

Mouraviev, Igor; Riedel, Frank; Sass, Linda

Working Paper

Kuhn's Theorem for Extensive Form Ellsberg Games

Center for Mathematical Economics Working Papers, No. 510

Provided in Cooperation with:

Center for Mathematical Economics (IMW), Bielefeld University

Suggested Citation: Mouraviev, Igor; Riedel, Frank; Sass, Linda (2014) : Kuhn's Theorem for Extensive Form Ellsberg Games, Center for Mathematical Economics Working Papers, No. 510, Bielefeld University, Center for Mathematical Economics (IMW), Bielefeld, <https://doi.org/10.2139/ssrn.2458913>

This Version is available at:

<https://hdl.handle.net/10419/103928>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Working Papers


Center for
Mathematical
Economics

510

June 2014

Kuhn's Theorem for Extensive Form Ellsberg Games

Igor Mouraviev, Frank Riedel and Linda Sass



IMW · Bielefeld University
Postfach 100131
33501 Bielefeld · Germany



email: imw@uni-bielefeld.de
<http://www.imw.uni-bielefeld.de/wp/2014.php>
ISSN: 0931-6558

Kuhn's Theorem for Extensive Form Ellsberg Games*

Igor Mouraviev, Frank Riedel and Linda Sass

Bielefeld University
Center for Mathematical Economics
33615 Bielefeld, Germany

June 18, 2014

Abstract

The paper generalizes Kuhn's Theorem to extensive form games in which players condition their play on the realization of ambiguous randomization devices and use a maxmin decision rule to evaluate the consequences of their decisions. It proves that ambiguous behavioral and ambiguous mixed strategies are payoff- and outcome equivalent only if the latter strategies satisfy a rectangularity condition. The paper also discusses dynamic consistency. In particular, it shows that not only the profile of ambiguous strategies must be appropriately chosen but also the extensive form must satisfy further restrictions beyond those implied by perfect recall in order to ensure that each player respects her ex ante contingent choice with the evolution of play.

Keywords: Kuhn's Theorem, Strategic Ambiguity, Maxmin Utility, Ellsberg Games

JEL Classification: C72, D81

*This paper is a revised version of IMW Working Paper 478 which was originally published by Linda Sass during her Ph.D. studies at the Center for Mathematical Economics. The authors gratefully acknowledge financial support through Grant Ri 1128-6-1 ("Ambiguity in Games: The Role of uncertainty in strategic interactions") by the German Research Foundation DFG and by the French Research Foundation ANR via their joint programme in social sciences and humanities.

1 Introduction

In this paper we develop a general theoretical framework to analyze dynamic games in which players employ ambiguity as strategy to influence the outcome of a game¹. We build on previous work by Riedel and Sass (2013) who study normal form games with *objective ambiguity*. As the authors conceptualize it, each player can create objective ambiguity by setting up an Ellsberg urn (i.e., an urn where the exact composition of balls of different colors is not known) and conditioning her strategy choice on a draw from this urn. In this paper, we extend this approach to finite extensive games in which players are allowed to choose compact and convex sets of mixed strategies. Such strategies are called Ellsberg strategies and games where players employ such strategies are called Ellsberg games (hence, the title of the paper). In our setting, players possess objective but imprecise information which we model by using sets of probability distributions over possible outcomes of a game. As shown by Gajdos et al. (2008), in such an environment each player’s preferences can be represented by the minimal expected utility evaluated with respect to all possible probability distributions over the outcomes of a game.

Our first contribution is a version of Kuhn’s Theorem for extensive-form games.² More precisely, we show that for each Ellsberg strategy satisfying a “rectangularity” condition there exists an outcome- and payoff-equivalent Ellsberg behavioral strategy, and vice versa. The rectangularity condition that we impose can be regarded as a generalization of the familiar condition used in the standard analysis which establishes the relationship between behavioral and mixed strategies. Specifically, with an Ellsberg strategy satisfying this condition, a player

¹In his paper on peace negotiations between two countries seeking to end the continuing military confrontation, Greenberg (2000) argues that the best strategy for the mediator is to “remain silent” regarding its policy to impose sanctions in case the negotiations break down. More recently, in his study of the Two China Problem D’Amato (2010) finds that the US deliberately relied on strategic ambiguity to deter China from using force against Taiwan. In particular, by signing the Taiwan Relations Act in 1979, which was neither a treaty nor a nontreaty, it actually left the two sides guessing at its willingness to intervene in a conflict. In a similar vein, the literature on communication provides numerous examples where agents pursuing conflicting goals opt to be vague or imprecise in their statements. As Eisenberg (1984) puts it: “Many different strategies are used to orient toward conflicting interactional goals; some examples include avoiding interaction altogether, remaining silent, or changing the topic.”

²Kuhn’s Theorem states that mixed and behavioral strategies are equivalent (in the sense that they induce the same probability distribution over outcomes of a game) if and only if the game satisfies perfect recall.

evaluates her *ex ante* expected utility by minimizing separately the conditional expectations on each of her information sets and then minimizing the expectation of these conditional expectations. Since, with an Ellsberg behavioral strategy, she evaluates her expected utility in exactly the same way, both strategies yield identical *ex ante* payoffs.

In contrast to Kuhn's Theorem, our proof does not rely on the notion of perfect recall. The reason is as follows. In Kuhn's setting, this property ensures that for *every* mixed strategy there exists a behavioral strategy which induces identical probability distribution over terminal nodes. In our setting, it is the rectangularity condition that plays a similar role. More precisely, if an Ellsberg strategy is rectangular then it comprises only those mixed strategies that can be replicated by behavioral strategies.

Our second main result is that dynamic consistency associated with the requirement that all players respect their *ex ante* contingent choices with the evolution of play generally does not hold in Ellsberg games even if they satisfy perfect recall. A similar point is made by Aryal and Stauber (2013) in their note on Kuhn's Theorem with ambiguity averse players. In particular, they provide an example of a *three*-player extensive game in which it is impossible to achieve dynamic consistency if one assumes consistency and independency of ambiguous beliefs/strategies across the players. Although their example is instructive, it does not fully capture the role of perfect recall in ensuring dynamic consistency. In this paper, we focus on this issue and show by means of examples that dynamic consistency may be violated even in (*two*-player) extensive games satisfying perfect recall.

Our third main result is that, in order to achieve dynamic consistency, the extensive game must satisfy further restrictions beyond those implied by perfect recall. One restriction implies that every time when a player is called upon to move she needs to make a conjecture about the actions of only one of her opponents who moved at the immediately preceding stage.³ Another restriction implies that in doing so she can disregard everything what happened before that stage. We show that if the game satisfies those restrictions then dynamic consistency can be restored by suitably choosing the profile of ambiguous strategies.

The paper is organized as follows. In the next section, we introduce notation

³Note that this condition is stronger than the one of perfect recall because it requires a player to remember not only what she did but also what her opponents did before the previous period.

and definitions. In Section 3, we prove a version of Kuhn's Theorem for extensive form Ellsberg games. In Section 4, we give examples of Ellsberg games satisfying perfect recall in which dynamic consistency cannot be achieved. We use these examples to justify the restrictions that we impose on the ambiguous strategies and the structure of a game in order to ensure dynamic consistency. In Section 5, we offer concluding remarks.

2 Notation and Definitions

We find it convenient to employ the tree-based definition of extensive game which was originally introduced by Kuhn (1953) and, recently, has been adopted in a number of related studies.⁴ Formally, an extensive game comprises the following constituents.

- A finite rooted tree $\langle \mathcal{T}, \rightarrow, t_0 \rangle$ which specifies the order of moves. Here, \mathcal{T} denotes a finite collection of nodes, $t_0 \in \mathcal{T}$ is the root, and for any $t', t'' \in \mathcal{T}$, the relation $t' \rightarrow t''$ means that t' is *immediate predecessor* of t'' or t'' is *immediate successor* of t' . A *path* from node t' to node t'' is a sequence (t_1, \dots, t_n) of nodes starting at $t_1 = t'$ and ending at $t_n = t''$ such that t_k is immediate successor of t_{k-1} for all $k = 2, \dots, n$. We denote by \prec the *precedence relation* on \mathcal{T} with $t' \prec t''$ if there is a path from t' to t'' . We call then t' a *predecessor* of t'' . As usual, $t' \preceq t''$ means $t' = t''$ or $t' \prec t''$. Recall that the precedence relation is a *well-ordered partial order*; each node $t \in \mathcal{T}$ (except t_0) has exactly one immediate predecessor and the set of predecessors is linearly ordered. The nodes that are not predecessors of any other nodes are called *terminal nodes*, and the set of these nodes is denoted by \mathcal{Z} . Each terminal node z induces a unique path \bar{z} from the root to z ; such paths are called *plays*. We also denote by $\mathcal{X} \equiv \mathcal{T} \setminus \mathcal{Z}$ the set of *decision nodes*.
- A set of *players* $\mathcal{N} = \{1, \dots, N\}$ and a *partition* of \mathcal{X} into subsets $\mathcal{X}_1, \dots, \mathcal{X}_N$ where each \mathcal{X}_i denotes the set of player i 's decision nodes.
- An *equivalence relation* \sim defined on the product spaces $\mathcal{X}_i \times \mathcal{X}_i$ such that for $x' \sim x'' \in \mathcal{X}_i$ then the number of immediate successors of x' is equal

⁴See, for example, Brandenburger (2007) and Bonnano (2004).

to the number of immediate successors of x'' . Thus, $x' \sim x''$ means that player i does not have the information to distinguish between x' and x'' . The equivalence relation allows us to further partition each \mathcal{X}_i into subsets called *information sets*. We denote by \mathcal{H}_i the *partition* of \mathcal{X}_i and by $H_i \in \mathcal{H}_i$ a typical information set of player i .

- A *choice partition* which partitions the branches out of nodes in a given information set of a player into that player's choices in that information set. We write $x' \rightarrow_c x''$ if c is the action or choice that leads from x' to its immediate successor x'' . The choice partition satisfies the following two conditions: (i) if $x' \rightarrow_c t$ and $x'' \rightarrow_c t$ then $x' = x''$, and (ii) if $x' \rightarrow_c t'$ and $x'' \sim_i x'$ then there exists t'' such that $x'' \rightarrow_c t''$. We denote by $\mathcal{C}(H_i)$ the set of choices available at the information set H_i .
- A player i 's *payoff* associated with each terminal node which is represented by a function $u_i : \mathcal{Z} \rightarrow \mathbb{R}$ (where \mathbb{R} is the set of real numbers).

As is common in the literature, we shall consider only those extensive games in which every decision node has at least two outgoing branches, and in which no path from the root to a terminal node crosses any information set more than once. In the latter case, we require the game to satisfy the so called 'no absentmindedness' condition. Formally, it states:

(NAM) If $x' \sim x''$ and $x' \lesssim x''$ then $x' = x''$ for all $x', x'' \in \mathcal{X}$.

Within the class of extensive games that we consider, it will prove useful to distinguish the games which satisfy *perfect recall*. This property, which was first introduced by Kuhn (1953), requires that no player ever forgets what she knew and what she did in the past. Formally, the game is of perfect recall if

(PR) For every player $k \in \mathcal{N}$, for all nodes $x_1, x'_1, x'_2 \in \mathcal{X}_k$ and $t_1 \in \mathcal{T}$ and for every choice c , if $x_1 \rightarrow_c t_1$, $t_1 \lesssim x'_1$ and $x'_1 \sim x'_2$ then there exists nodes $x_2 \in \mathcal{X}_k$ and $t_2 \in \mathcal{T}$ such that $x_2 \sim x_1$, $x_2 \rightarrow_c t_2$ and $t_2 \lesssim x'_2$.

We now turn to the definitions of strategies.

Definition 1 A pure strategy of player i is a function s_i that maps player i 's information sets into choices available at those sets such that $s_i(H_i) \in \mathcal{C}(H_i)$ for each $H_i \in \mathcal{H}_i$.

We write $s_i(H_i) = c$ if s_i specifies that player i makes a choice c at H_i . The set of pure strategies of player i is given by $S_i \equiv \times_{H_i \in \mathcal{H}_i} \mathcal{C}(H_i)$.

Definition 2 A mixed strategy of player i is a probability distribution σ_i on S_i so that $\sigma_i(s_i)$ is the probability of playing $s_i \in S_i$.

Definition 3 A behavioral strategy of player i is a function β_i that assigns to each $H_i \in \mathcal{H}_i$ a probability distribution β_{H_i} over the set of choices available at H_i .

Let $\Delta(A)$ denote the set of all probability distributions on some set A . The set of player i 's mixed strategies is then given by $\Delta(S_i)$ while the set of player i 's behavioral strategies is given by $\times_{H_i \in \mathcal{H}_i} \Delta(\mathcal{C}(H_i))$.

Denote by $P_z(\sigma_i, \sigma_{-i})$ the probability of reaching a terminal node $z \in \mathcal{Z}$ induced by the strategy profile (σ_i, σ_{-i}) .⁵

Definition 4 The strategies σ'_i and σ''_i are equivalent if $P_z(\sigma'_i, \sigma_{-i}) = P_z(\sigma''_i, \sigma_{-i})$ for all $z \in \mathcal{Z}$ and all $\sigma_{-i} \in \times_{j \neq i} \Delta(S_j)$.

The definition thus states that two (mixed) strategies of player i are equivalent if they induce the same probability distribution over terminal nodes regardless of how the opponents play.

Denote by \bar{z} the path from the root t_0 to a terminal node $z \in \mathcal{Z}$. Note that for the class of games that we consider every strategy profile $s \equiv (s_1, \dots, s_N) \in \times_{i \in \mathcal{N}} S_i$ induces a unique \bar{z} . This allows us to state the following definitions.

Definition 5 Given a play \bar{z} , the strategy $s_i \in S_i$ is compatible with \bar{z} if there exists $s_{-i} \in S_{-i}$ such that the profile (s_i, s_{-i}) induces \bar{z} .

Definition 6 Given information set $H_i \in \mathcal{H}_i$, the strategy s_i is compatible with H_i if there exists $s_{-i} \in S_{-i}$ such that the path induced by the profile (s_i, s_{-i}) passes through node $x \in H_i$.

⁵We shall refer to all players other than player i as “player i 's opponents” and denote them by “ $-i$ ”.

For each player i denote by S_{H_i} the set of pure strategies that are compatible with H_i and by $S_{H_i}(c)$ the set of pure strategies that are compatible with H_i and $c \in \mathcal{C}(H_i)$, i.e.,

$$\begin{aligned} S_{H_i} &\equiv \{s_i \in S_i : s_i \text{ is compatible with } H_i\}, \\ S_{H_i}(c) &\equiv \{s_i \in S_i : s_i \in S_{H_i} \text{ and } s_i(H_i) = c\}. \end{aligned}$$

For given probability distribution $\sigma_i \in \Delta(S_i)$, define the probability $\bar{\sigma}_{H_i}$ which is the restriction of σ_i on S_{H_i} , i.e.,

$$\bar{\sigma}_{H_i} = \sum_{s_i \in S_{H_i}} \sigma_i(s_i), \quad (1)$$

and the probability $\sigma_{H_i}(c)$ which is the restriction of σ_i on $S_{H_i}(c)$, i.e.,

$$\sigma_{H_i}(c) = \sum_{s_i \in S_{H_i}(c)} \sigma_i(s_i). \quad (2)$$

Also, for $\bar{\sigma}_{H_i} > 0$, define the probability $\theta_{H_i}(c)$ of making a choice $c \in \mathcal{C}(H_i)$ conditional on H_i as

$$\sigma_{H_i}(c) = \theta_{H_i}(c) \bar{\sigma}_{H_i}. \quad (3)$$

If $\bar{\sigma}_{H_i} = 0$ then the information set H_i is a “null event” in the sense that it will never be reached regardless of the strategies chosen by player i ’s opponents. In this case, the conditional probability $\theta_{H_i}(c)$ cannot be derived from the initial probability distribution σ_i . Nevertheless, this feature does not affect our results since what matters for the analysis is the strategies chosen by player i at the information sets that can occur with strictly positive probabilities. We thus take θ_{H_i} to be the uniform distribution over available choices in this case.

The following example illustrates these concepts.

Example 1. Consider the game-tree in Figure 1.⁶ Suppose that, by the rules of the game, player 1 is allowed to move at x_1 and x_3 while player 2 is allowed to move at x_2 , x_4 and x_5 . In which case, player 1’s information sets are given

⁶Hereafter, each letter placed along the branch out of a node denotes the choice made at this node. Note that the choices available at x_4 and x_5 are the same which reflects the fact that these nodes are equivalent.




Figure 1: Game-tree in Example 1.

by $H_{11} = \{x_1\}$ and $H_{12} = \{x_3\}$ while player 2's information sets are given by $H_{21} = \{x_2\}$ and $H_{22} = \{x_4, x_5\}$. This implies that $S_1 = \{s_{AE}, s_{AF}, s_{BE}, s_{BF}\}$ and $S_2 = \{s_{ae}, s_{af}, s_{be}, s_{bf}\}$.

Consider H_{11} . Since every $s_1 \in S_1$ is compatible with H_{11} then $S_{H_{11}} = S$ and, therefore, $\bar{\sigma}_{H_{11}} = 1$. Next, the only strategies which are compatible with H_{11} and $c = A$ are s_{AE} and s_{AF} . Hence, $S_{H_{11}}(A) = \{s_{AE}, s_{AF}\}$. Likewise, it can be shown that $S_{H_{11}}(B) = \{s_{BE}, s_{BF}\}$. Therefore,

$$\sigma_{H_{11}}(c) = \sigma_1(s_{cE}) + \sigma_1(s_{cF}) = \theta_{H_{11}}(c) \text{ for each } c = \{A, B\}. \quad (4)$$

Consider now H_{12} . Since this set can be reached only if player 1 chooses A at H_{11} , then $S_{H_{12}} = \{s_{AE}, s_{AF}\}$ implying that $\bar{\sigma}_{H_{12}} = \sigma_1(s_{AE}) + \sigma_1(s_{AF})$. By applying a similar reasoning as before, it can be shown that $S_{H_{12}}(E) = \{s_{AE}\}$ and $S_{H_{12}}(F) = \{s_{AF}\}$. As a result, we have $\sigma_{H_{12}}(c) = \sigma_1(s_{Ac})$ and

$$\sigma_1(s_{Ac}) = \theta_{H_{12}}(c)(\sigma_1(s_{Ac}) + \sigma_1(s_{Ac})) \text{ for each } c = \{E, F\}. \quad (5)$$

Let us now turn to player 2 and consider H_{21} . As before, it can be shown that $S_{H_{21}} = S_2$, $S_{H_{21}}(a) = \{s_{ae}, s_{af}\}$ and $S_{H_{21}}(b) = \{s_{be}, s_{bf}\}$ which yields

$$\sigma_{H_{21}}(c) = \sigma_2(s_{ce}) + \sigma_2(s_{cf}) = \theta_{H_{21}}(c) \text{ for each } c = \{a, b\}. \quad (6)$$

Finally, consider H_{22} . In which case, every $s_2 \in S_2$ is compatible with H_{22} . Indeed, s_{be} and s_{bf} are compatible with H_{22} because the paths induced by (s_{AF}, s_{be}) and (s_{AF}, s_{bf}) cross H_{22} at x_4 . Similarly, s_{ae} and s_{af} are compatible with H_{22} because the paths induced by (s_{BF}, s_{ae}) and (s_{BF}, s_{af}) cross H_{22} at x_5 . Since $S_{H_{22}}(e) = \{s_{ae}, s_{be}\}$ and $S_{H_{22}}(f) = \{s_{af}, s_{bf}\}$, we have

$$\sigma_{H_{22}}(c) = \sigma_2(s_{ac}) + \sigma_2(s_{ac}) = \theta_{H_{22}}(c) \text{ for each } c = \{e, f\}, \quad (7)$$

which completes the discussion. ■

Following the approach of Riedel and Sass (2013), we allow each player to condition her strategy choice on the realization of an ambiguous randomization device such as an Ellsberg urn. Formally, we state:⁷

Definition 7 *An Ellsberg strategy of player i is a compact and convex set $\Sigma_i \subseteq \Delta(S_i)$.*

The definition thus implies that each element $\sigma_i \in \Sigma_i$ is a mixed strategy of player i . By analogy with the rectangularity condition imposed by Epstein and Schneider (2003) on decision maker's ambiguous beliefs, we impose a rectangularity condition on the Ellsberg strategy Σ_i . Specifically, we require that for every θ'_{H_i} corresponding to some $\sigma'_i \in \Sigma_i$ and every $\bar{\sigma}''_{H_i}$ corresponding to some $\sigma''_i \in \Sigma_i$, the combination $\theta'_{H_i} \bar{\sigma}''_{H_i}$ corresponds to some $\sigma'''_i \in \Sigma_i$, provided that $\bar{\sigma}_{H_i} > 0$ for

⁷We will discuss the appropriateness of this definition in Footnote 10.

every $\sigma_i \in \Sigma_i$.⁸ More formally, define the following sets:

$$\begin{aligned}\bar{\Sigma}_{H_i} &= \{\bar{\sigma}_{H_i} : \sigma_i \in \Sigma_i\}, \\ \Sigma_{H_i} &= \{\sigma_{H_i} : \sigma_i \in \Sigma_i\}, \\ \Theta_{H_i} &= \{\theta_{H_i} : \sigma_i \in \Sigma_i\}.\end{aligned}$$

Denoting by \mathcal{J}_{H_i} the collection of sets $S_{H_i}(c)$, i.e., $\mathcal{J}_{H_i} \equiv \{S_{H_i}(c)\}_{c \in \mathcal{C}(H_i)}$, we state:

Definition 8 *An Ellsberg strategy Σ_i is $\{\mathcal{J}_{H_i}\}$ -rectangular if for every $H_i \in \mathcal{H}_i$*

$$\Sigma_{H_i} = \{\theta_{H_i} \bar{\sigma}_{H_i} : \theta_{H_i} \in \Theta_{H_i}, \bar{\sigma}_{H_i} \in \bar{\Sigma}_{H_i}\}.$$

A rectangular set Σ_i can be constructed by, first, specifying the set Θ_{H_i} for every $H_i \in \mathcal{H}_i$ and, then, using recursively (1)-(3) to determine the relationship between the corresponding probability measures so as to satisfy the rectangularity condition.

An alternative way to define a rectangular Ellsberg strategy could rely on the event-based approach proposed by Epstein and Schneider (2003) to modeling ambiguous beliefs. Specifically, one could view an extensive game as an event-tree in which each event is represented by a subset of the set of pure strategies compatible with a choice at some information set. The interpretation is that a player does not observe the strategy chosen as the realization of a draw from an Ellsberg urn; what she observes is the choices made at the information sets which allow her to refine the knowledge of the strategy actually played. In this setting, the information structure would be represented by the filtration $\{\mathcal{J}_{H_i}\}_{H_i \in \mathcal{H}_i}$ where $\mathcal{J}_{H_i} \equiv \{S_{H_i}(c)\}_{c \in \mathcal{C}(H_i)}$ for each $H_i \in \mathcal{H}_i$, and a rectangular set Σ_i could be constructed by specifying the sets of one-step-ahead distributions for each $H_i \in \mathcal{H}_i$. Although such a specification is equivalent to our specification of Θ_{H_i} 's and both approaches yield identical rectangular sets Σ_i 's (if player i has perfect recall), we do not pursue such a definition because it relies on the notion of "time line" along

⁸If $\bar{\sigma}'_{H_i} = 0$ for some $\sigma'_i \in \Sigma_i$, then θ'_{H_i} can be specified arbitrarily. In particular, it can be set equal to $\hat{\theta}_{H_i}$ which corresponds to some $\hat{\sigma}_i \in \Sigma_i$ that induces $\hat{\sigma}_{H_i} > 0$. In this case, the requirement that the combination $\theta'_{H_i} \bar{\sigma}'_{H_i}$ corresponds to some $\sigma_i''' \in \Sigma_i$ is equivalent to the requirement that the combination $\hat{\theta}_{H_i} \bar{\sigma}'_{H_i}$ corresponds to some $\sigma_i''' \in \Sigma_i$.

which the information sets can be ordered. While other approaches to define the rectangularity condition for Ellsberg strategies are possible, we view our approach as the most parsimonious departure from the standard analysis under expected utility maximization, which relies on condition (3) to establish the relationship between mixed and behavioral strategies.

In particular, by analogy with the standard analysis, we allow each player i to employ stochastically independent Ellsberg urns governing her choices at each of her information sets instead of employing a unique Ellsberg urn governing her choice of a mixed strategy (associated with the strategic form of the extensive game). Formally, we state:

Definition 9 *An Ellsberg behavioral strategy of player i is a function that assigns to each information set $H_i \in \mathcal{H}_i$ a compact and convex set $B_{H_i} \subseteq \Delta(\mathcal{C}(H_i))$.*

The definition thus implies that player i 's Ellsberg behavioral strategy is the collection of sets $B = \times_{H_i \in \mathcal{H}_i} B_{H_i}$ with typical element $\beta_i = (\beta_{H_i})_{H_i \in \mathcal{H}_i}$ where $\beta_{H_i} \in B_{H_i}$. The set of player i 's behavioral strategies is $\mathcal{B}_i \equiv \times_{H_i \in \mathcal{H}_i} \Delta(\mathcal{C}(H_i))$.

As we explained in the Introduction, our modeling approach captures the idea of introducing objective ambiguity in a setting of strategic interaction. Players' preferences in such an environment have been characterized by Gajdos et al. (2008) who show that in case of strict ambiguity aversion they can be represented by the utility functionals in the spirit of Gilboa and Schmeidler (1989). In our game-theoretical setting, this means that each player i evaluates her expected utility not only with respect to the worst probability distribution induced by her opponents' ambiguous strategy profile Σ_{-i} but also with respect to the worst probability distribution induced by her own ambiguous strategy Σ_i , i.e., her expected utility is given by⁹

$$U_i(\Sigma_i, \Sigma_{-i}) \equiv \min_{\substack{\sigma_i \in \Sigma_i \\ \sigma_{-i} \in \Sigma_{-i}}} \mathbf{E}_{\sigma_i, \sigma_{-i}}[u_i(z)],$$

where $\mathbf{E}_{\sigma_i, \sigma_{-i}}$ is the expectation operator associated with the probability distributions σ_i and σ_{-i} .¹⁰

⁹Throughout the paper, it will be assumed that each player i has complete knowledge of her opponents' strategies Σ_{-i} .

¹⁰The above condition, in particular, implies that defining Ellsberg strategies as compact and convex sets of probability measures is without loss of generality. Indeed, if $co(cl(\Sigma_i))$ denotes the

The following definition captures the notion of relative amounts of ambiguity associated with playing the strategies Σ'_i and Σ''_i .

Definition 10 *Given the strategy profile Σ_{-i} , the strategy Σ'_i is more ambiguous than the strategy Σ''_i if for any $\sigma''_i \in \Sigma''_i$ there exists $\sigma'_i \in \Sigma'_i$ such that $P_z(\sigma'_i, \sigma_{-i}) = P_z(\sigma''_i, \sigma_{-i})$ for all $z \in \mathcal{Z}$ and all $\sigma_{-i} \in \Sigma_{-i}$.*

Note that the above definition implies that if Σ'_i is more ambiguous than Σ''_i then it must be $U_i(\Sigma'_i, \Sigma_{-i}) \leq U_i(\Sigma''_i, \Sigma_{-i})$. Finally, we state:

Definition 11 *Given the strategy profile Σ_{-i} , the strategies Σ'_i and Σ''_i are payoff equivalent if $U_i(\Sigma'_i, \Sigma_{-i}) = U_i(\Sigma''_i, \Sigma_{-i})$. The strategies Σ'_i and Σ''_i are called outcome equivalent if they are payoff equivalent for every specification of terminal payoffs u_i .*

Note that we do not have a natural notion of outcome distribution in our Knightian setting; we thus define Ellsberg strategies to be outcome equivalent if their minimal payoffs do not depend on the particular terminal payoffs.

3 Version of Kuhn's Theorem

In his seminal paper, Kuhn (1953) establishes the relationship between mixed and behavioral strategies for finite extensive games. More precisely, he proves the following result.

Theorem (Kuhn, 1953). *An extensive game satisfies (PR) if and only if, for any mixed strategy σ_i of every player $i \in \mathcal{N}$ there exists an outcome equivalent behavioral strategy β_i , and vice versa.*

In this section, we generalize Kuhn's Theorem to extensive form Ellsberg games with ambiguity averse players. We proceed by stating two preliminary results.

convex hull of $cl(\Sigma_i)$, then we must have (using that $\mathbf{E}_{\sigma_i, \sigma_{-i}}[u_i(z)]$ is a linear function of σ_i and σ_{-i})

$$\min_{\sigma_i \in cl(\Sigma_i)} \mathbf{E}_{\sigma_i, \sigma_{-i}}[u_i(z)] = \min_{\sigma_i \in co(cl(\Sigma_i))} \mathbf{E}_{\sigma_i, \sigma_{-i}}[u_i(z)],$$

for every $\sigma_{-i} \in \Sigma_{-i}$.

Theorem 1. *Suppose that an extensive game satisfies (NAM) and the strategy profile Σ_{-i} is given. Then, for any behavioral Ellsberg strategy B_i there exists a more ambiguous rectangular Ellsberg strategy Σ_i .*

Proof. Given $B_i = \times_{H_i \in \mathcal{H}_i} B_{H_i}$, set $\Theta_{H_i} = B_{H_i}$ for every $H_i \in \mathcal{H}_i$. Define the rectangular strategy Σ_i such that

$$\Sigma_{H_i} = \{ \beta_{H_i} \bar{\sigma}_{H_i} : \beta_{H_i} \in B_{H_i}, \bar{\sigma}_{H_i} \in \bar{\Sigma}_{H_i} \},$$

for every $H_i \in \mathcal{H}_i$.

Fix a path \bar{z} from the root to a terminal node $z \in \mathcal{Z}$. Suppose that \bar{z} crosses the information sets H_{j1}, \dots, H_{jK_j} that belong to player j where $K_j \geq 1$ for $j = i$ and $K_j \geq 0$ for $j \neq i$. Denote by c_{jk} the choice of player j at H_{jk} such that the branch associated with c_{jk} is contained in \bar{z} .

Next, choose $\beta_i \in B_i$ and $\sigma_{-i} \in \Sigma_{-i}$. For given β_i and σ_{-i} , define the probability $\tilde{\beta}_{H_{jk}}(c_{jk})$ that player j chooses c_{jk} at H_{jk} . If $j = i$ then $\tilde{\beta}_{H_{jk}}(c_{jk}) = \beta_{H_{jk}}(c_{jk})$, while if $j \neq i$ then $\tilde{\beta}_{H_{jk}}(c_{jk})$ is determined by σ_j . Since condition (NAM) implies that no information set can appear more than once along any path from the root to a terminal node, we have

$$P_z(\beta_i, \sigma_{-i}) = \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \beta_{H_{ik}}(c_{ik}),$$

where

$$\bar{P}_z(\sigma_{-i}) \equiv \prod_{j \neq i} \prod_{k=0}^{K_j} \tilde{\beta}_{H_{jk}}(c_{jk}). \quad (8)$$

Consider now the mixed strategy σ_i which assigns to the pure strategy $s_i = (c_{H_i})_{H_i \in \mathcal{H}_i}$ the probability

$$\sigma_i(s_i) = \prod_{H_i \in \mathcal{H}_i} \beta_{H_i}(c_{H_i}). \quad (9)$$

This strategy is well defined, since (using $S_i = \times_{H_i \in \mathcal{H}_i} \mathcal{C}(H_i)$)

$$\sum_{s_i \in S_i} \sigma_i(s_i) = \prod_{H_i \in \mathcal{H}_i} \left(\sum_{c_{H_i} \in \mathcal{C}(H_i)} \beta_{H_i}(c_{H_i}) \right) = 1.$$

Denote by $\widehat{S}_{H_i}(c_{H_i})$ the set of strategies that require player i to choose c_{H_i} at H_i , i.e.,¹¹

$$\widehat{S}_{H_i}(c_{H_i}) \equiv \{s_i : s_i \in S_i \text{ and } s_i(H_i) = c_{H_i}\}.$$

Using (9), we have

$$\begin{aligned} \sum_{s_i \in \widehat{S}_{H_i}(c_{H_i})} \sigma_i(s_i) &= \beta_{H_i}(c_{H_i}) \prod_{H'_i \in \mathcal{H}_i \setminus H_i} \left(\sum_{c_{H'_i} \in \mathcal{C}(H'_i)} \beta_{H'_i}(c_{H'_i}) \right) \\ &= \beta_{H_i}(c_{H_i}). \end{aligned} \quad (10)$$

It can be verified that the set of player i 's pure strategies which are compatible with \bar{z} is $\cap_{k=1}^{K_i} \widehat{S}_{H_{ik}}(c_{ik})$. Thus, by using (10) and applying a similar argument as before, we have

$$\begin{aligned} P_z(\sigma_i, \sigma_{-i}) &= \bar{P}_z(\sigma_{-i}) \sum_{s_i \in \cap_{k=1}^{K_i} \widehat{S}_{H_{ik}}(c_{ik})} \sigma_i(s_i) \\ &= \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \left(\sum_{s_i \in \widehat{S}_{H_{ik}}(c_{ik})} \sigma_i(s_i) \right) \\ &= \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \beta_{H_{ik}}(c_{ik}) \\ &= P_z(\beta_i, \sigma_{-i}). \end{aligned}$$

It remains to show that σ_i is an element of the rectangular Ellsberg strategy Σ_i . Note first that $\sigma_i(s_i)$ can be written as (again, using $s_i = (c_{H_i})_{H_i \in \mathcal{H}_i}$)

$$\begin{aligned} \sigma_i(s_i) &= \beta_{H_i}(c_{H_i}) \prod_{H'_i \in \mathcal{H}_i \setminus H_i} \beta_{H'_i}(c_{H'_i}) \\ &= \beta_{H_i}(c_{H_i}) \sum_{c_{H_i} \in \mathcal{C}(H_i)} \prod_{H'_i \in \mathcal{H}_i} \beta_{H'_i}(c_{H'_i}) \\ &= \beta_{H_i}(c_{H_i}) \sum_{c_{H_i} \in \mathcal{C}(H_i)} \sigma_i(s_i). \end{aligned}$$

¹¹Note that, by construction, $S_{H_i}(c) \subset \widehat{S}_{H_i}(c)$ for each $c \in \mathcal{C}(H_i)$.

Using the above condition, we have

$$\begin{aligned}
\sigma_{H_i}(c) &= \sum_{s_i \in S_{H_i}(c)} \sigma_i(s_i) \\
&= \sum_{s_i \in S_{H_i}(c)} \beta_{H_i}(c) \sum_{c_{H_i} \in \mathcal{C}(H_i)} \sigma_i(s_i) \\
&= \beta_{H_i}(c) \sum_{c_{H_i} \in \mathcal{C}(H_i)} \sum_{s_i \in S_{H_i}(c)} \sigma_i(s_i) \\
&= \beta_{H_i}(c) \bar{\sigma}_{H_i},
\end{aligned}$$

which implies that σ_{H_i} can be represented as the product of β_{H_i} and $\bar{\sigma}_{H_i}$. Since the set Σ_{H_i} comprises all the combinations $\beta_{H_i} \bar{\sigma}_{H_i}$ where $\beta_{H_i} \in B_{H_i}$ and $\bar{\sigma}_{H_i} \in \bar{\Sigma}_{H_i}$, σ_{H_i} must belong to Σ_{H_i} . This completes the proof. ■

At this point, it is worth clarifying the role of the rectangularity condition for the result obtained. Specifically, if the Ellsberg strategy Σ_i induces the set of probability measures B_{H_i} at each $H_i \in \mathcal{H}_i$ but it is not rectangular then there exists $\beta_i \in B_i$ such that neither $\sigma_i \in \Sigma_i$ induces the same probability of reaching each terminal node $z \in \mathcal{Z}$ as β_i does. The following example illustrates the point.

Example 1 (continued). As before, suppose that player 1 is allowed to move at x_1 and x_3 . Since there are only two choices available at x_1 and x_3 , her Ellsberg behavioral strategy is given by $B_1 = [\underline{\beta}_A, \bar{\beta}_A] \times [\underline{\beta}_F, \bar{\beta}_F]$ where $[\underline{\beta}_A, \bar{\beta}_A]$ and $[\underline{\beta}_F, \bar{\beta}_F]$ are the sets of probabilities which correspond to choosing A at H_{11} and F at H_{12} , respectively. Define the rectangular Ellsberg strategy

$$\begin{aligned}
\Sigma_1 &= \{(\sigma_1(s_{AE}), \sigma_1(s_{AF}), \sigma_1(s_{BE}), \sigma_1(s_{BF})) : \sigma_1(s_{BE}) + \sigma_1(s_{BF}) = 1 - \beta_A, \\
\sigma_1(s_{AE}) &= \beta_A(1 - \beta_F), \sigma_1(s_{AF}) = \beta_A\beta_F, \beta_A \in [\underline{\beta}_A, \bar{\beta}_A] \text{ and } \beta_F \in [\underline{\beta}_F, \bar{\beta}_F]\}.
\end{aligned}$$

It is straightforward to verify that for each $\beta_1 \in B_1$ there exists $\sigma_1 \in \Sigma_1$ which induces identical probabilities of reaching each $z \in \mathcal{Z}$. In Figure 2, the projection of the set Σ_1 onto the $(\sigma_1(s_{AE}), \sigma_1(s_{AF}))$ plane is illustrated by the area $QTKM$, and the set B_{H_2} is represented by the thick black line connecting the points $(1 - \bar{\beta}_F, \bar{\beta}_F)$ and $(1 - \underline{\beta}_F, \underline{\beta}_F)$.

Consider now the set $\widehat{\Sigma}_1$ represented by the shaded area $QYKW$ in Figure 2. This set is not rectangular: even though for any $\sigma_1 \in \widehat{\Sigma}_1$ there exists $\beta_1 \in B_1$ such




Figure 2: Projection of the set Σ_1 onto the $(\sigma_{AE}, \sigma_{AF})$ plane.

that $\sigma_1(s_{AE}) = \beta_A(1 - \beta_F)$, $\sigma_1(s_{AF}) = \beta_A\beta_F$ and $\sigma_1(s_{BE}) + \sigma_1(s_{BF}) = 1 - \beta_A$, the converse is not true. For example, for $\beta_1 = (\bar{\beta}_A, \bar{\beta}_F)$ there is no $\sigma_1 \in \Sigma_1$ which induces the same probability distribution over the terminal nodes as β_1 does. ■

We now state our second preliminary result.

Theorem 2. *Suppose that an extensive game satisfies (NAM) and the strategy profile Σ_{-i} is given. Then, for any rectangular Ellsberg strategy Σ_i there exists a more ambiguous behavioral Ellsberg strategy B_i .*

Proof. Fix a path \bar{z} from the root to a terminal node $z \in \mathcal{Z}$. Suppose that \bar{z} crosses the information sets H_{j1}, \dots, H_{jK_j} that belong to player j where $K_j \geq 1$ for $j = i$ and $K_j \geq 0$ for $j \neq i$. Denote by c_{jk} the choice of player j at H_{jk} such that the branch associated with c_{jk} is contained in \bar{z} .

Choose $\sigma_i \in \Sigma_i$ and $\sigma_{-i} \in \Sigma_{-i}$. For given σ_i and σ_{-i} , define the probability $\tilde{\beta}_{H_{jk}}(c_{jk})$ that player j chooses c_{jk} at H_{jk} . Note that, since the Ellsberg strategy Σ_i is rectangular, then there exist $\theta_{H_i} \in \Theta_{H_i}$ and $\bar{\sigma}_{H_i} \in \bar{\Sigma}_{H_i}$ such that $\sigma_{H_i} = \theta_{H_i}\bar{\sigma}_{H_i}$ for every $H_i \in \mathcal{H}_i$. Thus, if $j = i$ then $\tilde{\beta}_{H_{jk}}(c_{jk}) = \theta_{H_{jk}}(c_{jk})$, while if $j \neq i$ then $\tilde{\beta}_{H_{jk}}(c_{jk})$ is determined by σ_j . Condition (NAM) implies that no information set can appear more than once along any path from the root to a terminal node. Thus, the probability of reaching any terminal node $z \in \mathcal{Z}$ is equal to

$$P_z(\sigma_i, \sigma_{-i}) = \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \theta_{H_{ik}}(c_{ik}),$$

where $\bar{P}_z(\sigma_{-i})$ is given by (8).

Define the strategy $B_i = \times_{H_i \in \mathcal{H}_i} B_{H_i}$ such that $B_{H_i} = \Theta_{H_i}$ for every $H_i \in \mathcal{H}_i$. By applying a similar argument as before, it can be shown that

$$P_z(\beta_i, \sigma_{-i}) = \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \beta_{H_{ik}}(c_{ik}) = \bar{P}_z(\sigma_{-i}) \prod_{k=1}^{K_i} \theta_{H_{ik}}(c_{ik}) = P_z(\sigma_i, \sigma_{-i}),$$

for any $\beta_i \in B_i$, $\sigma_{-i} \in \Sigma_{-i}$ and all $z \in \mathcal{Z}$. This completes the proof. ■

Once again, it is worth emphasizing the role of the rectangularity condition since in case of expected utility maximization an analogous result is obtained under the assumption of perfect recall. In particular, if an extensive game does not satisfy perfect recall then there are mixed strategies that cannot be generated

by either behavioral strategy, i.e., they cannot be represented as the product of the marginal probability $\bar{\sigma}$ and the conditional probability θ . As Definition 8 implies, our rectangularity condition guarantees that such representation holds *always* true. This is because it requires considering only *those* mixed strategies which can be decomposed in terms of marginals and conditionals. Consequently, the assumption of perfect recall becomes superfluous in this case. To see this more clearly, consider the following example.

Example 1 (continued). Consider player 2 who is allowed to move at x_2 , x_4 and x_5 which implies that she does not have perfect recall. Since she has only two choices available at $H_{21} = \{x_2\}$ and $H_{22} = \{x_4, x_5\}$, her Ellsberg behavioral strategy is given by $B_2 = [\underline{\beta}_a, \bar{\beta}_a] \times [\underline{\beta}_f, \bar{\beta}_f]$ where $\beta_a \in [\underline{\beta}_a, \bar{\beta}_a]$ and $\beta_f \in [\underline{\beta}_f, \bar{\beta}_f]$ are the probabilities that she chooses a at H_{21} and f at H_{22} , respectively. Note that (6) and (7) imply that $\sigma_2(s_{c\tilde{c}}) = \theta_{H_{21}}(c)\theta_{H_{22}}(\tilde{c})$ for $c = a, b$ and $\tilde{c} = e, f$. Set $B_{H_{2i}} = \Theta_{H_{2i}}$ for every $i = 1, 2$ and define the rectangular Ellsberg strategy

$$\begin{aligned} \Sigma_2 &= \{(\sigma_2(s_{ae}), \sigma_2(s_{af}), \sigma_2(s_{be}), \sigma_2(s_{bf})) : \sigma_2(s_{c\tilde{c}}) = \beta_c\beta_{\tilde{c}} \text{ for } c = a, b \text{ and } \tilde{c} = e, f \\ \beta_b &= 1 - \beta_a, \beta_e = 1 - \beta_a, \beta_a \in [\underline{\beta}_a, \bar{\beta}_a] \text{ and } \beta_f \in [\underline{\beta}_f, \bar{\beta}_f]\}. \end{aligned}$$

Consider $(\sigma_2(s_{ae}), \sigma_2(s_{af}), \sigma_2(s_{be}), \sigma_2(s_{bf})) = (1/2, 0, 0, 1/2)$. This mixed strategy cannot be generated by either behavioral strategy. However, it cannot belong to any rectangular set Σ_2 either because it cannot be decomposed into marginal and conditional probabilities. ■

We now state the main result of this section.

Theorem 3. *Suppose that an extensive game satisfies (NAM) and the strategy profile Σ_{-i} is given. Then, for any rectangular Ellsberg strategy Σ_i there exists an outcome and payoff equivalent behavioral Ellsberg strategy B_i , and vice versa.*

Proof. Choose the rectangular Ellsberg strategy Σ_i and the behavioral Ellsberg strategy B_i such that $B_{H_i} = \Theta_{H_i}$ for every $H_i \in \mathcal{H}_i$. As shown in the proof of Theorem 1, the strategy Σ_i is then more ambiguous than the strategy B_i implying that $U_i(\Sigma_i, \Sigma_{-i}) \leq U_i(B_i, \Sigma_{-i})$. Likewise, as shown in the proof of Theorem 2, the strategy B_i is then more ambiguous than the strategy Σ_i implying that $U_i(B_i, \Sigma_{-i}) \leq U_i(\Sigma_i, \Sigma_{-i})$. Taken together, the two conditions yield $U_i(B_i, \Sigma_{-i}) = U_i(\Sigma_i, \Sigma_{-i})$. ■

Theorem 3 thus states that rectangular Ellsberg strategies and behavioral Ellsberg strategies are payoff-equivalent. Note however that this result applies only to the ex ante choice of these strategies and, therefore, it cannot guarantee that a player will respect this choice during the course of play. We address this issue in the following section.

4 Dynamic (in)consistency

It is well established in the literature that dynamic consistency may not hold in settings where decision maker's preferences conform to the multiple priors model. One solution to this problem is proposed by Epstein and Schneider (2003) who show that dynamic consistency can be restored if decision maker's *beliefs* about the overall uncertainty satisfy a certain rectangularity condition.¹² Literally speaking, this condition requires that the initial set of priors can be constructed by recursively combining *all* conditional and marginal probabilities corresponding to the information available at each point of time. The key difference of our setting from the one developed by Epstein and Schneider is that each player's beliefs are represented by the opponents' ambiguous strategies which are assumed to be chosen independently. Given such an assumption, we can ask whether it is possible to construct the system of beliefs/strategies satisfying the rectangularity condition identified by Epstein and Schneider in extensive games with perfect recall.

We begin with the example which demonstrates that this may not be possible to do if one maintains the assumption that players choose their strategies in a non-cooperative way.

Example 2. Consider the game in Figure 3 in which each player has perfect recall. In this game, player 1 moves at the information sets H_{11} and H_{12} while players 2 and 3 move at the information set H_{21} and H_{31} , respectively. Suppose

¹²The consistency property in dynamic choice situations appears extensively in the literature. For example, in a setting where a single decision maker faces multi-stage decision problems, Sarin and Wakker (1998) illustrate how consistency can be preserved for a class of nonexpected utility models. Using a somewhat similar approach but focusing on financial trading, Riedel (2004) shows that consistency in risk assessment can be achieved if risk measures satisfy the axioms of coherence and the axiom of dynamic consistency.




Figure 3: Extensive game in Example 2.

that the Ellsberg behavioral strategies of players 2 and 3 are given by

$$B_2 = \{(\beta_A, \beta_B) : \beta_B = 1 - \beta_A \text{ and } \beta_A \in [\underline{\beta}_A, \bar{\beta}_A] \subseteq [0, 1]\}, \quad (11)$$

$$B_3 = \{(\beta_E, \beta_F) : \beta_F = 1 - \beta_E \text{ and } \beta_E \in [\underline{\beta}_E, \bar{\beta}_E] \subseteq [0, 1]\}. \quad (12)$$

Thus, if upon arriving at H_{11} player 1 chooses b with probability one, the set of the induced probabilities λ_{AF} , λ_{AE} and λ_B corresponding to the events that the play will continue along the paths AF , AE and B is given by

$$\begin{aligned} \Lambda &= \{(\lambda_{AE}, \lambda_{AF}, \lambda_B) : \lambda_{AE} = \beta_A \beta_E, \lambda_{AF} = \beta_A (1 - \beta_E), \\ \lambda_B &= 1 - \beta_A, \beta_A \in [\underline{\beta}_A, \bar{\beta}_A] \text{ and } \lambda_E \in [\underline{\beta}_E, \bar{\beta}_E]\}. \end{aligned}$$

In Figure 4, this set is represented by the trapezium $VQSU$ (for the case $\underline{\beta}_E < \bar{\beta}_A \bar{\beta}_E$).

To show that player 1 is not dynamically consistent (i.e., her ex ante choice at H_{12} is different from the one that she actually makes upon arriving at this set),

it suffices to establish that the set Λ is not rectangular relative to the information filtration $\mathcal{F} \equiv (\mathcal{F}_1, \mathcal{F}_2)$ where $\mathcal{F}_1 = \{\{AE, AF, B\}\}$ and $\mathcal{F}_2 = \{\{AE\}, \{AF, B\}\}$. Denote by λ_{AF}^u and λ_B^u the updated probabilities that she assigns to the events that the play has evolved along the paths AF and B upon arriving at H_{12} . λ_{AF}^u and λ_B^u are obtained by updating all probability measures in Λ according to the Bayes' rule which gives rise to the following set of updated priors

$$\Lambda^u = \{(\lambda_{AF}^u, \lambda_B^u) : \lambda_{AF}^u = 1 - \lambda_B^u \text{ and } \lambda_B^u \in [\underline{\lambda}_B^u, \bar{\lambda}_B^u]\},$$

where

$$\underline{\lambda}_B^u = \frac{1 - \bar{\beta}_A}{1 - \bar{\beta}_A \bar{\beta}_E} \text{ and } \bar{\lambda}_B^u = \frac{1 - \underline{\beta}_A}{1 - \underline{\beta}_A \underline{\beta}_E}.$$

In Figure 4, the set Λ^u is given by the thick black segment of the line having slope -1 . Next, as it has been noticed by Epstein and Schneider (2002), in order to construct a rectangular set of priors that induces Λ^u , the marginal probability of reaching H_{12} can be specified in an arbitrary way. Since this probability is equal to $1 - \lambda_{AE}$, we can write the rectangular set of priors as

$$\begin{aligned} \{(\lambda_{AE}, \lambda_{AF}, \lambda_B) \quad : \quad \lambda_{AF} = \lambda_{AF}^u(1 - \lambda_{AE}), \lambda_B = \lambda_B^u(1 - \lambda_{AE}), \\ \lambda_{AE} \in [\underline{\beta}_A \underline{\beta}_E, \bar{\beta}_A \bar{\beta}_E] \text{ and } (\lambda_{AF}^u, \lambda_B^u) \in \Lambda^u\}. \end{aligned}$$

In Figure 4, this set is represented by the trapezium $PTYJ$ which includes the trapezium $VQSU$ (i.e., the set Λ) and the four triangles VTQ , QYS , SJU and UPV . The figure makes it clear that the rectangularity condition necessarily fails in this example.

Remark: Our conclusion would continue to hold if instead of having two players moving at H_{21} and H_{31} we had a single player employing stochastically independent ambiguous randomization devices at those information sets. ■

Example 2 shows that dynamic consistency implicit in Kuhn's Theorem cannot be achieved in every extensive game with imperfect information even if it satisfies perfect recall.¹³ Given that each player's information filtration is induced by the

¹³As Brandenburger (2007) shows (see, Example 3.4), if players are allowed to use only unambiguous strategies then dynamic consistency may be violated if the game does not satisfy perfect recall. Allowing players to use ambiguous strategies in such games makes it even more difficult to achieve dynamic consistency.




Figure 4: Rectangular set of priors in Example 2.

structure of the game, we can ask whether it is possible to derive restrictions on the structure of the game so that dynamic consistency can be implemented by suitably choosing the profile of ambiguous strategies.

As Example 2 suggests, this may be impossible to do in games in which the information structure requires a player to form beliefs over the histories of play in some information set by using the probability distributions over her opponents' actions in more than one information set. To rule out such situations, we impose the following condition ('RI' stands for 'Relevant Information'):

(RI) For each player $k \neq l$, all nodes $x', x'' \in \mathcal{X}_k$ and $x \in \mathcal{X}_l$ and for every choice c' , if $x \rightarrow_{c'} x'$ and $x' \sim x''$ then there exists a choice c'' such that $x \rightarrow_{c''} x''$.

The above condition states that *every* node in player k 's information set is an immediate successor of the *same* node in player l 's information set. By analogy with the temporal precedence relation over nodes, we can say that this condition induces a temporal precedence relation over information sets in the following sense: set H_l is an immediate predecessor of set H_k or set H_k is immediate successor of set H_l if there exists node $x \in H_l$ which is an immediate predecessor of *every* node in H_k .

Note also that recursively applying this condition to all information sets of each player yields that any two paths ending at player k 's information set must pass through the same sequence of nodes of that player and contain identical actions at all such nodes. This in turn implies that player k must have perfect recall. The following lemma confirms this intuition.

Lemma 1. *If an extensive game satisfies (RI) then it satisfies (PR).*

Proof. Fix any $x'_1, x'_2 \in \mathcal{X}_k$ such that $x'_1 \sim x'_2$. If (PR) is not satisfied, then there must exist $x_1, x_2 \in \mathcal{X}_k$ and $t_1, t_2 \in \mathcal{T}$ such that $x_i \rightarrow_{c_i} t_i$, $t_i \succsim x'_i$ for each $i = 1, 2$ and $c_1 \neq c_2$. Since each node has exactly one immediate predecessor, then the paths from x_1 to x'_1 and from x_2 to x'_2 must differ. As a result, there must exist $t'_1, t'_2 \in \mathcal{T}$ such that $t'_i \rightarrow_{c'_i} x'_i$ for each $i = 1, 2$ and $c'_1 \neq c'_2$. This contradicts (RI). ■

Although condition (RI) guarantees that every time when a player is called upon to move she needs to make a conjecture about her opponent's moves in




Figure 5: Extensive game in Example 3.

only one immediately preceding information set, still it might not suffice to ensure dynamic consistency. The following example illustrates this point.

Example 3. The game in Figure 5 satisfies (RI). In this game, player 1 moves at H_{11} , H_{12} and H_{13} while player 2 moves at H_{21} . A distinct feature of this game is that player 1's choice at H_{11} influences her information filtration \mathcal{F} . In particular, \mathcal{F} is given by $(\mathcal{F}_1, \mathcal{F}_2)$ where $\mathcal{F}_1 = \{A, B, C\}$ and $\mathcal{F}_2 = \{\{A\}, \{B, C\}\}$ if she chooses a , and it is given by $(\mathcal{F}'_1, \mathcal{F}'_2)$ where $\mathcal{F}'_1 = \mathcal{F}_1$ and $\mathcal{F}'_2 = \{\{A, B\}, \{C\}\}$ if she chooses b . The dynamic consistency fails in this example because the requirement that player 2's strategy $B_2 = B_{H_{21}}$ is rectangular with respect to the filtration $(\mathcal{F}_1, \mathcal{F}_2)$ is incompatible with the requirement that it is rectangular with respect to the filtration $(\mathcal{F}'_1, \mathcal{F}'_2)$.

To show this, consider Figure 6 which gives a two-dimensional representation of any probability measure in B_2 . In the figure, the horizontal axis corresponds to the probability of state C , denoted by β_C , while the vertical axis corresponds to the probability of state B , denoted by β_B . By applying a similar argument as the one in Example 2, it can be shown that typical sets which are rectangular relative to the filtrations $(\mathcal{F}_1, \mathcal{F}_2)$ and $(\mathcal{F}'_1, \mathcal{F}'_2)$ would be given by the trapezia $QYJP$ and $VTSU$, respectively. The figure makes it clear that the set B_2 cannot satisfy both




Figure 6: Rectangular sets of priors in Example 3.

rectangularity conditions. ■

Guided by this example, we impose another condition which we call ‘Consistency of Information Filtration’ (hereafter, ‘CIF’).

(CIF) For each player $k \neq l$, all nodes $x_1, x_2, \hat{x}_1, \hat{x}_2 \in \mathcal{X}_k$ and $x, \hat{x} \in \mathcal{X}_l$ and for all choices c_1 and c_2 , if $x \sim \hat{x}$, $x_1 \sim x_2$, $x \rightarrow_{c_1} x_1$, $x \rightarrow_{c_2} x_2$, $\hat{x} \rightarrow_{c_1} \hat{x}_1$ and $\hat{x} \rightarrow_{c_2} \hat{x}_2$ then $\hat{x}_1 \sim \hat{x}_2$.

Condition (CIF) states that if player k cannot distinguish between two choices of player l at one node of the immediately preceding information set then the same must hold true for every other node in that information set. It thereby implies that there is no reason for player k to base her conjecture, as to what player l did right before the time when she is called upon to move, on the history of play up to that time (hence, the term ‘Consistency of Information Filtration’).

In what follows, we shall restrict attention to the class of extensive games that satisfy conditions (RI) and (CIF). This class is sufficiently more restrictive than

the class of extensive games with imperfect information but is sufficiently less restrictive than the class of extensive games with perfect information.

As is explained above, condition (RI) implies that what matters for player k upon arriving at the information set H_k is the knowledge of the probability distribution over player l 's choices at the immediately preceding information set H_l . In turn, condition (CIF) implies that the partition of the event space which comprises all the choices available at H_l is invariable across the nodes in H_l .

Denote by $\mathcal{H}_k(H_l)$ the collection of H_k 's which are immediate successors of H_l and by $\mathcal{C}_k(H_l) \subseteq \mathcal{C}(H_l)$ the set of choices in H_l which lead to $H_k \in \mathcal{H}_k(H_l)$, i.e.,

$$\mathcal{C}_k(H_l) \equiv \{c : c \in \mathcal{C}(H_l) \text{ and } x \rightarrow_c x' \text{ for some } x \in H_l \text{ and all } x' \in H_k\}.$$

Also, denote by $\mathcal{C}_O(H_l)$ the set of choices in H_l which lead to terminal nodes, i.e.,

$$\mathcal{C}_O(H_l) \equiv \mathcal{C}(H_l) \setminus \bigcup_{H_k \in \mathcal{H}_k(H_l)} \mathcal{C}_k(H_l).$$

Thus, at any point of time before the game reaches any $H_k \in \mathcal{H}_k(H_l)$ the information structure of player k (regarding continuation play starting at H_l) can be represented by the filtration $\mathcal{F}_{H_l} \equiv (\mathcal{F}_{H_l}^0, \mathcal{F}_{H_l}^1)$ where

$$\mathcal{F}_{H_l}^0 = \mathcal{C}(H_l) \text{ and } \mathcal{F}_{H_l}^1 = \{\{\mathcal{C}_O(H_l)\}, \{\mathcal{C}_k(H_l)\}_{H_k \in \mathcal{H}_k(H_l)}\}.$$

For any probability measure $\beta_{H_l} \in B_{H_l}$ define a marginal, or one-step-ahead, distribution $\beta_{H_l}^{0+}$ which is the restriction of β_{H_l} to $\mathcal{F}_{H_l}^0$. Denote also by $\beta_{H_l}^1$ the conditional distribution (i.e., the Bayesian update) given $\mathcal{F}_{H_l}^1$. By applying the standard decomposition in terms of marginals and conditionals, we have

$$\beta_{H_l} = \int \beta_{H_l}^1 d\beta_{H_l}^{0+}. \quad (13)$$

Denote by $B_{H_l}^{0+}$ the set of one-step-ahead distributions and by $B_{H_l}^1$ the set of conditional distributions induced by all the probability measures in B_{H_l} . Following Epstein and Schneider (2003), we define a rectangular set of priors.

Definition 12 The set B_{H_l} is \mathcal{F}_{H_l} -rectangular if

$$B_{H_l} = \left\{ \int \beta_{H_l}^1 d\beta_{H_l}^{0+} : \beta_{H_l}^1 \in B_{H_l}^1 \text{ and } \beta_{H_l}^{0+} \in B_{H_l}^{0+} \right\}.$$

Using this, we define a rectangular Ellsberg behavioral strategy.

Definition 13 The Ellsberg behavioral strategy $B_l = \times_{H_l \in \mathcal{H}_l} B_{H_l}$ is rectangular if B_{H_l} is \mathcal{F}_{H_l} -rectangular for every $H_l \in \mathcal{H}_l$.

Definition 14 The profile of Ellsberg behavioral strategies $B = (B_1, \dots, B_N)$ is rectangular if B_l is rectangular for every $l \in \mathcal{N}$.

Since, by assumption, each player k is averse to uncertainty created by her opponents playing the strategy profile B_{-k} , her ex ante expected utility from using strategy $\beta_k \in \mathcal{B}_k$ is given by

$$U_k(\beta_k, B_{-k}) = \min_{\beta_{-k} \in B_{-k}} \sum_{z \in \mathcal{Z}} P_z(\beta_k, \beta_{-k}) u_k(z),$$

where $P_z(\beta_k, \beta_{-k})$ is the probability of reaching a terminal node $z \in \mathcal{Z}$ induced by the strategy profile (β_k, β_{-k}) .

Definition 15 Given the opponents' strategy profile B_{-k} , the strategy $\hat{\beta}_k$ is ex ante optimal for player k if $U_k(\hat{\beta}_k, B_{-k}) \geq U_k(\beta_k, B_{-k})$ for all $\beta_k \in \mathcal{B}_k$.

Denote by β_{H_k+} the continuation strategy induced by β_k at H_k and all the information sets that follow H_k , by B_{-k, H_k+} the components of B_{-k} that correspond to the information sets that follow H_k , and by Υ_{H_k} the set-valued beliefs of player k over the histories in H_k . Given Υ_{H_k} and B_{-k, H_k+} , define player k 's expected utility from playing β_{H_k+} at H_k as

$$U_k(\beta_{H_k+}, \Upsilon_{H_k}, B_{-k, H_k+} | H_k) = \min_{\substack{\beta_{-k, H_k+} \in B_{-k, H_k+} \\ \gamma_{H_k} \in \Upsilon_{H_k}}} \sum_{z \in \mathcal{Z}} P_z(\beta_{H_k+}, \gamma_{H_k}, \beta_{-k, H_k+} | H_k) u_k(z),$$

where $P_z(\beta_{H_k+}, \gamma_{H_k}, \beta_{-k, H_k+} | H_k)$ is the probability of reaching a terminal node $z \in \mathcal{Z}$ conditional on H_k .

Definition 16 Given the opponents' strategy profile B_{-k} , the strategy $\widehat{\beta}_k$ is optimal at $H_k \in \mathcal{H}_k$ if $U_k(\widehat{\beta}_{H_k+}, \Upsilon_{H_k}, B_{-k, H_k+} | H_k) \geq U_k(\beta_{H_k+}, \Upsilon_{H_k}, B_{-k, H_k+} | H_k)$ for all $\beta_k \in \mathcal{B}_k$.

We now state the main result of this section.

Theorem 4. Suppose that (i) an extensive game satisfies (RI) and (CIF), and (ii) the strategy profile B_{-k} is rectangular. Then, $\widehat{\beta}_k$ is ex ante optimal if it is optimal at each $H_k \in \mathcal{H}_k$.

Proof. Suppose that $\widehat{\beta}_k$ is optimal at each $H_k \in \mathcal{H}_k$ but it is not ex ante optimal. Then, there must exist β_k such that $\beta_{H_k+} \neq \widehat{\beta}_{H_k+}$ for some $H_k \in \mathcal{H}_k$ and $U_k(\beta_k, B_{-k}) > U_k(\widehat{\beta}_k, B_{-k})$.

Let $\beta \equiv (\beta_k, \beta_{-k})$, $\beta_{H_k+} \equiv (\beta_{H_k+}, \beta_{-k, H_k+})$ and

$$\begin{aligned} \mathbf{E}_\beta[u_k(z)] &\equiv \sum_{z \in \mathcal{Z}} P_z(\beta) u_k(z), \\ \mathbf{E}_{\gamma_{H_k}, \beta_{H_k+}}[u_k(z) | H_k] &\equiv \sum_{z \in \mathcal{Z}} P_z(\gamma_{H_k}, \beta_{H_k+} | H_k) u_k(z). \end{aligned}$$

Since the game satisfies (RI), there must exist information set $H_l \in \mathcal{H}_l$ which immediately precedes H_l . Denote by \mathcal{Z}_{H_l} the set of terminal nodes such that for every $z \in \mathcal{Z}_{H_l}$ the path \bar{z} passes through H_l .¹⁴ In which case, the expression for $\mathbf{E}_\beta[u_k(z)]$ can be written as

$$\mathbf{E}_\beta[u_k(z)] = \sum_{z \in \mathcal{Z}_{H_l}} P_z(\beta) u_k(z) + \sum_{z \in \mathcal{Z} \setminus \mathcal{Z}_{H_l}} P_z(\beta) u_k(z). \quad (14)$$

Denote by $z_c \in \mathcal{Z}_{H_l}$ the terminal node that follows $c \in \mathcal{C}_O(H_l)$, by $P_{H_l}(\beta)$ the probability of reaching H_l ,¹⁵ and by $P_z(\beta_{H_k+}, \beta_{-k, H_k+} | c)$ the probability of reaching z conditional on H_l and $c \in \mathcal{C}_k(H_l)$.

Since the game satisfies (CIF), the partition of $\mathcal{C}(H_l)$ relative to the information sets that immediately follow H_l is invariable across the nodes in H_l . Hence, the

¹⁴Note that for the class of games that we consider the path \bar{z} is unique for every $z \in \mathcal{Z}$ and it crosses any information set only once.

¹⁵We write $P_{H_l}(\beta)$ for notational simplicity, even though only those components of β that correspond to the information sets which appear along the paths ending at H_l will affect this probability.

first term in (14) can be written as

$$\begin{aligned} \sum_{z \in \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}) u_k(z) &= P_{H_l}(\boldsymbol{\beta}) \left(\sum_{c \in \mathcal{C}_O(H_l)} \beta_{H_l}(c) u_k(z_c) \right. \\ &\quad \left. + \sum_{H_k \in \mathcal{H}_k(H_l)} \sum_{c \in \mathcal{C}_k(H_l)} \beta_{H_l}(c) \sum_{z \in \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}_{H_k+|c}) u_k(z) \right). \end{aligned} \quad (15)$$

Using the fact that β_{H_l} can be decomposed into marginals and conditionals, according to (13), we have

$$\begin{aligned} \sum_{z \in \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}) u_k(z) &= P_{H_l}(\boldsymbol{\beta}) \left(\beta_{H_l}^{0+}(O) \sum_{c \in \mathcal{C}_O(H_l)} \beta_{H_l}^1(c) u_k(z_c) \right. \\ &\quad \left. + \sum_{H_k \in \mathcal{H}_k(H_l)} \beta_{H_l}^{0+}(H_k) \sum_{c \in \mathcal{C}_k(H_l)} \beta_{H_l}^1(c) \sum_{z \in \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}_{H_k+|c}) u_k(z) \right). \end{aligned} \quad (16)$$

The rectangularity condition also implies that $\Upsilon_{H_k} = B_{H_k}^1$ because each element in $B_{H_k}^1$ corresponds to the Bayesian update of the histories in H_k . Hence,

$$\mathbf{E}_{\beta_{H_k}^1, \beta_{H_k+}} [u_k(z) | H_k] = \sum_{c \in \mathcal{C}_k(H_l)} \beta_{H_l}^1(c) \sum_{z \in \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}_{H_k+|c}) u_k(z). \quad (17)$$

By supposition, the Ellsberg behavioral strategy B_{-k} is rectangular which implies that the set B_{H_l} is \mathcal{F}_{H_l} -rectangular. Due to this rectangularity condition, the problem of minimizing any function over the set B_{H_l} can be decomposed into two independent problems of minimizing this function over the sets $B_{H_l}^1$ and $B_{H_l}^{0+}$. Denoting by $B_{-k \setminus H_l}$ the collection of the components of B_{-k} except B_{H_l} , the expression for $U_k(\beta_k, B_{-k})$ can then be written as

$$\begin{aligned} U_k(\beta_k, B_{-k}) &= \min_{\beta_{-k} \in B_{-k}} \mathbf{E}_{\beta_k, \beta_{-k}} [u_k(z)] \\ &= \min_{\beta_{-k \setminus H_l} \in B_{-k \setminus H_l}} \min_{\substack{\beta_{H_l}^1 \in B_{H_l}^1 \\ \beta_{H_l}^{0+} \in B_{H_l}^{0+}}} \mathbf{E}_{\beta_k, \beta_{H_l}^1, \beta_{H_l}^{0+}, \beta_{-k \setminus H_l}} [u_k(z)], \end{aligned} \quad (18)$$

where the minimand in the above expression is given by (using (14), (16) and (17))

$$\begin{aligned}
& \mathbf{E}_{\beta_k, \beta_{H_l}^1, \beta_{H_l}^{0+}, \beta_{-k \setminus H_l}} [u_k(z)] \tag{19} \\
&= P_{H_l}(\boldsymbol{\beta}) \left(\beta_{H_l}^{0+}(O) \sum_{c \in \mathcal{C}_O(H_l)} \beta_{H_l}^1(c) u_k(z_c) \right. \\
&\quad \left. + \sum_{H_k \in \mathcal{H}_k(H_l)} \beta_{H_l}^{0+}(H_k) \mathbf{E}_{\beta_{H_l}^1, \beta_{H_k+}} [u_k(z) | H_k] \right) \\
&\quad + \sum_{z \in \mathcal{Z} \setminus \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}) u_k(z).
\end{aligned}$$

As we explained above, $\Upsilon_{H_k} = B_{H_k}^1$ for every $H_k \in \mathcal{H}_k(H_l)$. Using (18), (19) and the expression for $U_k(\beta_{H_k+}, \Upsilon_{H_k}, B_{-k, H_k+} | H_k)$, we thus have¹⁶

$$\begin{aligned}
U_k(\beta_k, B_{-k}) &= \min_{\beta_{-k \setminus H_l} \in B_{-k \setminus H_l}} \min_{\substack{\beta_{H_l}^1 \in B_{H_l}^1 \\ \beta_{H_l}^{0+} \in B_{H_l}^{0+}}} P_{H_l}(\boldsymbol{\beta}) \left(\beta_{H_l}^{0+}(O) \sum_{c \in \mathcal{C}_O(H_l)} \beta_{H_l}^1(c) u_k(z_c) \right. \\
&\quad \left. + \sum_{H_k \in \mathcal{H}_k(H_l)} \beta_{H_l}^{0+}(H_k) U_k(\beta_{H_k+}, B_{H_k}^1, B_{-k, H_k+} | H_k) \right) \\
&\quad + \sum_{z \in \mathcal{Z} \setminus \mathcal{Z}_{H_l}} P_z(\boldsymbol{\beta}) u_k(z).
\end{aligned}$$

From the above expression it follows that if $\beta_{H_k+} \neq \widehat{\beta}_{H_k+}$ for some $H_k \in \mathcal{H}_k(H_l)$ and $\widehat{\beta}_{H_k+}$ maximizes U_k then it must be

$$U_k(\beta_{-H_k}, \beta_{H_k+}, B_{-k}) \leq U_k(\widehat{\beta}_k, B_{-k}),$$

where β_{-H_k} stands for all components of β_k other than β_{H_k+} . This contradicts the supposition that $U_k(\beta_k, B_{-k}) > U_k(\widehat{\beta}_k, B_{-k})$. ■

¹⁶We write $\beta_{-k \setminus H_l} \in B_{-k \setminus H_l}$ in the minimization problem for notational simplicity, even though the components that correspond to the information sets that follow H_l are reflected in $U_k(\beta_{H_k+}, B_{H_k}^1, B_{-k, H_k+} | H_k)$ for every $H_k \in \mathcal{H}_k(H_l)$.

5 Concluding remarks

In this paper, we make three primary points. First, we prove that finite extensive games with objective ambiguity can equivalently be analyzed using ambiguous behavioral and ambiguous mixed strategies, as long as the latter strategies satisfy the rectangularity condition. In contrast to the standard Kuhn's Theorem, this result does not rely on the assumption of perfect recall. However, it applies only to the ex ante choice of ambiguous strategies and, as such, it cannot be relied upon in ensuring dynamic consistency. Second, we show by means of examples that even in games satisfying perfect recall it might be impossible to achieve dynamic consistency if one maintains the assumptions that players choose their (ambiguous) strategies in a non-cooperative way. Finally, we argue that, in order to ensure dynamic consistency, one must impose restrictions not only on the ex ante choice of ambiguous strategies but also on the structure of a game. Overall, our results lay the foundation for studying equilibrium notions in games where players are capable of creating ambiguity.

References

- [1] Aryal, G., and R., Stauber, (2013), "A Note on Kuhn's Theorem with Ambiguity Averse Players", *Working paper*, Australian National University.
- [2] Brandenburger, A., (2007), "A Note on Kuhn's Theorem", In: van Benthem, J., Gabbay, D., Lowe, B. (Eds.), *Interactive Logic. Proceedings of the 7th Augustus de Morgan Workshop*, London. Texts in Logic and Games 1. Amsterdam University Press, 71 - 88.
- [3] Bonanno, G., (2004), "Memory and Perfect Recall in Extensive Games", *Games and Economic Behavior*, 47 (2), 237 - 256.
- [4] D'Amato, A., (2010), "Purposeful Ambiguity as International Legal Strategy: The Two China Problem", *Working paper*, Northwestern University School of Law.
- [5] Eisenberg, E., (1984), "Ambiguity as Strategy in Organizational Communication", *Communication monographs*, 51 (3), 227 - 242.

- [6] Epstein, L., and M., Schneider, (2013), “Recursive Multiple-Priors”, *Journal of Economic Theory*, 113 (1), 1 - 31.
- [7] Gajdos, T., Hayashi, T., Tallon, J., and J. Vergnaud, (2008), “Attitude Toward Imprecise Information”, *Journal of Economic Theory*, 140 (1), 27 - 65.
- [8] Gilboa, I., and D., Schmeidler, (1989), “Maxmin Expected Utility with Non-Unique Prior”, *Journal of Mathematical Economics*, 18 (2), 141 - 153.
- [9] Hanany, E., and P., Klibanoff, (2007), “Updating Preferences with Multiple Priors”, *Theoretical Economics*, 2, 261 - 298.
- [10] Kuhn, H., (1953), “Extensive Games and the Problem of Information”, *Contributions to the Theory of Games II*, 193 - 216.
- [11] Ma, C., (2000), “Uncertainty Aversion and Rationality in Games of Perfect Information”, *Journal of Economic Dynamics and Control*, 24(3), 451 - 482.
- [12] Riedel, F., (2004), “Dynamic Coherent Risk Measures”, *Stochastic Processes and their Applications*, 112(92), 185 - 200.
- [13] Riedel, F., and L. Sass, (2013), “Ellsberg Games”, *Theory and Decision*, forthcoming.
- [14] Ritzberger, K., (1999), “Recall in Extensive Games”, *International Journal of Game Theory*, 28 (1), 69 - 87.
- [15] Sarin, R., and P., Wakker, (1998), “Dynamic Choice and NonExpected Utility”, *Journal of Risk and Uncertainty*, 17, 87 - 119.