j)$ is given by the set $(1-\varepsilon)\beta_j(I_j)+\varepsilon(\Delta A(I_j))$, which is monotonic in ε. To see why, note that this set contains the point $\beta_j(I_j) = (1-\varepsilon)\beta_j(I_j) + \varepsilon\beta_j(I_j)$,[12] and that the ε-contamination always shrinks monotonically towards this point as $\varepsilon \to 0$. Thus, the projection of the ε-contamination of $\beta_j(I_j)$ to coordinates that correspond to actions belonging to some history in I_i, is also monotonic in ε.

The unconditional probabilities over histories in I_i resulting from the Cartesian product over all such projections across all information sets I_j leading to I_i, are therefore also monotonic in ε. If $\beta_{-i}[I_i] > 0$, the corresponding set of unconditional probabilities does not contain the zero vector for any $\varepsilon \in (0,1)$, and if $\beta_{-i}[I_i] = 0$, this set contains the zero vector for every $\varepsilon \in (0,1)$. It follows that the set of induced conditional probabilities $\mu_{I_i}(\beta^\varepsilon)$, which is the continuous image of all non-zero vectors of unconditional probabilities, is also monotonic in ε. □

[11] Note that this result also implies that $\mu_{I_i}(\beta^\varepsilon)$ converges to $\mu_{I_i}(\beta)$ in terms of net convergence, where the corresponding net is indexed by $\varepsilon \in (0,1)$.

[12] In Figure 2, this point is represented by the black dot in the grey shaded triangle.

We next propose three different definitions of equilibrium for our environment. The first considers the case where players' beliefs are derived from an ε-contamination of their opponents' strategies, and the equilibrium strategies are optimal at each information set given these beliefs. We call the resulting equilibria ε-Perfect Maxmin Equilibria (ε-PME). The second definition considers the set limit as $\varepsilon \to 0$ of the set of all ε-PME, and defines any strategy profile in this limit to be a Perfect Maxmin Equilibrium (PME). The last definition considers first the set limit $\mu_{I_i}(\beta)$ of the beliefs $\mu_{I_i}(\beta^\varepsilon)$ induced by ε-contaminations as $\varepsilon \to 0$, and defines a strategy profile β to be a Strong Perfect Maxmin Equilibrium (SPME) if it prescribes strategies that are optimal for $\mu_{I_i}(\beta)$ at each information set I_i. In the following, the notation $\beta_i(I_i+)$ is used to denote the continuation strategy induced by β_i at all information sets of player i that follow I_i.

Definition 2 (ε-Perfect Maxmin Equilibrium)**.** A strategy profile $\beta \in \mathcal{B}$ is an *ε-perfect maxmin equilibrium* if for every i and $I_i \in \mathcal{I}_i$,

$$\beta_i(I_i) \in \arg \max_{\beta_i'(I_i) \in \Delta(A(I_i))} \inf\{u_i(\mathcal{O}((\beta_i'(I_i), \beta_i(I_i+), \tilde{\beta}_{-i}), \tilde{\mu}_{I_i})) \,|\, \tilde{\beta}_{-i} \in \beta_{-i}^\varepsilon,\, \tilde{\mu}_{I_i} \in \mu_{I_i}(\beta^\varepsilon)\}.$$

Definition 3 (Perfect Maxmin Equilibrium)**.** Let $\mathcal{B}(\varepsilon\text{-PME})$ denote the set of ε-PME. A strategy profile $\beta \in \mathcal{B}$ is a *perfect maxmin equilibrium* if $\beta \in \limsup_{\varepsilon \to 0} \mathcal{B}(\varepsilon\text{-PME})$.

Note that an ε-PME strategy β is required to be optimal given beliefs derived from an ε-contamination of β, for a particular ε. A strategy β constitutes a PME if it can be approximated by a sequence of ε-PMEs as $\varepsilon \to 0$, using the respective Euclidean metric. Thus, we can view a PME as approximating the equilibrium behavior resulting from an ε-PME, when the probability of a mistake is arbitrarily small. An equilibrium definition that is stronger than that of an ε-PME, could require that an equilibrium strategy β is optimal for beliefs $\mu_{I_i}(\beta^\varepsilon)$ derived from β, for *all* values of ε that are smaller than some cutoff value (which would imply that β is an ε-PME for all small enough ε). Although we will not pursue such an equilibrium definition in the paper, we introduce next a closely related definition that requires optimality of an equilibrium strategy β for beliefs defined by $\mu_{I_i}(\beta) = \lim_{\varepsilon \to 0} \mu_{I_i}(\beta^\varepsilon)$. Clearly, if β is optimal for all $\mu_{I_i}(\beta^\varepsilon)$ with ε small enough, it will be optimal for the resulting limit beliefs $\mu_{I_i}(\beta)$, but the converse may not hold. However, as we will show, even the weaker definition is too strong to guarantee existence of a corresponding equilibrium.

Definition 4 (Strong Perfect Maxmin Equilibrium)**.** A strategy profile $\beta \in \mathcal{B}$ is a *strong perfect maxmin equilibrium* if for every i and $I_i \in \mathcal{I}_i$,

$$\beta_i(I_i) \in \arg \max_{\beta_i'(I_i) \in \Delta(A(I_i))} \min\{u_i(\mathcal{O}((\beta_i'(I_i), \beta_i(I_i+), \beta_{-i}), \tilde{\mu}_{I_i})) \mid \tilde{\mu}_{I_i} \in \mu_{I_i}(\beta)\}.$$

Our equilibrium definitions implicitly assume that players maintain the assumption that their opponents will follow their equilibrium strategies with high probability, even after having observed a deviation from the equilibrium strategies. Thus, any deviation is interpreted as a small probability mistake that has no implications regarding future play. Equilibria are defined by behavioral strategies profiles β—for ε-PME and SPME, the equilibrium strategies must be optimal for the corresponding belief systems defined by $\mu_{I_i}(\beta^\varepsilon)$ and $\mu_{I_i}(\beta)$, but strategies corresponding to a PME β may not necessarily be optimal for the belief system defined by $\mu_{I_i}(\beta)$.[13]

An alternative way to define an equilibrium, which does not rely on ε-contaminations, could require optimality of an equilibrium strategy β for beliefs defined by $\mu_{I_i}(\beta)$ at information sets I_i with $\beta_{-i}[I_i] > 0$ (in which case beliefs would be unambiguous by Lemma 1, and could just be derived from β_{-i} using Bayes' rule), and by ΔI_i at information sets I_i with $\beta_{-i}[I_i] = 0$ (so beliefs reflect "full" ambiguity). We do not pursue such a definition for two principal reasons: As a practical reason, such an equilibrium may not always exist—in Example 3 below, it would also be an SPME, and the non-existence of SPME in the example implies that such an equilibrium would not exist either. More importantly, the previous discussion of the game from Figure 3 shows that even if $\beta_{-i}[I_i] = 0$, the knowledge of equilibrium strategies chosen before I_i can significantly restrict the feasible beliefs at I_i, in which case the approach using ε-contaminations provides a tractable and intuitive way to derive the resulting beliefs. While other approaches to define equilibria that incorporate some form of ambiguity, especially off the equilibrium path, are possible, we view our approach based on ε-contaminations as the the most parsimonious departure from the standard analysis under expected utility maximization, which maintains the usual assumption that players know the equilibrium strategy profile, and also has an intuitive interpretation based on trembles modeled by ε-contaminations.

[13]The last paragraph of Example 3 illustrates why this is the case.

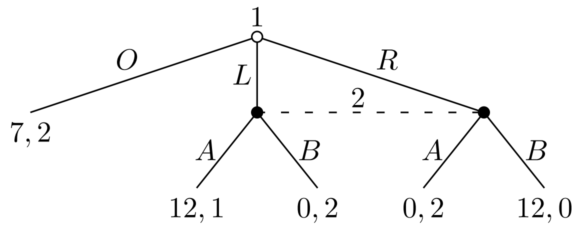

Figure 4: Example 3.

4 (Non-)Existence

For the following example, no SPME exists, but an ε-PME exists for every $\varepsilon \in (0, 1)$:

Example 3. For the game described in Figure 4, consider the possibility of an SPME where $\beta_1(L) > 0$ and $\beta_1(R) = 0$. Then 2's beliefs assign probability one to L, and thus his best response is to play B, in which case R is optimal for player 1. A similar argument shows that there is no SPME with $\beta_1(R) > 0$ and $\beta_1(L) = 0$. For an SPME to exist where $\beta_1(L) > 0$ and $\beta_1(R) > 0$, 1 must be indifferent between these two actions, which can only be the case when $\beta_2(A) = \beta_2(B) = \frac{1}{2}$. But then O is a best response for player 1. Consider as a last option the possibility of an SPME where $\beta_1(O) = 1$. Then 2's information set is not reached according to 1's strategy, and thus his set of beliefs is given by all possible distributions over $\{L, R\}$.[14] Letting μ denote the probability assigned to the history L by 2's beliefs, and p denote the probability that 2 assigns to his action A, 2's expected utility is given by

$$\begin{aligned} u_2(\beta_1, p) &= \min_{\mu \in [0,1]} \{\mu[p + 2(1-p)] + (1-\mu)2p\} \\ &= \min_{\mu \in [0,1]} \{\mu(2 - 3p) + 2p\} \\ &= \begin{cases} 2p, & \text{if } p \leq \frac{2}{3}, \\ 2 - p & \text{if } p > \frac{2}{3}. \end{cases} \end{aligned}$$

It follows that $p = \frac{2}{3}$ is the unique optimal strategy for player 2, in which case the payoff

[14]The set of unconditional probabilities over $\{L, R\}$ induced by an ε-contamination is the convex hull of $\{(0,0), (0,\varepsilon), (\varepsilon, 0)\}$. The corresponding set of conditional distributions is equal to $\Delta\{L, R\}$ for any $\varepsilon \in (0, 1)$, and therefore so is the set limit as $\varepsilon \to 0$.

1 receives from L is equal to 8, which is higher than the payoff of 7 he receives from playing O. Hence, the game has no SPME.

The case where 1 plays O and 2's beliefs are given by the set of all $\mu \in [0,1]$, illustrates that a player with maxmin preferences can strictly prefer randomizing over playing the corresponding pure strategies. As the minimizing value of μ is not constant across p, and hence the resulting expected payoffs $u_2(\beta_1, p)$ are not linear in p, the (unique) optimal choice of $p = \frac{2}{3}$ allows player 2 to completely hedge against the ambiguity he faces, and achieve a strictly higher expected payoff than he can get from any pure strategy or any alternative randomization. In contrast, since SE is based on standard expected utility preferences, 2's expected payoffs are linear in p for any SE, and hence player 2 only randomizes over his actions if he is indifferent between the pure actions and all corresponding randomizations. As a result, there exist a continuum of SE where 1 plays O and 2 chooses any randomization with $p \in \left[\frac{5}{12}, \frac{7}{12}\right]$, based on precise beliefs given by $\mu = \frac{2}{3}$.

As indicated previously, the game always possesses ε-PME, which are derived in Appendix A. In particular, we show that for $\varepsilon < \frac{1}{8}$, there is a unique ε-PME given by

$$\beta_1(L) = \frac{2\varepsilon}{(1-\varepsilon)}, \quad \beta_1(R) = 0, \quad \text{and} \quad \beta_2(A) = \frac{7}{(1-\varepsilon)12},$$

which implies that the PME of this game is the limit of the ε-PME as $\varepsilon \to 0$, and is given by $\beta_1(O) = 1$ and $\beta_2(A) = \frac{7}{12}$. In this example, the set of PME is a strict subset of the set of SE (which in this case is the same as the set of WPBE). The SE of the game are summarized by $\beta_1(O) = 1$, $\beta_2(A) \in \left[\frac{5}{12}, \frac{7}{12}\right]$, and $\mu_2(L) = \frac{2}{3}$.

To see what happens at the limit as $\varepsilon \to 0$, denote the ε-PME derived in this example by $\beta_{[\varepsilon]}$. Then $\beta_{[\varepsilon]1}(L) = \frac{2\varepsilon}{(1-\varepsilon)} \to 0$, but the set-valued conditional beliefs of player 2 derived from $\beta_{[\varepsilon]1}$ are constant across ε and given by $\mu_2(L) \in \left[\frac{2}{3}, 1\right]$, and hence also converge to $\left[\frac{2}{3}, 1\right]$ as $\varepsilon \to 0$. For ambiguous beliefs defined by $\left[\frac{2}{3}, 1\right]$, player 2 maximizes his expected utility by choosing any $\beta_2(A) \in \left[0, \frac{2}{3}\right]$, so the PME strategy that assigns probability $\frac{7}{12}$ to action A is optimal. However, if we consider an ε-contamination of player 1's PME strategy, which assigns probability 1 to action O, and take the limit of the induced beliefs as $\varepsilon \to 0$, then this limit is given by all $\mu_2(L) \in [0,1]$, in which case player 2's optimal strategy is to play action A with probability $\frac{2}{3}$. $\lhd$

The existence result from the example generalizes to any finite extensive game:

Theorem 1. *Every finite extensive game with perfect recall has an ε-PME for $\varepsilon \in (0,1)$.*

Proof. The consistent planning assumption implies that each information set I_i defines a distinct individual of player i who maximizes the infimum of his expected utility over the conditional beliefs over histories in I_i derived from an ε-contamination β^{ε}_{-i} of his opponents' strategy profile, *and* over all distributions over terminal histories following I_i, as induced by $\beta_i(I_i+)$ and β^{ε}_{-i}. We can therefore consider the best responses of each player at every one of his information sets separately.

Fix an information set I_i of player i, and consider an information set I_j for which there exists an action $a_j \in A(I_j)$ such that for some $h \in I_i$, a_j is either contained in the history h, or is the last action in a history that has h as a sub-history. Then I_j is either an information set that "leads to I_i," in the first case, or one that "follows I_i," in the second case.[15] Similarly, we will refer to the corresponding action a_j as described above as an action that leads to I_i or an action that follows I_i. Given a strategy profile β_{-i}, let $\beta'_{-i} \in \beta^{\varepsilon}_{-i}$ denote a strategy profile for which $\beta'_{-i}[I_i] > 0$.[16] Then the strategies constituting β'_{-i} that correspond to information sets that lead to I_i define a conditional distribution over histories in I_i using Bayesian updating. Let μ'_{I_i} denote the corresponding conditional beliefs. If player i uses a strategy $\sigma \in \Delta A(I_i)$ at I_i, and uses strategies given by $\beta_i(I_i+)$ at any of his information sets that follow I_i, the probability of a terminal history $(h, a_i, a_{j1}, \ldots, a_{jk})$ with $h \in I_i$, conditional on I_i, is given by $\mu'_{I_i}(h)\sigma(a_i)\prod_{l=1}^{l=k}\tilde{\beta}_{jl}(a_{jl})$, where $\tilde{\beta}_{jl}$ is either defined by β'_{-i} if $jl \neq i$, or by $\beta_i(I_i+)$ if $jl = i$. This is a consequence of perfect recall, since no information set can appear more than once along any particular history. Let $U_i(\sigma, \beta_i(I_i+), \beta'_{-i}, \mu'_{I_i})$ denote the expected utility of player i conditional on I_i, as induced by these probabilities over terminal histories. Note that we write U_i as a function of β'_{-i} for notational simplicity, even though only those components of β'_{-i} that correspond to information sets that follow I_i will affect this expected utility, as the components that lead to I_i are reflected in μ'_{I_i}.

Given the strategy profile β and information set I_i, let $\Gamma_{I_i}(\beta^{\varepsilon})$ denote the set of all pairs $(\beta'_{-i}, \mu'_{I_i})$ derived from strategies $\beta'_{-i} \in \beta^{\varepsilon}_{-i}$ with $\beta'_{-i}[I_i] > 0$. Then the expected

[15] Note that it is possible for an information set I_j to both lead to I_i and follow I_i.

[16] Recall that $\beta_{-i}[I_i]$ denotes the probability of reaching I_i under β_{-i}, assuming that player i plays the unique action sequence that leads to I_i.

utility of player i at I_i corresponding to β and a strategy $\sigma \in \Delta A(I_i)$ is given by

$$\inf_{(\beta'_{-i},\mu'_{I_i})\in\Gamma_{I_i}(\beta^\varepsilon)} U_i(\sigma, \beta_i(I_i+), \beta'_{-i}, \mu'_{I_i}),^{17}$$

and his best response at I_i to an ε-contamination of β is defined by

$$BR_{I_i}(\beta_i(I_i+), \beta^\varepsilon_{-i}) := \arg\max_{\sigma\in\Delta A(I_i)} \inf_{(\beta'_{-i},\mu'_{I_i})\in\Gamma_{I_i}(\beta^\varepsilon)} U_i(\sigma, \beta_i(I_i+), \beta'_{-i}, \mu'_{I_i}).$$

We will show that the closure of $\Gamma_{I_i}(\beta^\varepsilon)$, $\text{cl}(\Gamma_{I_i}(\beta^\varepsilon))$, is continuous in β. We then have

$$\begin{aligned}\bar{U}_{I_i}(\sigma,\beta) &:= \inf_{(\beta'_{-i},\mu'_{I_i})\in\Gamma_{I_i}(\beta^\varepsilon)} U_i(\sigma, \beta_i(I_i+), \beta'_{-i}, \mu'_{I_i}) \\ &= \min_{(\beta'_{-i},\mu'_{I_i})\in\text{cl}(\Gamma_{I_i}(\beta^\varepsilon))} U_i(\sigma, \beta_i(I_i+), \beta'_{-i}, \mu'_{I_i}),\end{aligned}$$

and since U_i is linear in σ for every $(\beta'_{-i}, \mu'_{I_i})$, it follows that $\bar{U}_{I_i}$ is concave in σ for every β. Furthermore, since U_i is continuous in each of its arguments, if $\text{cl}(\Gamma_{I_i}(\beta^\varepsilon))$ is continuous in β, the Maximum Theorem implies that $\bar{U}_{I_i}$ is continuous in (σ,β).

We can then apply the Maximum Theorem to the problem $\max_{\sigma\in\Delta A(I_i)} \bar{U}_{I_i}(\sigma,\beta)$ that defines the best response $BR_{I_i}(\beta_i(I_i+), \beta^\varepsilon_{-i})$. Since $\Delta A(I_i)$ is compact and convex, and independent of β, the Maximum Theorem implies that BR_{I_i} is convex-valued, upper hemi-continuous and closed-valued, and hence closed. A standard application of Kakutani's fixed point theorem, where each information set of a player is considered as a separate agent of this player, then shows that an ε-PME exists.

It remains to show that $\text{cl}(\Gamma_{I_i}(\beta^\varepsilon))$ is continuous in β. We start by proving a preliminary lemma. Consider an arbitrary vector $\rho \in \Delta^{k-1}$ in the $(k-1)$-dimensional simplex, so ρ can be interpreted as a distribution over the action set $A(I_j)$ at some information set I_j. Then the ε-contamination of ρ, defined by

$$\rho^\varepsilon = \{(1-\varepsilon)\rho + \varepsilon q \,|\, q \in \Delta^{k-1}\},$$

is compact, since it is the sum of $(1-\varepsilon)\rho$ and $\varepsilon\Delta^{k-1}$, and the following lemma shows that ρ^ε is continuous in ρ:

Lemma 3. *If $(\rho_n) \subset \Delta^{k-1}$ is a convergent sequence such that $\rho_n \to \rho$, then $\rho^\varepsilon_n \to \rho^\varepsilon$ in the corresponding Hausdorff metric over subsets of Δ^{k-1}.*

[17] We assume that player i has perfect recall and considers his future actions to be determined by $\beta_i(I_i+)$. The results would not change if we would instead assume that i's beliefs about his future actions are given by $\beta^\varepsilon_i(I_i+)$.

Proof. Since $(1-\varepsilon)\rho_n + \varepsilon q \to (1-\varepsilon)\rho + \varepsilon q$ for every $q \in \Delta^{k-1}$, every $r \in \rho^\varepsilon$ is the limit of a sequence $r_n \in \rho_n^\varepsilon$. Furthermore, every sequence $(1-\varepsilon)\rho_{n_k} + \varepsilon q_{n_k}$, if convergent, must converge to a point in ρ^ε, since $\rho_{n_k} \to \rho$, which implies that q_{n_k} must converge to some element of Δ^{k-1}. Thus, every element of ρ^ε is a limit point of ρ_n^ε, and every cluster point of ρ_n^ε is an element of ρ^ε. Since the set of limit points is a subset of the set of cluster points, this implies that $\rho^\varepsilon \subset \liminf \rho_n^\varepsilon \subset \limsup \rho_n^\varepsilon \subset \rho^\varepsilon$, and hence $\rho^\varepsilon = \liminf \rho_n^\varepsilon = \limsup \rho_n^\varepsilon$, so the result follows from Theorem 0.7 in Nadler (1978). $\square$

To show that $\mathrm{cl}(\Gamma_{I_i}(\beta^\varepsilon))$ is continuous in β, let $\beta^n \to \beta$, and consider any point $(\tilde\beta_{-i}, \tilde\mu_{I_i}) \in \mathrm{cl}(\Gamma_{I_i}(\beta^\varepsilon))$. Then there exists a sequence $(\tilde\beta_{-i}^k, \tilde\mu_{I_i}^k) \to (\tilde\beta_{-i}, \tilde\mu_{I_i})$ such that $(\tilde\beta_{-i}^k, \tilde\mu_{I_i}^k) \in \Gamma_{I_i}(\beta^\varepsilon)$ for every k, and hence $\tilde\beta_{-i}^k[I_i] > 0$ for every k. Since β^ε is continuous in β, every such $(\tilde\beta_{-i}^k, \tilde\mu_{I_i}^k)$ is the limit of a sequence $((\tilde\beta_{-i}^{n,k}, \tilde\mu_{I_i}^{n,k})) \subset \Gamma_{I_i}(\beta^{n,\varepsilon})$, which implies that

$$\mathrm{cl}(\Gamma_{I_i}(\beta^\varepsilon)) \subset \liminf \mathrm{cl}(\Gamma_{I_i}(\beta^{n,\varepsilon})).$$

Now consider any sequence $(\tilde\beta_{-i}^{n_l}, \tilde\mu_{I_i}^{n_l}) \to_l (\tilde\beta_{-i}, \tilde\mu_{I_i})$ such that $(\tilde\beta_{-i}^{n_l}, \tilde\mu_{I_i}^{n_l}) \in \mathrm{cl}(\Gamma_{I_i}(\beta^{n_l,\varepsilon}))$. Since each $\beta_{-i}^{n_l,\varepsilon}$ is closed, $\tilde\beta_{-i}^{n_l} \in \beta_{-i}^{n_l,\varepsilon}$ for every n_l, and hence Lemma 3 implies that $\tilde\beta_{-i} \in \beta^\varepsilon$. If $\tilde\beta_{-i}[I_i] > 0$, we must have $(\tilde\beta_{-i}, \tilde\mu_{I_i}) \in \Gamma_{I_i}(\beta^\varepsilon)$. If $\tilde\beta_{-i}[I_i] = 0$, there exists a sequence $(\tilde\beta_{-i}^m) \subset \beta_{-i}^\varepsilon$ such that $\tilde\beta_{-i}^m \to \tilde\beta_i$ and $\tilde\beta_{-i}^m[I_i] > 0$ for all m. But then $(\tilde\beta_{-i}^m, \tilde\mu_{I_i}^m) \in \Gamma_{I_i}(\beta^\varepsilon)$ and $(\tilde\beta_{-i}^m, \tilde\mu_{I_i}^m) \to (\tilde\beta_{-i}, \tilde\mu_{I_i})$, which implies

$$\limsup \mathrm{cl}(\Gamma_{I_i}(\beta^{n,\varepsilon})) \subset \mathrm{cl}(\Gamma_{I_i}(\beta^\varepsilon)),$$

and hence, since $\liminf \mathrm{cl}(\Gamma_{I_i}(\beta^{n,\varepsilon})) \subset \limsup \mathrm{cl}(\Gamma_{I_i}(\beta^{n,\varepsilon}))$, the continuity of the closure follows from Theorem 0.7 in Nadler (1978). $\square$

Since a PME is defined as a cluster point of the set of ε-PME as $\varepsilon \to 0$, and the set of behavioral strategy profiles is compact, Theorem 1 yields the following:

Corollary 4. *Every finite extensive game with perfect recall has a PME.*

5 Relation to other equilibrium notions

Since SPME and PME are defined by (unambiguous) behavioral strategy profiles, we can ask whether such strategy profiles also constitute a Nash equilibrium of the game in consideration, when the corresponding utility functions u_i are interpreted as standard von Neumann-Morgenstern utility functions. The answer to this question is indeed affirmative:

Proposition 5. *Every SPME and PME is a Nash equilibrium.*

Proof. Consider first an SPME defined by a strategy profile β. If $\beta_{-i}[I_i] = 0$ for an information set I_i, then the information set I_i will never be reached with positive probability given the strategies of i's opponents, β_{-i}, no matter what actions player i chooses at any of his information sets. Thus, no deviation by player i at such an information set I_i has any effect on the induced distribution over terminal histories, and therefore every strategy at I_i is optimal with respect to the strategic form of the game. Now consider an information set I_i with $\beta_{-i}[I_i] > 0$. Then Lemma 1 implies that i's beliefs over histories in I_i in the SPME are given by the singleton set containing the probability distribution derived from β_{-i} using Bayesian updating, and thus that these beliefs correctly capture the actions induced by β_{-i} whenever i chooses an action sequence that leads to I_i. Thus, the fact that β is a SPME implies that β_i prescribes an optimal strategy at every I_i with $\beta_{-i}[I_i] > 0$, given precise beliefs over histories in I_i that are induced by β_{-i}, and assuming that β determines all players' strategies at information sets other than I_i. The "one-shot-deviation principle" for extensive games with perfect recall (Hendon et al., 1996) then implies that β_i is sequentially rational for i at all I_i with $\beta_{-i}[I_i] > 0$, and hence that β_i constitutes an optimal strategy for the strategic form.

Now let β denote a PME. Then β is the limit of a sequence of ε-PME strategy profiles $\beta_{[\varepsilon]}$ converging to β as $\varepsilon \to 0$. Since $\beta_{-i[\varepsilon]}[I_i] \to \beta_{-i}[I_i]$, it follows that if $\beta_{-i}[I_i] > 0$, then $\beta_{-i[\varepsilon]}[I_i] > 0$ for small enough ε, and furthermore $\beta_{-i[\varepsilon]}[I_i]$ is also bounded away from 0 for small enough ε. This implies that the conditional beliefs $\mu_{I_i}(\beta^{\varepsilon}_{[\varepsilon]})$ converge to a singleton set given by the probability distribution over I_i derived from β_{-i} using Bayesian updating. Since the strategies induced by $\beta_{[\varepsilon]}$ at I_i are optimal given $\mu_{I_i}(\beta^{\varepsilon}_{[\varepsilon]})$ and the continuation strategies defined by $\beta_{[\varepsilon]}$, the limit strategies induced by β at I_i are optimal given the limit of $\mu_{I_i}(\beta^{\varepsilon}_{[\varepsilon]})$. Therefore, the same arguments we made for an SPME imply that β is a Nash equilibrium. □

In addition to being a Nash equilibrium, every SPME or PME strategy profile also defines a subgame perfect equilibrium (SPE):

Proposition 6. *Every SPME and PME is a subgame perfect equilibrium.*

Proof. Consider an SPME β and a history h^* such that all histories following h^* define a subgame of the original game. Denote this subgame by $G(h^*)$, and the restriction of any strategy β_i to $G(h^*)$ by $\beta_{i|G(h^*)}$. Furthermore, for any information set I_i in $G(h^*)$, let $\beta_{-i|G(h^*)}[I_i]$ be the probability of reaching I_i conditional on h^*, given strategies $\beta_{-i|G(h^*)}$ and assuming that player i's actions in $G(h^*)$ are consistent with I_i.

To show that β induces a NE in $G(h^*)$, we can ignore i's actions at every I_i with $\beta_{-i|G(h^*)}[I_i] = 0$, since no such information set will be reached with positive probability in the subgame given $\beta_{-i|G(h^*)}$. Consider then any I_i with $\beta_{-i|G(h^*)}[I_i] > 0$, and let $\mu_{I_i}(\beta'_{-i})$ denote the beliefs over histories in I_i induced by some $\beta'_{-i} \in \beta^{\varepsilon}_{-i}$ with $\beta'_{-i}[I_i] > 0$. If $\beta'_{-i}[I_i] > 0$, it must also be the case that $\beta'_{-i|G(h^*)}[I_i] > 0$, and furthermore the beliefs $\mu_{I_i}(\beta'_{-i})$ must be independent of the probability of reaching h^* under β'_{-i}, since this probability always yields a common factor in the corresponding unconditional probabilities. Hence, $\mu_{I_i}(\beta'_{-i})$ is defined by the Bayesian update of $\beta_{-i|G(h^*)}$, assuming that i's actions are consistent with I_i. Since our contaminations are independent across information sets, it follows that in order to derive $\mu_{I_i}(\beta^{\varepsilon})$, we only need to consider the unconditional probabilities over histories in I_i induced by β^{ε} in the subgame $G(h^*)$, i.e., conditional on h^*. An application of Lemma 1 to the subgame then implies that $\mu_{I_i}(\beta)$ is a singleton set defined by the Bayesian update of $\beta_{-i|G(h^*)}$ in $G(h^*)$, and the argument used in the proof of Proposition 5 then implies that β induces a NE in $G(h^*)$.

To show that a PME β is also a SPE, note that if $\beta_{[\varepsilon]} \to \beta$, then $\beta_{-i|G(h^*)[\varepsilon]} \to \beta_{-i|G(h^*)}$ for any subgame $G(h^*)$, and hence $\beta_{-i|G(h^*)[\varepsilon]}[I_i] \to \beta_{-i|G(h^*)}[I_i]$. It follows that whenever $\beta_{-i|G(h^*)}[I_i] > 0$, $\beta_{-i|G(h^*)[\varepsilon]}[I_i] > 0$ for small enough ε. Since beliefs induced by $\beta^{\varepsilon}_{-i[\varepsilon]}$ at I_i only depend on $\beta^{\varepsilon}_{-i|G(h^*)[\varepsilon]}$, the argument use in the proof of Proposition 5 shows that β also induces a NE in $G(h^*)$. □

The proof of the previous proposition shows that any player i's beliefs corresponding to an SPME or PME are precise at all information sets *except* those information sets I_i for which there does *not* exist a subgame $G(h^*)$ such that $\beta_{-i|G(h^*)}[I_i] > 0$. Hence, non-trivial ambiguous beliefs can only arise at an information set I_i if $\beta_{-i|G(h^*)}[I_i] = 0$ for every subgame containing I_i. Even though this property significantly restricts the possibility of non-precise beliefs, as we show next, there exist games where SPME and

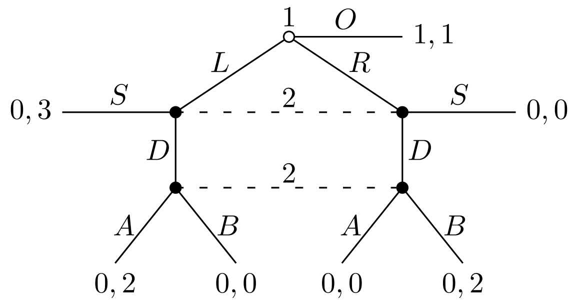

Figure 5: Example 4.

PME strategies are not included in any WPBE or SE. Hence, the "belief-system-based" refinements of WPBE and SE may not always capture outcomes that may arise if players are ambiguity averse, even though in some games they may yield much larger sets of predictions, such as in Example 3. The following example shows that an SPME/PME strategy profile may not yield a WPBE:

Example 4. Consider the game depicted in Figure 5. Since player 1 has a dominant strategy to play O, he must assign probability 1 to O in any type of equilibrium. Let μ denote the probability assigned to history L at player 2's first information set, and let δ denote the probability assigned to history LD at player 2's second information set. The fact that 1 must play O in any equilibrium, implies that any ε-contamination of 1's strategy will yield conditional beliefs at the two information sets of player 2 that are given by all $\mu \in [0,1]$ and $\delta \in [0,1]$. Consider first the second information set of player 2, and denote the probability he assigns to action A by q. Then his expected utility conditional on reaching this information set is given by

$$u_{2^2}(\beta_1, q) = \min_{\delta \in [0,1]} \{\delta(2q) + (1-\delta)2(1-q)\} = \min_{\delta \in [0,1]} \{\delta(4q-2) + 2(1-q)\}$$
$$= \begin{cases} 2q, & \text{if } q \leq \frac{1}{2}, \\ 2-2q & \text{if } q > \frac{1}{2}, \end{cases}$$

and thus his unique optimal strategy is to set $q = \frac{1}{2}$. The consistent planning assumption implies that at player 2's first information set, he takes the optimal strategy $q = \frac{1}{2}$ at his second information set as given, and chooses a probability p to assign to action S that maximizes

$$u_{2^1}(\beta_1, p) = \min_{\mu\in[0,1]} \{\mu[3p + (1-p)] + (1-\mu)(1-p)\}$$
$$= \min_{\mu\in[0,1]} \{\mu 3p + 1 - p\} = 1 - p.$$

Hence, the optimal strategy at 2's first information set is to set $p = 0$. It follows that the game has a unique SPME, ε-PME and PME, where 1 plays O, and 2 sets $p = 0$ and $q = \frac{1}{2}$. Clearly, this strategy profile constitutes a NE. The strategy $\left(p = 0, q = \frac{1}{2}\right)$ yields a WPBE if there exists a (precise) belief system $(\bar{\mu}, \bar{\delta})$ such that $q = \frac{1}{2}$ is sequentially rational at 2's second information set if $\bar{\delta} \in [0,1]$ is the probability assigned to history LD, and $\left(p = 0, q = \frac{1}{2}\right)$ is sequentially rational at 2's first information set if $\bar{\mu} \in [0,1]$ is the probability assigned to the history L. Clearly, for $\bar{\delta} = \frac{1}{2}$, 2 is indifferent between his two actions at his second information set, so $q = \frac{1}{2}$ is sequentially rational given such beliefs. To check sequential rationality at 2's first information set, notice that we need to find a value $\bar{\mu} \in [0,1]$ such that $p = 0$ and $q = \frac{1}{2}$ are jointly optimal at the first information set. Thus, for this strategy profile to yield a WPBE, there must exist a $\bar{\mu} \in [0,1]$ so that $\left(p = 0, q = \frac{1}{2}\right)$ solve

$$\max_{(p,q)} \{\bar{\mu}[3p + 2(1-p)q] + (1-\bar{\mu})[2(1-p)(1-q)]\}$$
$$= \max_{(p,q)} \{(4\bar{\mu} - 2)(1-p)q + (5\bar{\mu} - 2)p + 2 - 2\bar{\mu}\}.$$

For $p = 0$ and $q = \frac{1}{2}$ the corresponding payoff is 1, and is thus independent of $\bar{\mu}$. If $\bar{\mu} \geq \frac{1}{2}$, and we set $p = q = 1$, the resulting payoff is $3\bar{\mu} \geq \frac{3}{2} > 1$; similarly, if $\bar{\mu} < \frac{1}{2}$ we can set $p = q = 0$, which yields a payoff $2 - 2\bar{\mu} > 1$. It follows that no $\bar{\mu} \in [0,1]$ exists for which $p = 0$ and $q = \frac{1}{2}$ are sequentially rational at 2's first information set. $\triangleleft$

The main reason why the SPME/PME from the previous example is not part of a WPBE, is that sequential rationality does not hold at the first of two subsequent information sets of a player. The following proposition shows that in games where each player's information sets are not interdependent in such a way, every SPME/PME yields a WPBE:

Proposition 7. *Assume that no two distinct information sets I_i and I_i' of any player i have the property that I_i' contains histories that have sub-histories in I_i. Then if β is a* *strategy profile that defines an SPME or PME, there exists a belief system γ consisting of precise beliefs, such that (β, γ) is a WPBE.*

Proof. Note first that if the assumption of the proposition is satisfied, sequential rationality at any player's information set can be checked independently of his other information sets. Now let β denote an SPME or PME. Then the strategy prescribed by β at an information set I_i with $\beta_{-i}[I_i] > 0$ is optimal given precise beliefs that are derived as the limit of $\mu_{I_i}(\beta^\varepsilon)$ if β is an SPME, or the limit of $\mu_{I_i}(\beta^\varepsilon_{[\varepsilon]})$ if β is a PME defined as the limit of a sequence of ε-PME $\beta_{[\varepsilon]}$. This follows from Lemma 1 and the proof of Proposition 5, and from the fact that future actions are not contaminated at the limit for an SPME or PME. For the case where β is a PME, let $\bar{\mu}_{I_i}(\beta) := \lim_{\varepsilon \to 0} \mu_{I_i}(\beta^\varepsilon_{[\varepsilon]})$. By the definition of SPME and PME, if $\beta_{-i}[I_i] = 0$ at an information set I_i, the strategy prescribed by β at I_i, $\beta_i(I_i)$, is optimal given the ambiguous beliefs defined by $\mu_{I_i}(\beta)$ or $\bar{\mu}_{I_i}(\beta)$, respectively. To show that β also yields a WPBE, we need to show that there exist precise beliefs γ_{I_i} at I_i, under which $\beta_i(I_i)$ is sequentially rational. We prove the existence of such a γ_{I_i} when β is an SPME. The proof for the case where β is a PME follows by substituting $\bar{\mu}_{I_i}(\beta)$ for $\mu_{I_i}(\beta)$.

Let h_k denote an arbitrary history in I_i, let a_l denote an arbitrary action in $A(I_i)$, and let $U_i(h_k, a_l)$ denote the expected utility of player i induced by the strategy β conditional on h_k and a_l. For any $\mu'_{I_i} \in \mathrm{cl}(\mu_{I_i}(\beta))$ and $\sigma \in \Delta A(I_i)$, also define

$$U_i(\mu'_{I_i}, \sigma) := \sum_k \sum_l \mu'_{I_i}(h_k)\sigma(a_l)U_i(h_k, a_l).$$

Then the SPME strategy $\beta_{I_i} \equiv \beta_i(I_i)$ satisfies

$$\beta_{I_i} \in \arg \max_{\sigma \in \Delta A(I_i)} \min_{\mu'_{I_i} \in \mathrm{cl}(\mu_{I_i}(\beta))} U_i(\mu'_{I_i}, \sigma).$$

If $\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))$ denotes the convex hull of $\mathrm{cl}(\mu_{I_i}(\beta))$, we must have

$$\min_{\mu'_{I_i} \in \mathrm{cl}(\mu_{I_i}(\beta))} U_i(\mu'_{I_i}, \sigma) = \min_{\mu'_{I_i} \in \mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))} U_i(\mu'_{I_i}, \sigma) \tag{1}$$

for every $\sigma \in \Delta A(I_i)$. To see why, let $\mu^* \in \mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))$ be a solution to the second minimization problem above, so that $\mu^* = \alpha\mu' + (1-\alpha)\mu''$, with $\mu', \mu'' \in \mathrm{cl}(\mu_{I_i}(\beta))$. Then $U_i(\mu', \sigma) \geq U_i(\mu^*, \sigma)$, $U_i(\mu'', \sigma) \geq U_i(\mu^*, \sigma)$, and

$$U_i(\mu^*, \sigma) = \alpha U_i(\mu', \sigma) + (1 - \alpha)U_i(\mu'', \sigma),$$

which implies that the minima in equation (1) must be attained at an element of $\mathrm{cl}(\mu_{I_i}(\beta))$. It follows that

$$\max_{\sigma\in\Delta A(I_i)} \min_{\mu'_{I_i}\in\mathrm{cl}(\mu_{I_i}(\beta))} U_i(\mu'_{I_i},\sigma) = \max_{\sigma\in\Delta A(I_i)} \min_{\mu'_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))} U_i(\mu'_{I_i},\sigma),$$

and that β_{I_i} is also a solution to $\max_{\sigma\in\Delta A(I_i)} \min_{\mu'_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))} U_i(\mu'_{I_i},\sigma)$. Since both $\Delta A(I_i)$ and $\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))$ are compact and convex, and $U_i(\mu'_{I_i},\sigma)$ is linear in μ'_{I_i} and σ, we can apply the Minimax Theorem of Fan (1952) to conclude that

$$\max_{\sigma\in\Delta A(I_i)} \min_{\mu'_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))} U_i(\mu'_{I_i},\sigma) = \min_{\mu'_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))} \max_{\sigma\in\Delta A(I_i)} U_i(\mu'_{I_i},\sigma),$$

and that there exist solutions to the corresponding maxmin and minmax problems. We know that β_{I_i} solves the maxmin problem. If we let $\gamma_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))$ denote a solution to the minmax problem, then $(\gamma_{I_i},\beta_{I_i})$ must be a saddle point of U_i (see Bertsekas et al., 2003, pp. 131-132), and hence

$$U_i(\gamma_{I_i},\sigma) \le U_i(\gamma_{I_i},\beta_{I_i}) \le U_i(\mu'_{I_i},\beta_{I_i}), \tag{2}$$

for all $\sigma\in\Delta A(I_i)$ and $\mu'_{I_i}\in\mathrm{co}(\mathrm{cl}(\mu_{I_i}(\beta)))$. The saddle point equation (2) implies in particular that β_{I_i} is sequentially rational given the precise beliefs γ_{I_i}. Hence, if $\beta_{-i}[I_i]=0$, and thus I_i lies off the equilibrium path, we can use the beliefs γ_{I_i} to define a belief system that yields (β,γ) as a WPBE.[18] □

Clearly, since any SE is also a WPBE, an SPME or PME strategy profile may not always yield an SE. The following example shows that even when the assumption of Proposition 7 is satisfied, so every SPME or PME yields a WPBE, this may not be the case for an SE:

[18] The saddle point equation (2) also shows, as would be expected, that in cases where β_{I_i} assigns strictly positive probability to more than one action in $A(I_i)$, the corresponding precise beliefs γ_{I_i} must make player i indifferent between all those those actions to which β_{I_i} assigns strictly positive probability. To see this, assume that β_{I_i} only assigns strictly positive probability to actions a_l and a_m, and denote the strategies in $\Delta A(I_i)$ that assign probability 1 to these actions by δ_{a_l} and δ_{a_m}, respectively. Then the saddle point equation implies that $U_i(\gamma_{I_i},\delta_{a_l}) \le U_i(\gamma_{I_i},\beta_{I_i})$ and $U_i(\gamma_{I_i},\delta_{a_m}) \le U_i(\gamma_{I_i},\beta_{I_i})$. Furthermore, we have

$$U_i(\gamma_{I_i},\beta_{I_i}) = \beta_{I_i}(a_l)U_i(\gamma_{I_i},\delta_{a_l}) + \beta_{I_i}(a_m)U_i(\gamma_{I_i},\delta_{a_m}),$$

which implies that all three expected utilities must be equal, so player i is indifferent between a_l and a_m at I_i, given precise beliefs γ_{I_i}.

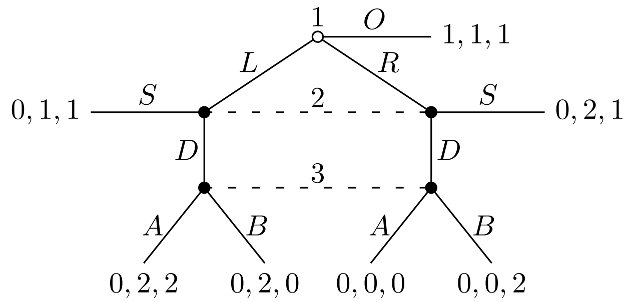

Figure 6: Example 5.

Example 5. Consider the game in Figure 6. Since player 1 has a dominant strategy to play O, he must assign probability 1 to O in any type of equilibrium. Let μ denote the probability assigned to history L at player 2's information set, and let δ denote the probability assigned to history LD at player 3's information set. The fact that 1 must play O in any equilibrium, implies that any ε-contamination of 1's strategy will yield conditional beliefs at the other players' information sets that are given by all $\mu \in [0,1]$ and $\delta \in [0,1]$. Since in addition, 2's payoffs do not depend on 3's actions, it follows that in any SPME or ε-PME (and hence PME), 2's and 3's equilibrium strategies must be optimal given beliefs defined by $\mu \in [0,1]$ and $\delta \in [0,1]$. Denoting the probability that 2 assigns to action S by p, his expected utility is given by

$$u_2(\beta_1, p) = \min_{\mu\in[0,1]} \{\mu[p + 2(1-p)] + (1-\mu)2p\} = \min_{\mu\in[0,1]} \{\mu(2-3p) + 2p\}$$
$$= \begin{cases} 2p, & \text{if } p \le \frac{2}{3}, \\ 2-p & \text{if } p > \frac{2}{3}. \end{cases}$$

Similarly, denoting the probability that 3 assigns to action A by q, yields the following expected utility at 3's information set:

$$u_3(\beta_1, q) = \min_{\delta\in[0,1]} \{\delta(2q) + (1-\delta)2(1-q)\} = \min_{\delta\in[0,1]} \{\delta(4q-2) + 2(1-q)\}$$
$$= \begin{cases} 2q, & \text{if } q \le \frac{1}{2}, \\ 2-2q & \text{if } q > \frac{1}{2}. \end{cases}$$

It follows that the game has a unique SPME, ε-PME and PME, where 1 plays O, 2 sets $p = \frac{2}{3}$, and 3 sets $q = \frac{1}{2}$. Clearly, this strategy profile constitutes a NE. To show that it also yields a WPBE, we must derive a belief system $(\bar{\mu}, \bar{\delta})$, consisting of precise beliefs, under which players 2 and 3 are indifferent between the actions available at their respective information sets (since they randomize in the unique SPME/PME). If $\bar{\mu}$ denotes the probability assigned to history L by player 2, then 2 is indifferent between S and D if

$$u_2(S) = \bar{\mu} + 2(1 - \bar{\mu}) = 2\bar{\mu} = u_2(D) \;\; \Leftrightarrow \;\; \bar{\mu} = \frac{2}{3}.$$

Similarly, if $\bar{\delta}$ denotes the probability assigned to history LD by player 3, then 3 is indifferent between A and B if

$$u_3(A) = 2\bar{\delta}3 = 2(1 - \bar{\delta}) = u_3(B) \;\; \Leftrightarrow \;\; \bar{\delta} = \frac{1}{2}.$$

Since WPBE imposes no restrictions on beliefs off the equilibrium path, these values for $\bar{\mu}$ and $\bar{\delta}$ yield a belief system under which the SPME/PME strategy profile is sequentially rational, and hence defines a WPBE. The same strategy profile only defines an SE if the same beliefs $\bar{\mu}$ and $\bar{\delta}$ are also consistent in the sense of SE. However, if $\bar{\mu}^n$ and $\bar{\delta}^n$ denote corresponding beliefs of players 2 and 3 that are derived from a completely mixed strategy profile β^n using Bayes' rule, then it must be that $\bar{\mu}^n = \bar{\delta}^n$ for all n, and hence, it is impossible for $\bar{\mu}^n$ to converge to $\frac{2}{3}$, while at the same time $\bar{\delta}^n$ converges to $\frac{1}{2}$. It follows that no consistent belief system exists under which the SPME/PME strategy profile is sequentially rational, and thus, it will not be played as part of any SE. Note also that the game has a continuum of SE that are induced by all belief systems $(\bar{\mu}, \bar{\delta})$ with $\bar{\mu} = \bar{\delta}$, in contrast to having a unique SPME, ε-PME and PME. $\lhd$

Together with Example 3, which shows that strategy profiles corresponding to WPBE or SE do not necessarily constitute an SPME or PME, our previous results imply that while SPME and PME can be viewed as refinements of NE and SPE, they do not yield refinements of WPBE or SE. Hence, with some abuse of notation, SPME/PME $\subsetneq$ NE/SPE, but SPME/PME $\nsubseteq$ WPBE/SE and WPBE/SE $\nsubseteq$ SPME/PME. Note however, that since every SPME or PME strategy profile is associated with a unique corresponding system of potentially ambiguous beliefs, the set of SPME or PME will in general be smaller than the set of WPBE or SE. This is the case in Example 3, where the set of PME does not include all WPBE and SE strategy profiles of the corresponding game. However, the excluded WPBE and SE only differ from the remaining ones in terms of equilibrium behavior off the equilibrium path. The following two examples show that PME and

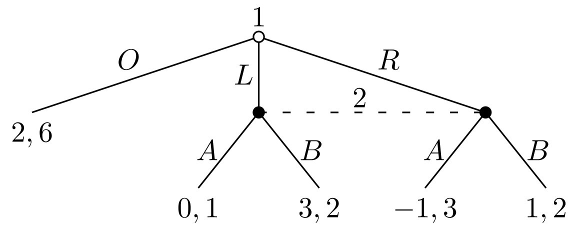

Figure 7: Example 6.

SPME can also exclude WPBE and SE that differ from non-excluded ones in terms of the behavior on the equilibrium path, and hence in terms of the outcomes of the game.

Example 6. Consider the game in Figure 7, and let μ denote the probability that player 2 assigns to history L at his information set. Then the set of SE (and WPBE) is given by (L, B) with $\mu = 1$, (O, A) with $\mu \in \left[0, \frac{1}{2}\right]$, and the strategy profile where 1 plays O, and 2 plays A with probability at least $\frac{1}{3}$, with $\mu = \frac{1}{2}$.

To derive the set of SPME, consider the possibility that 1 plays O in such an equilibrium. Then 2's information set is not reached, and thus 2 considers all $\mu \in [0, 1]$ as possible at his information set. Letting p denote the probability that 2 assigns to action A, 2's expected utility is given by

$$\begin{aligned} u_2(\beta_1, p) &= \min_{\mu \in [0,1]} \{\mu[p + 2(1-p)] + (1-\mu)[3p + 2(1-p)]\} \\ &= \min_{\mu \in [0,1]} \{\mu(-2p) + 2 + p\} = 2 - p. \end{aligned}$$

It follows that 2's best response is to set $p = 0$, and thus to play B, in which case 1 would deviate to L. Since L dominates R, 1 never plays R with strictly positive probability, and therefore the only SPME is (L, B) with $\mu = 1$. Hence, the set of SPME is strictly contained in the set of SE.

To derive the ε-PME of the game, denote the strategy of player 1 by

$$\beta_1 = (\beta_1(O), \beta_1(L), \beta_1(R)) = (1 - l - r, l, r),$$

denote the strategy of player 2 by

$$\beta_2 = (\beta_2(A), \beta_2(B)) = (p, 1 - p),$$

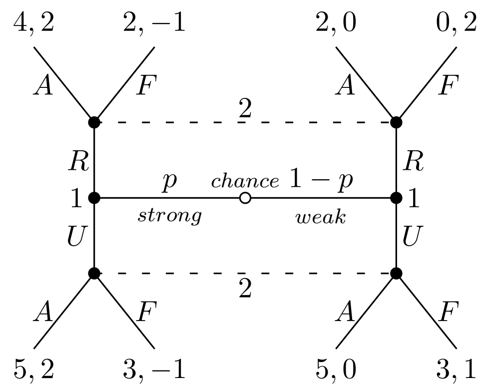

Figure 8: Example 7.

and let μ denote the probability player 2 assigns to history L according to some conditional belief. Then player 1's expected utility arising from β_1 and an ε-contamination of β_2 is given by

$$
\begin{aligned}
u_1(\beta_1, \beta_2) &= 2(1-l-r) + (1-\varepsilon)[l(1-p)3 - rp + r(1-p)] + \varepsilon \min_{q\in[0,1]}\{q(-3l-2r) + 3l + r\} \\
&= 2(1-l-r) + (1-\varepsilon)(1-p)3l + [(1-\varepsilon)(1-2p) - \varepsilon]r.
\end{aligned}
$$

It follows that 1's best response is to play L when

$$
p < \frac{1-3\varepsilon}{3(1-\varepsilon)}, \text{ or equivalently, when } 1-p > \frac{2}{3(1-\varepsilon)},
$$

to play O when the reverse inequalities hold, and mix over L and O using any distribution when the inequalities hold with equality. Therefore, when looking for any ε-PME, we can restrict the analysis to the case where $\beta_1(R) = 0$. Then if $\beta_1(O) = 1$, the set $\mu(\beta_1^\varepsilon)$ of conditional beliefs of player 2 is given by $\Delta\{L, R\}$ for any $\varepsilon > 0$, which yields B as a best response for player 2. If $\beta_1(L) > 0$, the minimum of player 2's expected utility over $\mu(\beta_1^\varepsilon)$ is still attained when $\mu = 1$, and thus 2's best response remains to play B, or equivalently, to set $p = 0$. Hence, in any ε-PME, player 2 always plays B, and as long as $\varepsilon < 1/3$, this implies that player 1 plays L. It follows that (L, B) is the unique ε-PME for $\varepsilon < 1/3$, and is therefore also the unique PME (which in this example is equal to the SPME of the game). $\triangleleft$

Example 7. Consider next the signaling game described in Figure 8,[19] and assume that p is common knowledge and belongs to $(0, 1)$, so chance is not viewed as a player in the game. Hence, in order to derive the players' beliefs, only the actions of their opponents are contaminated, but not the move of nature. Note also that a weak type of player 1 has a dominant strategy to play U, independently of whether the strategies of player 2 at any of his information sets are contaminated or not, and hence such a weak type plays U for all equilibrium concepts we consider. The game has a pure strategy separating equilibrium that is a SE, an SPME, and an ε-PME for small enough ε, and hence a PME. In this separating equilibrium, a strong type of player 1 plays R, and player 2 plays A at his top information set and F at his bottom information set.

To analyze the pooling equilibria, consider first the WPBE/SE, and note that if $p > \frac{1}{4}$, 2 always plays A at his bottom information set if he assigns probabilities p and $1 - p$ to the respective histories in this information set. Thus, there exist a continuum of pooling WPBE/SE where both types of 1 play U, and 2 is allowed to have any beliefs and corresponding optimal actions at his top information set. If $p < \frac{1}{4}$, 2 plays F at his bottom information set in a pooling WPBE/SE, and therefore sequential rationality of U for a strong type of player 1 requires 2 to play F at his top information set. This is optimal as long as the history (strong, R) is assigned probability at most $\frac{1}{4}$. All such beliefs yield a continuum of pooling WPBE/SE where 2 plays F at both his information sets. In contrast, we show in Appendix B that for the case of $p < \frac{1}{4}$, no corresponding pooling SPME or PME exist. $\triangleleft$

We close this section with a brief exploration of the relation between SPME and PME. If β is an SPME in pure strategies that is strict, in the sense that the strategy prescribed in equilibrium yields a strictly higher expected payoff than any alternative pure or mixed strategy, then the fact that β is optimal given $\mu_{I_i}(\beta)$, and that $\mu_{I_i}(\beta) = \lim_{\varepsilon \to 0} \mu_{I_i}(\beta^\varepsilon)$, implies that the strategies prescribed by β must also be optimal if beliefs are given by $\mu_{I_i}(\beta^\varepsilon)$ for small enough ε. This yields the following result:

Proposition 8. *Every strict SPME in pure strategies is a PME.*

The following example shows that the conditions from Proposition 8 are sufficient, but not necessary, for an SPME to be a PME. The associated game possesses an SPME in

[19] This example is borrowed from Osborne (2004).

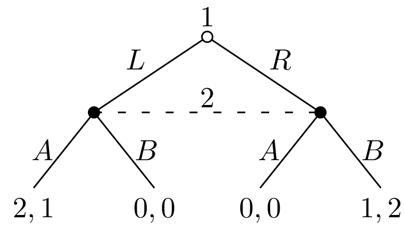

Figure 9: Example 8.

mixed strategies that is not strict, but is still a PME:

Example 8. Consider the game illustrated in Figure 9, which is a standard "battle of the sexes" game in extensive form. Since there are no information sets that lie off the equilibrium path, beliefs $\mu_{I_i}(\beta)$ induced by the limit of ε-contaminations of any strategy β are precise, and given by the corresponding Bayesian updates. Hence, the sets of SE and SPME coincide, and are given by the Nash equilibria of the game. In particular, there exists an SPME in mixed strategies where player 1 plays L with probability $\frac{2}{3}$ and player 2 plays A with probability $\frac{1}{3}$. We will show that this particular strategy profile is the limit of a sequence of ε-PME as $\varepsilon \to 0$, and is thus also a PME.

Denote the probability of playing L for player 1 by p, and the probability of playing A for player 2 by q. Then the expected utility of player 2, as induced by an ε-contamination of 1's strategy p is given by

$$u_2(p,q) = \min_{\mu \in [(1-\varepsilon)p,(1-\varepsilon)p+\varepsilon]} \{\mu q + (1-\mu)2(1-q)\} = \min_{\mu \in [(1-\varepsilon)p,(1-\varepsilon)p+\varepsilon]} \{\mu(3q-2) + 2(1-q)\}$$
$$= \begin{cases} [3(1-\varepsilon)p - 2 + 3\varepsilon]q + 2[1-(1-\varepsilon)p - \varepsilon], & \text{if } q \leq \frac{2}{3}, \\ [3(1-\varepsilon)p - 2]q + 2[1-(1-\varepsilon)p] & \text{if } q > \frac{2}{3}, \end{cases}$$

and the expected utility of player 1, as induced by an ε-contamination of 2's strategy q is given by

$$u_1(p,q) = \min_{\delta \in [(1-\varepsilon)q,(1-\varepsilon)q+\varepsilon]} \{\delta 2p + (1-\delta)(1-p)\} = \min_{\delta \in [(1-\varepsilon)q,(1-\varepsilon)q+\varepsilon]} \{\delta(3p-1) + (1-p)\}$$
$$= \begin{cases} [3(1-\varepsilon)q - 1 + 3\varepsilon]p + [1-(1-\varepsilon)q - \varepsilon], & \text{if } p \leq \frac{1}{3}, \\ [3(1-\varepsilon)q - 1]p + [1-(1-\varepsilon)q] & \text{if } p > \frac{1}{3}, \end{cases}$$

We get the following best responses:

$$BR_2(p) = \begin{cases} q = 0, & \text{if } p < \frac{2-3\varepsilon}{3(1-\varepsilon)}, \\ q \in \left[0, \frac{2}{3}\right], & \text{if } p = \frac{2-3\varepsilon}{3(1-\varepsilon)}, \\ q = \frac{2}{3}, & \text{if } p \in \left(\frac{2-3\varepsilon}{3(1-\varepsilon)}, \frac{2}{3(1-\varepsilon)}\right), \\ q \in \left[\frac{2}{3}, 1\right], & \text{if } p = \frac{2}{3(1-\varepsilon)}, \\ q = 1, & \text{if } p > \frac{2}{3(1-\varepsilon)}, \end{cases}$$

$$BR_1(q) = \begin{cases} p = 0, & \text{if } q < \frac{1-3\varepsilon}{3(1-\varepsilon)}, \\ p \in \left[0, \frac{1}{3}\right], & \text{if } q = \frac{1-3\varepsilon}{3(1-\varepsilon)}, \\ p = \frac{1}{3}, & \text{if } q \in \left(\frac{1-3\varepsilon}{3(1-\varepsilon)}, \frac{1}{3(1-\varepsilon)}\right), \\ p \in \left[\frac{1}{3}, 1\right], & \text{if } q = \frac{1}{3(1-\varepsilon)}, \\ p = 1, & \text{if } q > \frac{1}{3(1-\varepsilon)}. \end{cases}$$

In addition to yielding two ε-PME where both p and q are equal to either 0 or 1, when $\varepsilon < \frac{1}{3}$, these best responses also show that there exists an ε-PME with

$$p = \frac{2-3\epsilon}{3(1-\varepsilon)}, \text{ and } q = \frac{1}{3(1-\varepsilon)}.$$

Clearly, when $\varepsilon \to 0$, these ε-PME converge to the mixed strategy SPME identified previously, which is thus also a PME. ◁

Given a strategy profile β and $\varepsilon \in (0,1)$, let $\mu(\beta^\varepsilon)$ denote a profile of (possibly ambiguous) beliefs across all information sets corresponding to a game, as induced by an ε-contamination of β. Thus, $\mu(\beta^\varepsilon)$ collects the beliefs $\mu_{I_i}(\beta^\varepsilon)$ across all information sets I_i and players i. Furthermore, let $\mu(\beta) = \lim_{\varepsilon\to 0} \mu(\beta^\varepsilon)$.

Proposition 9. *Let β be a PME, with $\beta_{[\varepsilon]}$ denoting a sequence of ε-PMEs converging to β. If $\mu(\beta) = \lim_{\varepsilon\to 0} \mu(\beta^\varepsilon_{[\varepsilon]})$, then β is also an SPME.*[20]

Proof. If $\beta_{[\varepsilon]}$ is optimal given $\mu(\beta^\varepsilon_{[\varepsilon]})$, and $\beta_{[\varepsilon]} \to \beta$, then $\mu(\beta) = \lim_{\varepsilon\to 0} \mu(\beta^\varepsilon_{[\varepsilon]})$ implies that β is optimal given $\mu(\beta)$, and thus β is an SPME. □

[20] Note that $\lim_{\varepsilon\to 0} \mu(\beta^\varepsilon_{[\varepsilon]})$ denotes the limit of the beliefs induced by the ε-contaminations of the strategy profiles $\beta_{[\varepsilon]}$, whereas $\mu(\beta) = \lim_{\varepsilon\to 0} \mu(\beta^\varepsilon)$ denotes the limit of the beliefs induced by ε-contaminations of the fixed strategy profile β.

6 Extensions and generalizations

6.1 Alternative updating rules or contaminations

Our results so far have assumed that beliefs about past actions at every information set are derived based on full prior-by-prior Bayesian updating whenever possible. While this approach yields the most straightforward updating rule in the context of ambiguous beliefs modeled by sets of distributions, alternative updating methods are easy to define. For example, the best-known alternative to full Bayesian updating is the maximum likelihood rule, which in our setting would result in only updating those unconditional beliefs under which the probability of reaching the information set in question is maximized. More generally, an arbitrary abstract updating rule can be defined by specifying for each information set, a particular subset of unconditional beliefs corresponding to an ε-contamination of opponents' strategies, whose Bayesian updates are then assumed to define the beliefs of the player moving at that information set.[21] Given a strategy profile β, and a corresponding ε-contamination β^ε, we can denote the beliefs at an information set I_i that result from such an arbitrary updating rule by $\hat{\mu}_{I_i}(\beta^\varepsilon)$. The assumption that these beliefs are defined by only updating a subset of the unconditional beliefs defined by β^ε then implies that $\hat{\mu}_{I_i}(\beta^\varepsilon) \subset \mu_{I_i}(\beta^\varepsilon) \subset \Delta A(I_i)$. If an abstract updating rule excludes a particular strategy $\beta'_j \subset \beta^\varepsilon_j$ in the computation of beliefs $\hat{\mu}_{I_i}(\beta^\varepsilon)$ at player i's information set I_i, assume that following I_i, player i also excludes the possibility that j plays according to the strategy β'_j. Assume further that every such rule yields a non-empty set of beliefs at every I_i. If we then define

$$\hat{\mu}_{I_i}(\beta) := \limsup_{\varepsilon \to 0} \hat{\mu}_{I_i}(\beta^\varepsilon),$$

our prior definitions of ε-PME, PME and SPME apply directly to the case of an arbitrary updating rule, by replacing $\mu_{I_i}(\beta^\varepsilon)$ and $\mu_{I_i}(\beta)$ with $\hat{\mu}_{I_i}(\beta^\varepsilon)$ and $\hat{\mu}_{I_i}(\beta)$, respectively.

A different generalization of our model can be obtained by changing the nature of the contaminations, instead of (or in addition to) the updating rule. The type of ε-contaminations used so far seem to most appropriately capture the intuition that if mistakes are arbitrary and involuntary, then we should be agnostic about the nature of trembles that we can expect, and hence view all distributions over actions as equally possible

[21] General definitions of such updating rules in dynamic choice problems can be found in Gilboa and Schmeidler (1993) and Hanany and Klibanoff (2007).

in case a mistake occurs. We could however also consider alternative ε-contaminations, where mistakes still occur with small probability ε, but where the possible distributions that may generate such mistakes are restricted to subsets of the feasible strategy spaces. More generally, we could define a perturbation of a strategy not by an ε-contamination, but by the closure of an arbitrary open neighborhood containing the intended strategy. Denote a sequence of such arbitrary perturbations of a strategy profile β by $\beta^{[n]}$, where n is an index such that $\lim_{n\to\infty} \beta^{[n]} = \{\beta\}$. The resulting beliefs at an information set I_i can then be derived using either full Bayesian updating, or some alternative updating rule. If $\mu_{I_i}(\beta^{[n]})$ or $\hat{\mu}_{I_i}(\beta^{[n]})$ represent corresponding beliefs, we can let

$$\mu_{I_i}(\beta) := \limsup_{n\to\infty} \mu_{I_i}(\beta^{[n]}), \text{ or } \hat{\mu}_{I_i}(\beta) := \limsup_{n\to\infty} \hat{\mu}_{I_i}(\beta^{[n]}),$$

and our prior definitions of ε-PME, PME and SPME apply again using the appropriate adjustments.

Note that for certain simple games, such as the game analyzed in Example 3, any sequence of generalized contaminations that are updated using full Bayesian updating, can be equivalently replicated using a sequence of ε-contaminations (indexed by $\varepsilon \to 0$) that are updated using a more general updating rule, and vice versa. Furthermore, as shown in Appendix C, for this particular example the sets of SPME and PME resulting from any arbitrary such generalization are identical to the ones derived using the initial assumption of full Bayesian updating of beliefs derived from standard ε-contaminations.

For a generalized model yielding beliefs $\hat{\mu}_{I_i}(\beta^{[n]})$ at an information set I_i, which are derived from a perturbation $\beta^{[n]}$ using an arbitrary updating rule, define, analogously to the definition in the proof of Theorem 1, a set

$$\hat{\Gamma}_{I_i}(\beta^{[n]}) := \{(\beta'_{-i}, \mu'_{I_i}) \,|\, \mu'_{I_i} \in \hat{\mu}_{I_i}(\beta^{[n]}), \text{ and } \mu'_{I_i} \text{ is derived from } \beta'_{-i} \text{ using Bayesian updating}\}.$$

If the perturbations $\beta^{[n]}$ and the updating rule yielding $\hat{\mu}_{I_i}(\beta^{[n]})$ satisfy the property that $\mathrm{cl}(\hat{\Gamma}_{I_i}(\beta^{[n]}))$ is continuous in β for any particular n, then the proof of Theorem 1 implies the existence of an ε-PME (where the ε now corresponds to the index n) for this model. Note that while this continuity requirement will hold for a large class of perturbations and updating rules, it might not hold for all games if the maximum likelihood updating rule is used, since the set of maximizers at some information sets may not always be continuous in β.[22] There exist however games where the beliefs derived using the maximum likelihood updating rule are identical to the beliefs derived using full Bayesian updating,[23]so beliefs, and hence equilibrium outcomes, will be the same under both updating rules. Finally, the assumptions we made also imply that Lemma 1 holds for such general models, so $\hat{\mu}_{I_i}(\beta)$ is given by the Bayesian update of β_{-i} whenever $\beta_{-i}[I_i] > 0$. As a consequence, the proofs of Propositions 5 and 7 apply, and hence the associated PME and SPME are also NE and SPE.

6.2 Dynamic consistency of equilibrium strategies

We now discuss the possibility of attaining dynamic consistency in our model. The decision theory literature presents two alternative approaches that can guarantee dynamic consistency with maxmin preferences: Epstein and Schneider (2003) maintain full Bayesian updating, but restrict the class of feasible priors, whereas Hanany and Klibanoff (2007) allow arbitrary sets of priors and restrict the updating rule. Before addressing how these approaches might be used in our modeling framework, consider first our initial setting where beliefs are derived from ε-contaminations through full Bayesian updating. While preferences in this model may in general be dynamically inconsistent, this is not an issue if we are only concerned with the dynamic consistency of SPME or PME strategies—since any SPME or PME β yields precise beliefs derived through Bayesian updating at information sets I_i for which $\beta_{-i}[I_i] > 0$, as shown in the proof of Proposition 5, and equilibrium strategies are optimal given these beliefs, an application of the one-shot-deviation principle analogous to the proof of Proposition 5, implies that equilibrium strategies across information sets with $\beta_{-i}[I_i] > 0$ are also optimal from an ex-ante point of view, and are thus dynamically consistent. Furthermore, since information sets with $\beta_{-i}[I_i] = 0$ are reached with probability zero from an ex-ante viewpoint, any strategies chosen at such information sets are optimal ex-ante, and hence do not contradict dynamic

[22]It is (relatively) easy to construct games that generate such discontinuities. For instance, in Example 4, if we split player 2's information set, so that he observes 1's move, and 3's beliefs are derived using maximum likelihood updating, then 3's beliefs are discontinuous at any strategy profile where 2 assigns equal probabilities to D at both information sets.

[23]For example, this is the case if unconditional beliefs have a triangular shape as in Figure 2, in which case the hypotenuse of the triangle represents the set of relevant unconditional probabilities that are updated under the maximum likelihood rule.

consistency. It follows that every SPME or PME strategy profile is optimal both from an ex-ante point of view, and from an interim point of view conditional on each information set, and is therefore dynamically consistent.

Note however that an analogous dynamic consistency property may not hold for strategy profiles that define an ε-PME. Since any ε-PME requires strategies to be sequentially rational, i.e., conditionally optimal at each information set given beliefs derived from ε-contaminations and given the consistent planning assumption, such equilibrium strategies may not be optimal as ex-ante plans of action, which just reflects the potential dynamic inconsistency of maxmin preferences. However, this is mitigated by the fact that even though ε-PME strategies may not be optimal ex-ante, the potential ex-ante losses from such strategies are small when ε is small, and are converging to zero as $\varepsilon \to 0$. To see this, let β denote a PME that is the limit of a sequence of ε-PMEs $\beta_{[\varepsilon]}$. As noted in the proof of Proposition 5, along the equilibrium path induced by β, the PME strategy β_i is optimal for player i given precise beliefs derived from β_{-i} using Bayesian updating, and is furthermore ex-ante optimal as noted above. Since $\beta_i = \lim_{\varepsilon\to 0} \beta_{i[\varepsilon]}$, this implies that for small values of ε, $\beta_{i[\varepsilon]}$ is approximately optimal from an ex-ante viewpoint given beliefs derived from β_{-i}. Now if we consider player i's ex-ante beliefs as induced by the ε-contamination $\beta^{\varepsilon}_{-i[\varepsilon]}$, these beliefs place high probability on the opponents' strategy profile $\beta_{-i[\varepsilon]}$, which implies that $\beta_{i[\varepsilon]}$ is approximately ex-ante optimal since $\beta_{-i[\varepsilon]} \to \beta_{-i}$.

Example 9. To illustrate the properties discussed above, consider again the game from Example 3. Clearly, the strategy profile defining the PME is dynamically consistent, in the sense that if player 2 were able to commit to a strategy ex-ante, before player 1 chooses an action, then given 1's PME strategy $\beta_1(O) = 1$, any strategy of 2 is optimal, in particular his PME strategy $\beta_2(A) = \frac{7}{12}$. Furthermore, given 1's ε-PME strategy, we can derive 2's ex-ante expected utility as a function of the probability p assigned to action A, which yields

$$u_2^{EA}(\beta_{[\varepsilon]1}, p) = \begin{cases} 2(1-\varepsilon), & \text{if } p \leq \frac{2}{3}, \\ 2 - 3\varepsilon p, & \text{if } p > \frac{2}{3}. \end{cases}$$

It follows that in this particular case, 2's ε-PME strategy is also optimal at an ex-ante stage, and is therefore dynamically consistent. However, this property does not always hold. If we change player 2's payoff from history O to -1, then the ε-PME, and hence

PME, of the game do not change, and thus the PME strategy profile is still dynamically consistent. But given 1's ε-PME strategy, 2's ex-ante expected utility now becomes

$$u_2^{EA}(\beta_{[\varepsilon]1}, p) = 6\varepsilon - 1 - 2\varepsilon p,$$

which implies that $p = 0$ is uniquely optimal ex-ante, and hence that the ε-PME strategy $\beta_{[\varepsilon]2}(A) = \frac{7}{(1-\varepsilon)12}$ is not dynamically consistent. Player 2's ex-ante losses associated with his ε-PME strategy are given by $\frac{7\varepsilon}{6(1-\varepsilon)}$, a value that is small as long as ε is small, and that converges to zero as $\varepsilon \to 0$. $\triangleleft$

If the approaches of Epstein and Schneider (2003) or Hanany and Klibanoff (2007) apply in our game-theoretic setting, they should allow us to guarantee dynamic consistency across our model, and not just for equilibrium strategies. While their methods work for certain games, their application presents various difficulties, which we now discuss.

Epstein and Schneider (2003) show that maxmin preferences are dynamically consistent with full Bayesian updating if the initial set of priors satisfies a "rectangularity" assumption with respect to the decision-maker's information filtration. Roughly, rectangularity means that the initial set of priors can be constructed by recursively combining all corresponding conditional and marginal probabilities at each stage of the information filtration. In our game-theoretic setting, as the set of priors corresponding to each player's information filtration is derived from the ε-contaminations of his opponents' strategies, to attain rectangularity, we would need to consider a generalized set of contaminations (as described in Section 6.1) that is chosen in a way that yields rectangularity for each player. While this approach works in limited settings, we show in Aryal and Stauber (2013) that even in a simple three-player game where each player only moves once, it may be impossible to attain the rectangularity property by suitably choosing the shape of the contaminations. The basic intuition for this result is that since each player's information filtration is induced by the structure of the game relative to his own information partition, for the case of more than two players, the requirements on player 1's contamination that are needed to guarantee rectangularity for player 2's information filtration may be incompatible with those needed to guarantee rectangularity for player 3's information filtration. Hence, unless players 2 and 3 are allowed to have distinct beliefs about the contamination of 1's strategy, this contamination cannot be defined in a way that guarantees rectangularity for both of his opponents.

Hanany and Klibanoff (2007) (henceforth, HK) show that dynamic consistency can be achieved with maxmin preferences through the choice of an appropriate updating rule, which may depend on the given dynamic choice problem and the optimal ex-ante choice. Applying the updating rules described by HK in our game-theoretic setting is not straightforward, mainly because the strong non-null assumptions of HK, which require that each conditioning event is assigned strictly positive probability by every element of the prior set, are not always satisfied in our model. In an extensive game, the relevant conditioning events are defined by the players' information sets. If an information set I_i of player i is a "null event" in a sense that will be made more precise shortly, then every action at that information set is optimal ex-ante, even an action that is dominated across all histories in the information set—clearly, no updating rule exists under which such a dominated action is optimal. Whether an information set I_i is null will depend on the given strategy profile: For example, if an ex-ante optimal strategy of player i implies that I_i will never be reached, independently of his opponents' actions, I_i is null;[24] alternatively, I_i is null given an ε-contamination of his opponents' strategies, if the distribution in the induced beliefs of player i that attains an ex-ante maxmin assigns probability zero to I_i. If an information set is null given a strategy profile β, we cannot guarantee the existence of an updating rule for that information set which is dynamically consistent in the sense that any ex-ante optimal strategy is also conditionally optimal. However, since any conditionally optimal strategy at such an information set is also optimal ex-ante, no updating rule contradicts dynamic consistency. We can then apply the method of HK to choose dynamically consistent updating rules only at non-null information sets. HK show that such rules exist, but that they would depend on the initial set of priors, and potentially require updating different subsets of priors at different information sets. Although as discussed in Section 6.1, our equilibrium definitions still apply with arbitrary updating rules, to attain dynamic consistency following HK, the required updating rules may need to vary with the chosen strategies.[25] As a consequence, in equilibrium, both the null events and updating rules would depend on equilibrium strategies, and hence

[24]Hence, in contrast to the two-stage decision-theoretic framework of HK, whether an information set I_i can be identified as a null event also depends on the strategy chosen by player i.

[25]For example, a dynamically consistent updating rule for a given strategy could use full Bayesian updating at null information sets, and use the "ambiguity maximizing" updating rule of HK at non-null information sets.

any dynamically consistent updating rule would be endogenously determined as part of the equilibrium. A full analysis along these lines—which would also need to analyze whether dynamically consistent updating rules yield the continuity properties required for equilibrium existence—would provide an interesting game-theoretic application of HK, but is beyond the scope of the present paper.

7 Conclusion

We introduce equilibrium notions that capture how ambiguity averse players may interact in an extensive game. The two notions of Strong Perfect Maxmin Equilibrium (SPME) and Perfect Maxmin Equilibrium (PME) can both be viewed as refinements of NE and SPE, if expected utility preferences are interpreted as a subset of maxmin preferences. Hence, our model seems to indicate that NE and SPE are still valid solution concepts in some environments with ambiguity averse players, if non-trivial ambiguous beliefs are suitably restricted to off-the-equilibrium-path information sets. However, the predictions of SPME and PME are distinct from those of WPBE and SE, in the sense that there exist WPBE and SE that do not yield SPME or PME, and there exist SPME and PME that do not correspond to any WPBE or SE. Hence, allowing for ambiguity averse players can yield behavior that differs from the predictions of belief-system-based refinements under standard expected utility maximization, even if the ambiguity only results from allowing for small probability errors in the implementation of unambiguous strategies, and we consider the limiting case where the probability of making such errors converges to zero.

In addition to the generalizations we already discussed, a few other modeling assumptions could easily be varied. A limited modification of our setup could be achieved by assuming that each player also considers the possibility that he himself might tremble and make mistakes. Since players have perfect recall about their own past actions, and future mistakes have small probability, such a change would not have resulted in a significant difference to our results. Considering different types of preferences with ambiguity would most likely make a significant difference, but such considerations would require extensive changes that fall outside the scope of the current paper.

Appendices

A Derivation of ε-PME for Example 3

To derive the ε-PME of the game, denote the strategy of player 1 by

$$\beta_1 = (\beta_1(O), \beta_1(L), \beta_1(R)) = (1 - l - r, l, r),$$

denote the strategy of player 2 by

$$\beta_2 = (\beta_2(A), \beta_2(B)) = (p, 1 - p),$$

and let μ denote the probability player 2 assigns to history L according to some conditional belief. Then player 1's expected utility arising from β_1 and an ε-contamination of β_2 is given by

$$\begin{aligned}
u_1(\beta_1, \beta_2) &= (1 - \varepsilon)[7(1 - l - r) + 12lp + 12r(1 - p)] + \varepsilon \min_{q\in[0,1]} \{7(1 - l - r) + 12lq + 12r(1 - q)\} \\
&= 7(1 - l - r) + (1 - \varepsilon)12lp + (1 - \varepsilon)12r(1 - p) + \varepsilon \min_{q\in[0,1]} \{12q(l - r) + 12r\} \\
&= \begin{cases} 7 + [(1 - \varepsilon)12p - 7]l + [5 - 12(1 - \varepsilon)p]r, & \text{if } l \geq r, \\ 7 + [(1 - \varepsilon)12p + \varepsilon 12 - 7]l + [(1 - \varepsilon)12(1 - p) - 7]r, & \text{if } l < r. \end{cases}
\end{aligned}$$

Note that for every p, $u_1(\beta_1, \beta_2) = u_1((l, r), p)$ is a continuous and piecewise linear function of (l, r). The domain of this function is the "Machina triangle" defined by the convex hull of $\{(0, 0), (1, 0), (0, 1)\}$, and the function is linear on the two subsets of this domain defined by the points $\{(0, 0), (1, 0), (1/2, 1/2)\}$ and $\{(0, 0), (1/2, 1/2), (0, 1)\}$, respectively (see Figure 10). Hence, since player 1 maximizes a function that is continuous and piecewise linear on the union of two convex polytopes, the maximum over (l, r) must be attained at the extreme points of these two polytopes. We get

$$u_1((0, 0), p) = 7, \qquad u_1((1, 0), p) = (1 - \varepsilon)12p,$$
$$u_1((1/2, 1/2), p) = 6, \qquad u_1((0, 1), p) = (1 - \varepsilon)12(1 - p).$$

Comparing the payoffs resulting from these four points implies that when $\varepsilon > \frac{5}{12}$, playing O is a unique best response for player 1 for every p, and that when $\varepsilon \in \left(0, \frac{5}{12}\right)$,

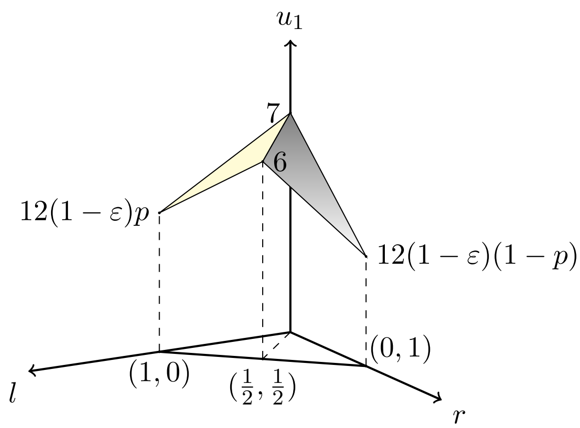

Figure 10: Graph of $u_1((l,r),p)$.

the best response of player 1 as a function of p is given by

$$BR_1(p) = \begin{cases} \{R\}, & \text{if } p \in \left[0, 1 - \frac{7}{(1-\varepsilon)12}\right), \\ \Delta\{R,O\}, & \text{if } p = 1 - \frac{7}{(1-\varepsilon)12}, \\ \{O\}, & \text{if } p \in \left(1 - \frac{7}{(1-\varepsilon)12}, \frac{7}{(1-\varepsilon)12}\right), \\ \Delta\{O,L\}, & \text{if } p = \frac{7}{(1-\varepsilon)12}, \\ \{L\}, & \text{if } p \in \left(\frac{7}{(1-\varepsilon)12}, 1\right]. \end{cases}$$

To characterize the best response of player 2, note that the extreme points of 2's conditional beliefs induced by a strategy (l,r) of player 1 are derived from the Bayesian updates of the unconditional probabilities given by $(1-\varepsilon)(l,r)+\varepsilon(1,0)$ and $(1-\varepsilon)(l,r)+\varepsilon(0,1)$. It follows that the probability μ that 2 assigns to the history L ranges from

$$\frac{(1-\varepsilon)l}{(1-\varepsilon)l+(1-\varepsilon)r+\varepsilon} \text{ to } \frac{(1-\varepsilon)l+\varepsilon}{(1-\varepsilon)l+(1-\varepsilon)r+\varepsilon}.$$

Denoting this set of conditional beliefs by $\mu^\varepsilon(l,r)$, we get

$$\begin{aligned} u_2((l,r),p) &= \min_{\mu\in\mu^\varepsilon(l,r)} \{\mu[p+2(1-p)]+(1-\mu)2p\} \\ &= \min_{\mu\in\mu^\varepsilon(l,r)} \{\mu(2-3p)+2p\} \\ &= \begin{cases} \frac{2(1-\varepsilon)l}{(1-\varepsilon)(l+r)+\varepsilon} + \left[2 - \frac{3(1-\varepsilon)l}{(1-\varepsilon)(l+r)+\varepsilon}\right]p, & \text{if } p \le \frac{2}{3}, \\ \frac{2(1-\varepsilon)l+2\varepsilon}{(1-\varepsilon)(l+r)+\varepsilon} + \left[2 - \frac{3(1-\varepsilon)l+3\varepsilon}{(1-\varepsilon)(l+r)+\varepsilon}\right]p, & \text{if } p > \frac{2}{3}. \end{cases} \end{aligned}$$

The best response of player 2 as a function of (l, r) is then

$$BR_2(l,r) = \begin{cases} p = 0, & \text{if } 2(1-\varepsilon)r < (1-\varepsilon)l - 2\varepsilon, \\ p \in \left[0, \frac{2}{3}\right], & \text{if } 2(1-\varepsilon)r = (1-\varepsilon)l - 2\varepsilon, \\ p = \frac{2}{3}, & \text{if } 2(1-\varepsilon)r > (1-\varepsilon)l - 2\varepsilon \text{ and } 2(1-\varepsilon)r < (1-\varepsilon)l + \varepsilon, \\ p \in \left[\frac{2}{3}, 1\right], & \text{if } 2(1-\varepsilon)r = (1-\varepsilon)l + \varepsilon, \\ p = 1, & \text{if } 2(1-\varepsilon)r > (1-\varepsilon)l + \varepsilon. \end{cases}$$

To find the set of ε-PME, consider the possibility that 1 plays O in such an equilibrium, so $l = r = 0$. Then 2's best response is to set $p = \frac{2}{3}$, in which case playing O is a best response for 1 if $\varepsilon \geq \frac{1}{8}$, but not if $\varepsilon < \frac{1}{8}$, in which case L yields a higher payoff. Thus, we get an ε-PME where 1 plays O and 2 sets $p = \frac{2}{3}$ as long as $\varepsilon \geq \frac{1}{8}$. Assume from here on that $\varepsilon < \frac{1}{8}$, and consider the possibility that 1 plays L, so $l = 1$ and $r = 0$. Then as long as $\varepsilon < \frac{1}{3}$, $p = 0$ is a best response for 2, in which case 1 would prefer to play R. Similarly, consider the possibility that 1 plays R, so $l = 0$ and $r = 1$. Then as long as $\varepsilon < \frac{2}{3}$, $p = 1$ is a best response for 2, in which case 1 would prefer to play L. It follows that if $\varepsilon < \frac{1}{8}$, 1 cannot play a pure strategy in any ε-PME, and therefore, p must be equal to either $1 - \frac{7}{(1-\varepsilon)12}$ or $\frac{7}{(1-\varepsilon)12}$. Since both these terms are less than $\frac{2}{3}$ when $\varepsilon < \frac{1}{8}$, the equilibrium values of l and r must satisfy $2(1-\varepsilon)r = (1-\varepsilon)l - 2\varepsilon$. Player 1's best response shows that he will only randomize over R and O, or over O and L. Setting $l = 0$ implies that $(1-\varepsilon)r = -\varepsilon$, which can never hold. Hence, it must be the case that $l > 0$ and $r = 0$, which implies that for $\varepsilon \in \left(0, \frac{1}{8}\right)$, a (unique) ε-PME exists where

$$l = \frac{2\varepsilon}{(1-\varepsilon)}, \quad r = 0, \quad \text{and} \quad p = \frac{7}{(1-\varepsilon)12}.$$

B Non-existence of pooling SPME or PME for Example 7

Let μ denote the probability assigned by player 2 to the history (strong, R) at his top information set, let q and r denote the probabilities 2 assigns to action A at his top and bottom information sets, respectively, and let α denote the probability of playing R for a strong type of player 1. Consider the possibility of a pooling SPME where $\alpha = 0$, so both types of player 1 play U, and $p < \frac{1}{4}$, so player 2 plays F at his bottom information set. Then the beliefs corresponding to such an SPME at the top information set are given by all $\mu \in [0,1]$, and hence player 2's utility at the top information set as a function of his strategy q is

$$u_{2,\text{top}}(q) = \min_{\mu \in [0,1]} \{\mu(2q - (1-q)) + (1-\mu)2(1-q)\}$$
$$= \begin{cases} 3q - 1, & \text{if } q \leq \frac{3}{5}, \\ -2q + 2 & \text{if } q > \frac{3}{5}. \end{cases}$$

It follows that 2's optimal strategy is to set $q = \frac{3}{5}$, in which case the expected payoff of a strong type of player 1 from playing R is $\frac{16}{5} = 3.2$, which is greater than the payoff of 3 he receives from playing U. Hence, no pooling SPME exists when $p < \frac{1}{4}$.

To characterize the PME when $p < \frac{1}{4}$, note that since U always dominates R for a weak type of player 1, the beliefs of player 2 at his top information set, as induced by an ε-contamination of 1's strategy, are given by all μ in the interval

$$\mu^{\varepsilon}(\alpha) := \left[\frac{(1-\varepsilon)p\alpha}{(1-\varepsilon)p\alpha + \varepsilon(1-p)}, 1\right],$$

and hence, 2's expected utility at his top information set is given by

$$u_{2,\text{top}}(q) = \min_{\mu \in \mu^{\varepsilon}(\alpha)} \{\mu(2q - (1-q)) + (1-\mu)2(1-q)\}$$
$$= \begin{cases} 3q - 1, & \text{if } q \leq \frac{3}{5}, \\ \left[\frac{5(1-\varepsilon)p\alpha}{(1-\varepsilon)p\alpha + \varepsilon(1-p)} - 2\right] q + \left[2 - \frac{3(1-\varepsilon)p\alpha}{(1-\varepsilon)p\alpha + \varepsilon(1-p)}\right], & \text{if } q > \frac{3}{5}. \end{cases}$$

This yields the following best response as a function of α:

$$BR_{2,\text{top}}(\alpha) = \begin{cases} q = \frac{3}{5}, & \text{if } \alpha < \frac{2\varepsilon(1-p)}{3(1-\varepsilon)p}, \\ q \in \left[\frac{3}{5}, 1\right], & \text{if } \alpha = \frac{2\varepsilon(1-p)}{3(1-\varepsilon)p}, \\ q = 1, & \text{if } \alpha > \frac{2\varepsilon(1-p)}{3(1-\varepsilon)p}. \end{cases}$$

This best response implies in particular that $q \geq \frac{3}{5}$ in any ε-PME. Furthermore, since the trembles of player 2 at his top and bottom information sets are independent, a strong type of player 1 will choose whichever of his actions R or U yields the largest worst-case expected payoff given 2's trembles, which yields the following best response in any ε-PME:

$$BR_{1,\text{strong}}(q, r) = \begin{cases} \alpha = 0, & \text{if } q < r + \frac{1}{2(1-\varepsilon)}, \\ \alpha \in [0,1], & \text{if } q = r + \frac{1}{2(1-\varepsilon)}, \\ \alpha = 1, & \text{if } q > r + \frac{1}{2(1-\varepsilon)}. \end{cases}$$

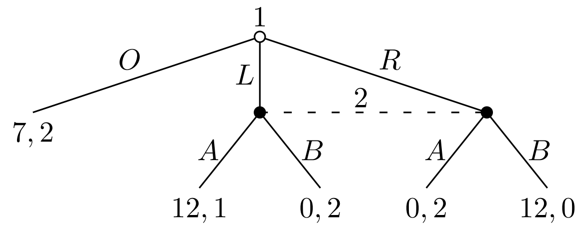

Figure 11: The game from Example 3.

To show that there exists no pooling PME in which player 2 plays F at his bottom information set, or equivalently, sets $r = 0$, note that if such a PME were to exist, it would be the limit of a sequence of ε-PME, indexed by n, with corresponding probabilities $r^n \to 0$, and contaminations defined by $\varepsilon^n \to 0$. For n large enough, the associated values of q^n, r^n and ε^n must satisfy

$$q^n \geq \frac{3}{5} > r^n + \frac{1}{2(1-\varepsilon^n)},$$

in which case the best response of a strong type of player 1 requires $\alpha = 1$, or equivalently, playing R, which contradicts the existence of a pooling PME with $r = 0$.

C Generalized contaminations in Example 3

Consider again the game from Example 3, reproduced in Figure 11, but assume instead a sequence of generalized contaminations $\beta^{[n]}$ defined for any strategy $\beta = (\beta_1, \beta_2)$, such that each $\beta_i^{[n]}$ is given by the closure of a convex open set[26] containing β_i, with $\beta_i^{[n+1]} \subseteq \beta_i^{[n]}$, and the diameter of $\beta_i^{[n]}$ converging to 0 as $n \to \infty$. Assume also that every $\beta_i^{[n]}$ is continuous in β_i. An equilibrium analogous to ε-PME is defined for each n by a strategies $\beta_1 = (1-l-r, l, r)$ and $\beta_2 = (p, 1-p)$ that are optimal given beliefs derived from $\beta^{[n]}$. To derive the beliefs of player 2 at his information set, we project $\beta_1^{[n]}$ onto the coordinates corresponding to actions L and R, and find all conditional beliefs corresponding to points in this projection. If $l = r = 0$, the assumption that $\beta_1^{[n]}$ is the closure of an open set containing β_1 implies that the resulting conditional beliefs are given by all $\mu \in [0, 1]$, where μ denotes the probability assigned to history L. If

[26] Assuming that openness is defined relative to the respective strategy set.

$1-l-r<1$, the resulting conditional beliefs are given by all μ in a (non-empty) interval $[\sigma_n, \tau_n] \subseteq [0,1]$. Player 2's expected utility can then be calculated as

$$u_2(\beta) = \min_{\mu \in [\sigma_n, \tau_n]} \{\mu(2-3p) + 2p\}$$
$$= \begin{cases} (2-3\sigma_n)p + 2\sigma_n, & \text{if } p \leq \frac{2}{3}, \\ (2-3\tau_n)p + 2\tau_n, & \text{if } p > \frac{2}{3}, \end{cases}$$

which yields the following best response:

$$BR_2^n(l,r) = \begin{cases} p = 0, & \text{if } \sigma_n > \frac{2}{3}, \\ p \in \left[0, \frac{2}{3}\right], & \text{if } \sigma_n = \frac{2}{3}, \\ p = \frac{2}{3}, & \text{if } \sigma_n < \frac{2}{3} \text{ and } \tau_n > \frac{2}{3}, \\ p \in \left[\frac{2}{3}, 1\right], & \text{if } \tau_n = \frac{2}{3}, \\ p = 1, & \text{if } \tau_n < \frac{2}{3}. \end{cases}$$

Similarly, from the point of view of player 1, the probability p assigned to action A by player 2 must lie in some interval $[a_n, b_n]$ defined by $\beta_2^{[n]}$, and hence his expected utility is given by

$$u_1(\beta) = 7(1-l-r) + \min_{p \in [a_n, b_n]} \{12p(l-r) + 12r\}$$
$$= \begin{cases} 7 + (12a_n - 7)l + (12(1-a_n) - 7)r, & \text{if } l \geq r, \\ 7 + (12b_n - 7)l + (12(1-b_n) - 7)r, & \text{if } l < r. \end{cases}$$

As in the case of ε-contaminations, we only need to compare the corresponding expected utilities at points where (l,r) is given by $(0,0)$, $\left(\frac{1}{2}, \frac{1}{2}\right)$, $(1,0)$ and $(0,1)$, which yields the following best response:

$$BR_1^n(p) = \begin{cases} \{R\}, & \text{if } a_n < b_n < \frac{5}{12}, \\ \Delta\{R, O\}, & \text{if } a_n < b_n = \frac{5}{12}, \\ \{O\}, & \text{if } a_n < \frac{7}{12} \text{ and } b_n > \frac{5}{12}, \\ \Delta\{O, L\}, & \text{if } b_n > a_n = \frac{7}{12}, \\ \{L\}, & \text{if } b_n > a_n > \frac{7}{12}. \end{cases}$$

For each n, let (l_n, r_n, p_n) denote an equilibrium point associated to the above best responses, and assume that $(l_n, r_n, p_n) \to (l, r, p)$ as $n \to \infty$. Then (l, r, p) corresponds to a PME for such generalized contaminations, and it must be the case that $[a_n, b_n] \to \{p\}$, that p is a best response given beliefs defined by $[\sigma, \tau]$, with $[\sigma_n, \tau_n] \to [\sigma, \tau]$, and that (l, r) is a best response to p. Player 1's best response BR_1^n implies that both l and r cannot be strictly positive at the same time. Hence, either $l > 0$ and $r = 0$, $r > 0$ and $l = 0$, or $l = r = 0$. If $l > 0$ and $r = 0$, we must have $\sigma = \tau = 1$, which implies $p = 0$, contradicting the fact that $l > 0$. Similarly, if $r > 0$ and $l = 0$, we must have $\sigma = \tau = 0$, which implies $p = 1$, contradicting the fact that $r > 0$. Thus, $l = r = 0$ must hold in any such equilibrium, which implies that for beliefs given by $[\sigma, \tau]$, the best response of player 1 is to either play O, mix between O and L, or mix between O and R. For O and R to be optimal at the limit, we need $p = \frac{5}{12}$, which requires $\sigma = \frac{2}{3}$. However, such limit beliefs $[\sigma, \tau]$ can only be attained if $l_n > 0$ and $r_n = 0$ for large enough n, which is only possible if $p_n \geq \frac{7}{12}$, which cannot converge to $p = \frac{5}{12}$. For O only to be optimal at the limit, we need $p \in \left(\frac{5}{12}, \frac{7}{12}\right)$, and since $\frac{7}{12} < \frac{8}{12} = \frac{2}{3}$, this again requires $\sigma = \frac{2}{3}$. However, $p \in \left(\frac{5}{12}, \frac{7}{12}\right)$ can only hold if $a_n < \frac{7}{12}$ and $b_n > \frac{5}{12}$ for large enough n, which implies $l_n = r_n = 0$, and thus, $[\sigma, \tau] = [0, 1]$, which contradicts the optimality of O. The only remaining option is that both O and L are optimal at the limit, which requires $p = \frac{7}{12}$. Since we already showed that $l = r = 0$, this yields the same unique PME that we found for the case based on ε-contaminations.

References

Aryal, G., Stauber, R., 2013. A note on Kuhn's theorem with ambiguity averse players, Working paper, Australian National University, http://cbe.anu.edu.au/media/15784/kuhn.pdf.

Aubin, J.-P., Frankowska, H., 1990. Set-Valued Analysis. Birkhäuser.

Azrieli, Y., Teper, R., 2011. Uncertainty aversion and equilibrium existence in games with incomplete information. Games and Economic Behavior 73 (2), 310–317.

Bade, S., 2011. Ambiguous act equilibria. Games and Economic Behavior 71 (2), 246–260.

Berger, J. O., 1985. Statistical Decision Theory and Bayesian Analysis, 2nd Edition. Springer Series in Statistics. Springer.

Bertsekas, D. P., Nedić, A., Ozdaglar, A. E., 2003. Convex Analysis and Optimization. Athena Scientific.

Bose, S., Daripa, A., 2009. A dynamic mechanism and surplus extraction under ambiguity. Journal of Economic Theory 144 (5), 2084–2114.

Bose, S., Ozdenoren, E., Pape, A., 2006. Optimal auctions with ambiguity. Theoretical Economics 1 (4), 411–438.

Epstein, L. G., Schneider, M., 2003. Recursive multiple-priors. Journal of Economic Theory 113 (1), 1–31.

Fan, K., 1952. Fixed-point and minimax theorems in locally convex topological linear spaces. Proceedings of the National Academy of Sciences 38 (2), 121–126.

Gajdos, T., Hayashi, T., Tallon, J.-M., Vergnaud, J.-C., 2008. Attitude toward imprecise information. Journal of Economic Theory 140 (1), 27–65.

Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. Journal of Mathematical Economics 18 (2), 141–153.

Gilboa, I., Schmeidler, D., 1993. Updating ambiguous beliefs. Journal of Economic Theory 59 (1), 33–49.

Hanany, E., Klibanoff, P., 2007. Updating preferences with multiple priors. Theoretical Economics 2 (3), 261–298.

Hendon, E., Jacobsen, H. J., Sloth, B., 1996. The one-shot-deviation principle for sequential rationality. Games and Economic Behavior 12 (2), 274–282.

Kajii, A., Ui, T., 2005. Incomplete information games with multiple priors. Japanese Economic Review 56 (3), 332–351.

Kopylov, I., 2008. Subjective probability and confidence, Working paper, UC Irvine.

Lo, K. C., 1996. Equilibrium in beliefs under uncertainty. Journal of Economic Theory 71 (2), 443–484.

Lo, K. C., 1999. Extensive form games with uncertainty averse players. Games and Economic Behavior 28 (2), 256–270.

Nadler, S. B., 1978. Hyperspaces of Sets. Marcel Dekker.

Nishimura, K. G., Ozaki, H., 2004. Search and Knightian uncertainty. Journal of Economic Theory 119 (2), 299–333.

Nishimura, K. G., Ozaki, H., 2006. An axiomatic approach to ε-contamination. Economic Theory 27 (2), 333–340.

Osborne, M. J., 2004. An Introduction to Game Theory. Oxford University Press.

Osborne, M. J., Rubinstein, A., 1994. A Course in Game Theory. MIT Press.

Riedel, F., Sass, L., 2013. Ellsberg games. Theory and Decision (forthcoming) (), .

Sass, L., 2013. Kuhn's theorem for extensive form Ellsberg games, Working Paper 478, Institute of Mathematical Economics, Bielefeld University.

Selten, R., 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4 (1), 25–55.

Siniscalchi, M., 2011. Dynamic choice under ambiguity. Theoretical Economics 6 (3), 379–421.

Stauber, R., 2011. Knightian games and robustness to ambiguity. Journal of Economic Theory 146 (1), 248–274.

Strotz, R. H., 1955-1956. Myopia and inconsistency in dynamic utility maximization. Review of Economics Studies 23 (3), 165–180.

Wasserman, L. A., Kadane, J. B., 1990. Bayes' theorem for Choquet capacities. The Annals of Statistics 18 (3), 1328–1339.