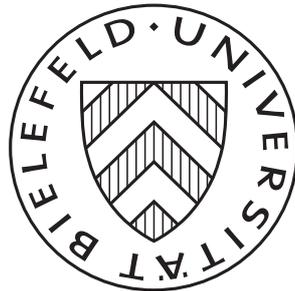


April 2013

Kuhn's Theorem for Extensive Form Ellsberg Games

Linda Sass



Kuhn's Theorem for Extensive Form Ellsberg Games

Linda Sass ^{*†}

Bielefeld University and Université Paris 1 Panthéon-Sorbonne

This Version: April 29, 2013

Riedel and Sass (2013) propose a framework for normal form games where players can use imprecise probabilistic devices. We extend this strategic use of objective ambiguity to extensive form games. We show that with rectangularity of Ellsberg strategies we have dynamic consistency in the sense of Kuhn (1953): rectangular Ellsberg strategies are equivalent to Ellsberg behavior strategies. We provide an example for our result and define *Ellsberg equilibrium* in such extensive form Ellsberg games.

Key words and phrases: Knightian Uncertainty in Games, Objective Ambiguity, Strategic Ambiguity, Extensive Form Ellsberg Games, Kuhn's Theorem, Rectangularity

JEL subject classification: C72, C73, D81

1 Introduction

Following the concept of strategic ambiguity in Ellsberg games introduced in Riedel and Sass (2013), it is straightforward that imprecise probabilistic devices in form of Ellsberg

* I thank Raffaël Giraud, Johannes Höwing, Peter Klibanoff and Frank Riedel for helpful discussions during the development of the paper, as well as Klaus Ritzberger and Jan-Henrik Steg for assistance with some technical details.

† Financial support through the German Research Foundation, International Research Training Group EBIM "Economic Behavior and Interaction Models", through the DFH-UFA, French-German University, through the Institute of Mathematical Economics at Bielefeld University and through the Rectorate of Bielefeld University is gratefully acknowledged.

urns can also be used as strategic instruments in dynamic games. In this paper we develop the theoretical framework to analyze such extensive form Ellsberg games. In analogy to mixed and behavioral strategies in classic extensive form theory, players can use such imprecise probabilistic devices either over their set of pure strategies, or they can place them at every single decision point in the process of the game. These strategies we call Ellsberg strategies, or Ellsberg behavior strategies.

The players in our extensive form Ellsberg games are averse to objective ambiguity according to Gajdos, Hayashi, Tallon, and Vergnaud (2008); they evaluate sets of probability distributions using a maxmin rule, as in maxmin expected utility by Gilboa and Schmeidler (1989). Now, considering dynamic games with ambiguity-averse players leads directly to the question of dynamic consistency of the players' preferences. We decided to adapt the notion of rectangularity by Epstein and Schneider (2003) to obtain Ellsberg strategies for which the preferences are dynamically consistent.

Under the assumption of rectangularity, we prove the main result (Theorem 7) of this paper, which is a generalization of Kuhn's Theorem (Kuhn (1953)): for every rectangular Ellsberg strategy there exists an Ellsberg behavior strategy that yields the same maxmin expected utility; and for every Ellsberg behavior strategy there exists at least one Ellsberg strategy with the same payoff. In a detailed example we provide intuition for rectangularity, and calculate the equivalent strategies.

Moreover, we extend the notion of Ellsberg equilibrium defined in Riedel and Sass (2013) to extensive form games. From Theorem 7 we have that Ellsberg equilibrium in Ellsberg strategies and Ellsberg equilibrium in Ellsberg behavior strategies are equivalent.

The paper is organized as follows. We define extensive form Ellsberg games in Section 2. We discuss dynamic consistency in Section 3 and develop a formalism to translate the notion of rectangularity by Epstein and Schneider (2003) to extensive form Ellsberg games. This property of Ellsberg strategies is then used to prove a generalization of Kuhn's Theorem (Kuhn (1953)) for extensive form Ellsberg games in Section 4. We present a detailed example to illustrate the result. In Section 5 thereafter, we define Ellsberg equilibrium in extensive form Ellsberg games and finally, in Section 6, we compare Ellsberg equilibria to other extensive form solution concepts. We conclude in Section 7.

2 Extensive Form Ellsberg Games

We focus on dynamic games that can be represented in finite game trees. This implies a finite number n of players and a finite number of moves for each player. For simplicity, we do not allow any chance moves. We use the model of Osborne and Rubinstein (1994) for extensive form games with imperfect information, and extend it to allow for imprecise

probabilistic devices.

Definition 1. A finite extensive form game is a tuple $(N, H, W, l, (\mathcal{I}_i), (u_i))$ whose components are defined by:

- A finite set N of players i ;
- A finite set H of sequences of actions $(a^k)_{k=1, \dots, K}$ with $\emptyset \in H$, which represent the histories;
- A set $W \subset H$ of terminal histories; for each $h \in H \setminus W$, $A(h) := \{a \mid (h, a) \in H\}$ defines the set of actions available after history h ;
- A player function $l : H \setminus W \rightarrow N$;
- For each player $i \in N$ an information partition \mathcal{I}_i of $\{h \in H \mid l(h) = i\}$ with the property that $A(h) = A(h')$ whenever $h, h' \in I_i \in \mathcal{I}_i$. For $I_i \in \mathcal{I}_i$ we denote by $A(I_i)$ the set $A(h)$ and by $l(I_i)$ the player $l(h)$ for any $h \in I_i$.
- For each player $i \in N$ the preferences on lotteries over W can be represented by a Von Neumann and Morgenstern expected utility function u_i .

The set of pure strategies of a player i is denoted S_i . The utility of a pure strategy profile $s \in S_1 \times \dots \times S_n$ is simply $u_i(w)$, where w is the terminal history that results when the profile s is played. The set of mixed strategies P_i is denoted ΔS_i . We assume a product structure on $\Delta S := \Delta S_1 \times \dots \times \Delta S_n$ and thus have stochastic independence of mixed strategies. Thus, if $P = (P_1, \dots, P_n) \in \Delta S$ and $s \in S$, then

$$P(s) = P(s_1, \dots, s_n) = \prod_{i=1}^n P_i(s_i) \text{ for all } s \in S.$$

The expected utility $u = (u_1, \dots, u_n) : S \rightarrow \mathbb{R}^n$ of a mixed strategy profile P (we use the same notation as for the utility of a pure strategy) is $u_i(P) = \sum_{s \in S} P(s)u_i(s)$. A behavioral strategy of player i is a function $\theta_i = (\theta_i(I_i))_{I_i \in \mathcal{I}_i}$ with $\theta_i(I_i) \in \Delta A(I_i)$ that assigns to each information set I_i of player i a probability distribution over the set of actions available at I_i . The set of behavioral strategies of player i is denoted \mathcal{O}_i , the set of profiles $\theta = (\theta_1, \dots, \theta_n)$ of behavioral strategies by $\mathcal{O} = \mathcal{O}_1 \times \dots \times \mathcal{O}_n$.

We assume that in addition to classic randomizing devices, players can use imprecise probabilistic devices to choose among their pure strategies or among their available actions at each information set. Therefore, in addition to classical pure, mixed and behavioral strategies, we define Ellsberg strategies and Ellsberg behavioral strategies as convex and compact sets of mixed and behavioral strategies.

Definition 2. An Ellsberg strategy of player i is a convex and compact set \mathcal{P}_i of probability distributions $P_i : S_i \rightarrow \mathbb{R}_+$, such that $\sum_{s_i \in S_i} P_i(s_i) = 1$.

Definition 3. An Ellsberg behavior strategy Θ_i of player i is a function that assigns to each information set I_i of player i a convex and compact set of probability distributions θ_i over the set of actions available at I_i .

A profile of Ellsberg behavior strategies is denoted $\Theta = (\Theta_1, \dots, \Theta_n)$, and Θ is constituted of profiles of behavior strategies $\theta = (\theta_1, \dots, \theta_n)$. For any history $h \in I_i \in \mathcal{I}_i$ and action $a \in A(h)$ we denote by $\theta_i(h)(a)$ the probability $\theta_i(I_i)(a)$ assigned by $\theta_i(I_i)$ to the action a .

For a profile $P = (P_1, \dots, P_n)$ of mixed strategies, we define the outcome Π_P of P to be the probability distribution over the terminal histories that results when each player i follows the precepts of P_i . The probability that Π_P assigns to a terminal history w is $\prod_{i \in N} \pi_i(w)$ where $\pi_i(w)$ is the sum of the probabilities according to P_i of all the pure strategies of player i that are *consistent* (i.e. that result in the terminal history w) with w . The same way we define Π_θ for a profile of behavioral strategies. The probability that Π_θ assigns to $w = (a^1, \dots, a^K)$ is $\prod_{k=0}^{K-1} \theta_{l(a^1, \dots, a^k)}(a^1, \dots, a^k)(a^{k+1})$, where for $k = 0$ the history (a^1, \dots, a^k) is the initial history.

For a profile of Ellsberg strategies \mathcal{P} or Ellsberg behavior strategies Θ , the outcomes $\Pi_{\mathcal{P}}$ and Π_Θ are sets of probabilities defined as

$$\Pi_{\mathcal{P}}(w) := \left\{ \Pi_P(w) \mid P \in \mathcal{P} \right\}, \quad (1)$$

$$\Pi_\Theta(w) := \left\{ \Pi_\theta(w) \mid \theta \in \Theta \right\}.$$

Finally we specify players' preferences over the new strategic devices. We assume that the players are ambiguity-averse according to Gajdos, Hayashi, Tallon, and Vergnaud (2008) in the special case $\phi = \text{id}$. The utility of both an Ellsberg and Ellsberg behavior strategy profile \mathcal{P}, Θ is then evaluated with a maxmin rule using the worst case probability distribution. We have stochastic independence of the Ellsberg strategies by assuming a product structure for every mixed strategy profile P contained in $\mathcal{P} = (\mathcal{P}_1, \dots, \mathcal{P}_n)$.¹ Thus,

¹ In the definition of a normal form Ellsberg game we assume that the Ellsberg urns $(\Omega_i, \mathcal{F}_i, \mathcal{P}_i)$ of all players $i \in N$ are stochastically independent. This is done by using product spaces as first suggested by Gilboa and Schmeidler (1989) (instead of Ellsberg urns they speak of “non-unique probability spaces” $(\Omega, \mathcal{F}, \mathcal{P})$, p. 150 therein). We define the product $(\Omega, \mathcal{F}, \mathcal{P})$ of n Ellsberg urns $(\Omega_i, \mathcal{F}_i, \mathcal{P}_i)$, $i = 1, \dots, n$, as follows: $\Omega := \Omega_1 \times \dots \times \Omega_n$, $\mathcal{F} := \mathcal{F}_1 \otimes \dots \otimes \mathcal{F}_n$, and \mathcal{P} is the closed convex hull of the set of product measures, $\mathcal{P} := \bar{\text{co}} \{ P_1 \otimes \dots \otimes P_n \mid P_1 \in \mathcal{P}_1, \dots, P_n \in \mathcal{P}_n \}$. This way the Ellsberg urns are stochastically independent. Different notions of stochastic independence in the context of ambiguity aversion have been discussed in the literature, see for example Klibanoff (2001), Bade (2011b) and Bade (2011a). In the present context of objective ambiguity in the form of Ellsberg urns the above notion seems the most natural.

the utility of an Ellsberg strategy profile \mathcal{P} for player i is

$$U_i(\mathcal{P}) = \min_{P \in \mathcal{P}} \sum_{w \in W} \Pi_P(w) u_i(w). \quad (2)$$

The utility of an Ellsberg behavior strategy profile Θ for a player i is then

$$U_i(\Theta) = \min_{\theta \in \Theta} \sum_{w \in W} \Pi_\theta(w) u_i(w). \quad (3)$$

With the ambiguity-averse preference representation at hand, we can now give the definition of an extensive form Ellsberg game.

Definition 4. *An extensive form Ellsberg game with n players is a tuple $(N, H, W, l, (\mathcal{I}_i)(U_i))$ whose components satisfy the conditions in Definition 1, and $U = (U_1, \dots, U_n)$ is the ambiguity-averse preference representation in (2) and (3).*

By specifying in the definition how players evaluate Ellsberg strategies we allow players to use Ellsberg urns as imprecise probabilistic devices. Of course, it is still possible to play classic mixed and behavioral strategies. Then U_i reduces to the Von Neumann and Morgenstern expected utility u_i .

3 Rectangular Ellsberg Strategies and Dynamic Consistency

Dynamic consistency requires that preferences over outcomes at some point in the game will not be reversed or contradicted at a later point. A compelling feature of classic extensive form games lies in the fact that for a rational player it does not matter whether he plans his strategy in advance or executes it whenever one of his information sets is reached; this means that additional information arriving at some point in the game does not ex post influence the choice of his strategy. This equivalence is captured by Kuhn's Theorem which we discuss and extend to Ellsberg games in the next section. When players are expected utility maximizers, Kuhn's Theorem establishes a connection between conditional preferences at different points in the game tree via Bayes' updating rule: when players update in that way, their choices following the expected utility model are dynamically consistent.

Updating of ambiguity-averse preference representations and its relation to dynamic consistency has been extensively analyzed. Gilboa and Schmeidler (1993) propose two updating rules for multiple-prior expected utilities, full Bayesian updating and maximum

likelihood updating, which are in general not dynamically consistent. Epstein and Schneider (2003) characterize a particular set of priors which they call *rectangular*, for which full Bayesian updating is dynamically consistent. They show, more precisely, that multiple-prior utilities like those used in Ellsberg games are dynamically consistent if and only if each set of priors is rectangular and it is updated by Bayes' Rule applied prior by prior. Riedel (2004) uses rectangularity to prove dynamic consistency of dynamic coherent risk measures. Later Hanany and Klibanoff (2007) use a weaker notion of dynamic consistency to obtain a larger set of dynamically consistent updating rules for maxmin expected utilities.

We translate the property of rectangularity in Epstein and Schneider (2003) to Ellsberg games and thus obtain dynamic consistency in our setting. We start by describing the information structure of an extensive form Ellsberg game $(N, H, W, l, (\mathcal{I}_i), (U_i))$ with the help of a filtration $\{\mathcal{W}_t\}_0^T$ on a state space W . This is only possible, when the game satisfies perfect recall, see Osborne and Rubinstein (1994) Definition 203.3. Thus, we assume that the games in this section satisfy perfect recall. The information structure of the game is captured in form of a tree with root $\emptyset \in H$. The root \emptyset represents time 0, when the game has not yet started. The sequence of decisions taken by the players defines the information available: at every time t there is a finite information partition \mathcal{J}_t consisting of the information sets J_t available at that time.

Let now the set W of terminal histories be the state space. Then, staying with the model of Epstein and Schneider (2003), we assume that \mathcal{W}_0 is trivial (that is, consists only of W and \emptyset) and that for each t , \mathcal{W}_t is generated by the finite information partition \mathcal{J}_t at time t . Then $\mathcal{W}_t(w)$ denotes the partition component J_t at which the terminal history w is still possible, along with a number of other terminal histories. At $t + 1$, the number of possible terminal histories is narrowed down again. At time T , $\mathcal{W}_T(w)$ is the unique terminal history that has materialized.

Epstein and Schneider (2003) consider lotteries over adapted consumption processes, and preferences over \mathcal{W}_t -measurable acts from W into such lotteries. Then, for simplicity, they assume full support (Axiom 5 therein): every non-empty event in \mathcal{W}_T is considered possible at time 0. To apply the result by Epstein and Schneider (2003) we assume the same condition. For the game $(N, H, W, l, (\mathcal{I}_i), (U_i))$ this implies that we only consider strategies with full support. This is without loss of generality. To allow for strategies without full support, we apply the construction to a suitable subset of plays. It is up to further research to determine whether this procedure has implications for the strategic analysis of extensive form Ellsberg games.

In the description of the information structure $\{\mathcal{W}_t\}_0^T$, we drop the information of *who*

takes the decision in some information set J_t at some time t . The filtration only represents what is known to all players at some time t , therefore it does not define a game, but only an event tree. When we define a rectangular Ellsberg strategy profile \mathcal{P} , we therefore only impose restrictions on the set of induced distributions over plays: as long as the set of induced realization probabilities $\Pi_{\mathcal{P}}$ (see (1)) is rectangular, any Ellsberg strategy profile \mathcal{P} which induces $\Pi_{\mathcal{P}}$ is admissible. We specify this in Definition 5. For the extension of Kuhn’s theorem to extensive form Ellsberg games this leaves some degree of freedom for the choice of an Ellsberg strategy profile \mathcal{P} given some set of realization probabilities $\Pi_{\mathcal{P}}$.

We have already found a way to describe Ellsberg strategy profiles, that is, “the big urn” over all pure strategies, in the setting of Epstein and Schneider (2003): they are represented by the outcomes $\Pi_{\mathcal{P}}$. Now we see how Ellsberg behavior strategies are described in this setting. Define the set of \mathcal{W}_t -conditionals of measures π on (W, \mathcal{W}_T) as

$$\Pi_{\mathcal{P}}^t(w)(\cdot) = \{ \pi^t(w)(\cdot) := \pi(\cdot | \mathcal{W}_t)(w) \mid \pi \in \Pi_{\mathcal{P}} \},$$

this is the set of Bayesian updates at a time t . The set of conditional one-step-ahead measures is defined by

$$\Pi_{\mathcal{P},+1}^t(w) = \{ \pi_{+1}^t(w) \mid \pi \in \Pi_{\mathcal{P}} \},$$

where π_{+1}^t is the restriction of π^t to \mathcal{W}_{t+1} . The sets $\Pi_{\mathcal{P}}^t$ and $\Pi_{\mathcal{P},+1}^t$ can be viewed as realizations of \mathcal{W}_t -measurable correspondences into $\Delta(W, \mathcal{W}_T)$ and $\Delta(W, \mathcal{W}_{t+1})$, respectively. The name “conditional one-step-ahead measures” has the following intuition. Each measure in $\Pi_{\mathcal{P},+1}^t(w)$ is a measure on \mathcal{W}_{t+1} , thus one can think of $\Pi_{\mathcal{P},+1}^t(w)$ as the set of measures describing beliefs about the “next step”: the belief at time t about what will happen at time $t + 1$. Think of it as cutting the tree at time t and time $t + 1$, then one is only left with distributions over the choices at the information sets at time t , and the distributions are induced by $\Pi_{\mathcal{P}}$. From this explanation it is clear that the conditional one-step-ahead measures are the induced Ellsberg behavior strategies. Epstein and Schneider (2003) explain further that from any set of one-step-ahead conditionals, e.g. some set Π_{+1}^t , a rectangular set Π can be constructed with backward construction via

$$\Pi = \{ \pi \in \Delta(W, \mathcal{W}_T) \mid \pi_{+1}^t \in \Pi_{+1}^t(w) \text{ for all } t \text{ and } w \}.$$

We see that Π is the set of all measures π whose one-step ahead conditionals conform with the Π_{+1}^t .

When is an Ellsberg strategy dynamically consistent (rectangular)? The Bayesian theory

says, a probability distribution π on (W, \mathcal{W}_T) can be for every t decomposed into its conditionals and marginals in the form

$$\pi^t(w) = \sum_{w' \in W} \pi^{t+1}(w') \cdot \pi_{+1}^t(w). \quad (4)$$

The set $\Pi_{\mathcal{P}}$ of probability distributions (and thus \mathcal{P}) on (W, \mathcal{W}_T) is rectangular if it admits a corresponding decomposition. Details are explained in the proof of Theorem 7 and in Epstein and Schneider (2003). First, we define rectangularity precisely.

Definition 5. *Let $(N, H, W, l, (\mathcal{I}_i), (U_i))$ be an extensive form Ellsberg game. An Ellsberg strategy profile \mathcal{P} is rectangular, if the set of realization probabilities $\Pi_{\mathcal{P}}$ is rectangular in the following sense. Let the information structure of the game be described by the filtration $\{\mathcal{W}_t\}_0^T$. Then $\Pi_{\mathcal{P}}$ is $\{\mathcal{W}_t\}$ -rectangular, if for all w and all t*

$$\Pi_{\mathcal{P}}^t(w) = \left\{ \sum_{w' \in W} \pi^{t+1}(w') \cdot \pi_{+1}^t(w) \mid \pi^{t+1}(w') \in \Pi_{\mathcal{P}}^{t+1}(w') \text{ for all } w' \in W, \pi_{+1}^t(w) \in \Pi_{\mathcal{P}, +1}^t(w) \right\}. \quad (5)$$

An Ellsberg strategy \mathcal{P}_i for player i is rectangular if it is part of a rectangular Ellsberg strategy profile \mathcal{P} . When necessary, we denote the restriction of a set X to its rectangular subset with X^R .

Remark 6. *Definition 5 also applies to the outcome Π_{Θ} . Π_{Θ} is rectangular, if it has an analog decomposition as in (5). In difference to Definition 5 this is not a property of the Ellsberg behavior strategy profile Θ .*

Observe that the inclusion \subset in Definition 5 is always satisfied by applying (4) for every $\pi \in \Pi_{\mathcal{P}}$. In order to assure the inclusion in the other direction, $\Pi_{\mathcal{P}}$ is required to have the special “rectangular” form. Combinations of conditionals and marginals that arise from different $\pi, \pi' \in \Pi_{\mathcal{P}}$ have to lie in $\Pi_{\mathcal{P}}^t$. In Example 9 in the following section, we explain the geometric representation of a rectangular Ellsberg strategy. Epstein and Schneider (2003) point out some features of rectangularity. One of these is particularly important in the setting of extensive form Ellsberg games: rectangularity imposes no restrictions on one-step-ahead conditionals. This means that players can use any Ellsberg behavior strategy they wish. Only when one wants to represent this Ellsberg behavior strategy by an outcome-equivalent Ellsberg strategy one has to use the induced rectangular Ellsberg strategy. We show this in detail in the following section.

4 Kuhn's Theorem for Extensive Form Ellsberg Games

Kuhn (1953) proves that an extensive form satisfies perfect recall if and only if, for every probability distribution on terminal histories that is induced by some mixed strategy profile, there is a behavior strategy profile that induces the same distribution, and for every probability distribution on terminal histories that is induced by some behavior strategy profile, there is a mixed strategy profile that induces the same distribution. Under expected utility this implies that the strategies yield the same utility. We show that under the assumption of perfect recall and the condition that the Ellsberg strategies are rectangular we get an equivalent result of Kuhn's theorem for extensive form Ellsberg games, namely:

Theorem 7 (Generalization of Kuhn's Theorem). *In an extensive form Ellsberg game $(N, H, W, l, (\mathcal{I}_i), (U_i))$ satisfying perfect recall, every rectangular Ellsberg strategy profile \mathcal{P} induces an Ellsberg behavior strategy profile $\Theta^{\mathcal{P}}$ via prior-by-prior updating; every Ellsberg behavior strategy profile Θ induces a rectangular Ellsberg strategy profile \mathcal{P}^{Θ} such that prior-by-prior updating of \mathcal{P}^{Θ} yields Θ . The induced strategy profiles are payoff-equivalent, i.e.,*

$$U_i(\mathcal{P}) = U_i(\Theta^{\mathcal{P}}) \quad \text{and} \quad U_i(\Theta) = U_i(\mathcal{P}^{\Theta}).$$

Proof. We construct the induced strategies as follows. Note that we only consider strategies with full support according to Axiom 5 in Epstein and Schneider (2003). This is without loss of generality. To allow for strategies without full support, we apply the construction to a suitable subset of plays.

A rectangular Ellsberg strategy profile \mathcal{P} has by definition a rectangular set of realization probabilities $\Pi_{\mathcal{P}}$. Due to rectangularity, $\Pi_{\mathcal{P}}$ possesses at any point in time t a decomposition into marginals and one-step-ahead conditionals according to (5). The set of conditional one-step-ahead measures at t , $\Pi_{\mathcal{P},+1}^t$, defines the induced Ellsberg behavior strategy profile $\Theta^{\mathcal{P}}$ by setting

$$\Theta_i^{\mathcal{P},\alpha} := \Pi_{\mathcal{P},+1}^t, \tag{6}$$

when player i is at information set α at time t . The construction is graphically captured in the diagram in Figure 1.

The construction of an induced rectangular Ellsberg strategy profile \mathcal{P}^{Θ} is similar. Any Ellsberg behavior strategy profile Θ has a set of realization probabilities Π_{Θ} . This set can be restricted to its rectangular subset Π_{Θ}^R by requiring that the conditional one-step-ahead

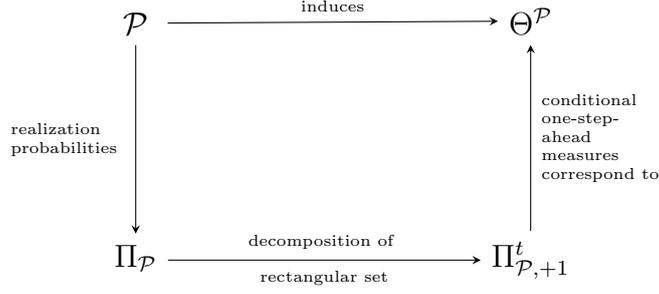


Figure 1: Rectangular \mathcal{P} induces Ellsberg behavior strategy $\Theta^{\mathcal{P}}$.

measures are again the Ellsberg behavior strategy profile Θ in the sense that

$$\Pi_{\Theta,+1}^{R,t} = \Theta_i^\alpha \quad (7)$$

when player i is at information set α at time t (dashed line in Figure 2). Then we define $\Pi_{\mathcal{P}}^R := \Pi_{\Theta}^R$. $\Pi_{\mathcal{P}}^R$ is uniquely determined by the process of conditional one-step-ahead correspondences $\Pi_{\Theta,+1}^{R,t}$. The rectangular set is the set of realization probabilities of the induced rectangular Ellsberg strategy profile \mathcal{P}^Θ . The construction is graphically captured in the diagram in Figure 2.

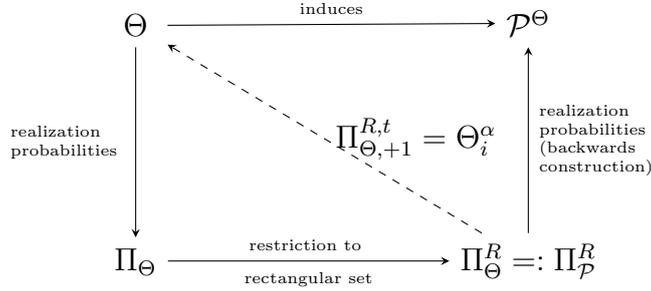


Figure 2: Θ induces rectangular Ellsberg strategy \mathcal{P}^Θ .

It remains to show that \mathcal{P} and $\Theta^{\mathcal{P}}$, as well as Θ and \mathcal{P}^Θ , yield the same payoff. Recall that

$$U_i(\mathcal{P}) = U_i(\Pi_{\mathcal{P}}) = \min_{P \in \mathcal{P}} \sum_{w \in W} \Pi_P(w) u_i(w),$$

where $\Pi_P(w) = \prod_{i \in N} \pi_i(w)$ with $\pi_i(w)$ the sum of all probabilities according to P_I of all the pure strategies of player i which are consistent with some terminal history $w \in W$.

Let $w = (a^1, \dots, a^K)$. Now consider the payoff of $\Theta^{\mathcal{P}}$.

$$\begin{aligned} U_i(\Theta^{\mathcal{P}}) &= \min_{\theta \in \Theta^{\mathcal{P}}} \sum_{w \in W} \Pi_{\theta}(w) u_i(w) \\ &= \min_{\theta \in \Theta^{\mathcal{P}}} \sum_{w \in W} \left(\prod_{k=0}^{K-1} \theta_{l(a^1, \dots, a^k)}(a^1, \dots, a^k)(a^{k+1}) \right) u_i(w) \end{aligned} \quad (8)$$

$$= \min_{\pi_{+1}^t \in \Pi_{\mathcal{P},+1}^t} \sum_{w \in W} \left(\prod_{t=0}^{T-1} \pi_{+1}^t(w) \right) u_i(w) \quad (9)$$

$$= \min_{P \in \mathcal{P}} \sum_{w \in W} \Pi_P(w) u_i(w). \quad (10)$$

From (8) to (9) we get by definition of θ_i in the construction of the proof, see (6). We can replace $\prod_{k=0}^{K-1} \theta_{l(a^1, \dots, a^k)}(a^1, \dots, a^k)(a^{k+1})$ by $\prod_{t=0}^{T-1} \pi_{+1}^t(w)$. The equality of (9) and (10) results from the rectangularity of \mathcal{P} . Because of rectangularity, $\prod_{t=0}^{T-1} \pi_{+1}^t(w)$ is an element of $\Pi_{\mathcal{P}}$, and $\prod_{t=0}^{T-1} \pi_{+1}^t(w) = \Pi_P(w)$. The equality of $U_i(\Theta)$ and $U_i(\mathcal{P}^{\Theta})$ is shown analogously. \square

Remark 8. *Note that the induced Ellsberg behavior strategy profile $\Theta^{\mathcal{P}}$ is unique, but the induced Ellsberg strategy profile \mathcal{P}^{Θ} is not. This is because the construction of \mathcal{P}^{Θ} from the set of realization probabilities $\Pi_{\mathcal{P}}^R$ is in general not unique.*

Theorem 7 says that for every Ellsberg behavior strategy profile there exists an Ellsberg strategy profile that yields the same utility for all players. This induced Ellsberg strategy profile has to be chosen rectangular if we want it to be dynamically consistent and equivalent to the original Ellsberg behavior strategy. Note that this does not impose any restriction on the choice of the Ellsberg behavior strategy profile (this is also pointed out in Epstein and Schneider (2003), p. 10), but only allows a restricted set of Ellsberg strategy profiles. Conversely, for every rectangular Ellsberg strategy, there exists an equivalent Ellsberg behavior strategy.

To gain some intuition about the nature of rectangular Ellsberg strategies and the inductions explained in the proof of Theorem 7, we now discuss a simple game.

Example 9. *Consider the following example taken from Osborne and Rubinstein (1994), p. 93. We believe this to be the simplest two-player information structure with which we can illustrate the equivalence of rectangular Ellsberg strategies and Ellsberg behavior strategies. The game is presented in Figure 3.*

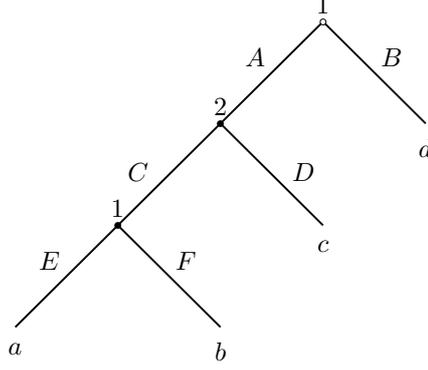


Figure 3: Game where player 1 moves before and after player 2.

We assume that $a, b, c, d \in \mathbb{R}$ are payoffs for player 1 with $d < a = b$. To keep things as simple as possible, we let player 2 choose C at all times, thus c will never be reached. Player 1 has four pure strategies: (A, E) , (A, F) , (B, E) and (B, F) , we denote the set of pure strategies of player 1 by S_1 . Player 2 only has C and D as pure strategies.

First, we start with Ellsberg behavior strategies $\Theta_1 = (\Theta_1^0, \Theta_1^1)$ and Θ_2 that specify a set of probability distributions at every information set of player 1 and 2, respectively. Suppose player 1 plays

$$\Theta_1^0 = \{(\theta_1^0, 1 - \theta_1^0) \mid \theta_1^0 \in [0, 1], 1/4 \leq \theta_1^0 \leq 3/4\},$$

$$\Theta_1^1 = \{(\theta_1^1, 1 - \theta_1^1) \mid \theta_1^1 \in [0, 1], 1/2 \leq \theta_1^1 \leq 2/3\},$$

with Θ_1^0 the set of distributions over A and B , and Θ_1^1 the set of distributions over E and F . For the purpose of the example, let player 2 have an Ellsberg behavior strategy Θ_2 that chooses C with probability 1, without creating any ambiguity.

The set of terminal histories W is then given by $W = \{E, F, B\}$ referring to the outcomes a, b and d . We derive the set of outcomes Π_Θ of the Ellsberg behavior strategy Θ ,

$$\Pi_\Theta(E) = \{\theta_1^0 \cdot \theta_1^1 \mid \theta_1^0 \in \Theta_1^0, \theta_1^1 \in \Theta_1^1\},$$

$$\Pi_\Theta(F) = \{\theta_1^0 \cdot (1 - \theta_1^1) \mid \theta_1^0 \in \Theta_1^0, \theta_1^1 \in \Theta_1^1\},$$

$$\Pi_\Theta(B) = \{1 - \theta_1^0 \mid \theta_1^0 \in \Theta_1^0\}.$$

Now we can calculate the maxmin expected utility for player 1 of the Ellsberg behavior

strategy Θ .

$$\begin{aligned}
U_1(\Theta) &= \min_{\theta \in \Theta} \sum_{w \in W} \Pi_{\theta}(w) u_i(w) \\
&= \min_{\theta \in \Theta} \Pi_{\theta}(E) \cdot a + \Pi_{\theta}(F) \cdot b + \Pi_{\theta}(B) \cdot d \\
&= \min_{\theta \in \Theta} \theta_1^0 \cdot \theta_1^1 \cdot a + \theta_1^0 \cdot (1 - \theta_1^1) \cdot b + (1 - \theta_1^0) \cdot d \\
&= 1/4 \cdot 1/2 \cdot a + 1/4 \cdot 1/2 \cdot b + 3/4 \cdot d \\
&= 1/8 \cdot a + 1/8 \cdot b + 3/4 \cdot d.
\end{aligned} \tag{11}$$

Now, what is the rectangular Ellsberg strategy profile \mathcal{P}^{Θ} induced by the Ellsberg behavior strategy profile Θ ? We follow the construction in Figure 2.

The set of induced realization probabilities Π_{Θ} is a hexagon (colored green in Figure 4), given by the surface that lies in the intersection of $\Pi_{\Theta}(E) = [1/8, 1/2]$, $\Pi_{\Theta}(F) = [1/12, 3/8]$ and $\Pi_{\Theta}(B) = [1/4, 3/4]$. It can be easily seen that Π_{Θ} is too large in the sense that it induces conditional probabilities over E and F which lie outside the set Θ_1^1 of original Ellsberg behavior strategies. Take for example the distribution $(1/8, 3/8, 1/2)$ (probability distribution over E, F, B) at the lower right corner of the green hexagon. This distribution yields a conditional one-step-ahead probability for E which is equal to $\frac{1/8}{1/2} = 1/4$, and $1/4$ does not lie in the set $\Theta_1^1(E)$. Hence we have to restrict Π_{Θ} to its rectangular subset Π_{Θ}^R , this is exactly the set that yields the correct conditionals.

To apply the construction of Epstein and Schneider (2003) we represent the game in Figure 9 by a state space $W = \{E, F, B\}$ (the set of plays) and a filtration $\{\mathcal{W}_t\}_0^2$ where

$$\begin{aligned}
\mathcal{W}_0 &= \{\{E, F, B\}\}, \\
\mathcal{W}_1 &= \{\{E, F\}, \{B\}\}, \\
\mathcal{W}_2 &= \{\{E\}, \{F\}, \{B\}\}.
\end{aligned}$$

$\Pi_{\Theta}^R \subset \Pi_{\Theta}$ is the set of distributions on (W, \mathcal{W}_2) that for all $t \in \{0, 1, 2\}$ and all $w \in W$ admits a decomposition

$$\begin{aligned}
\Pi_{\Theta}^{R,t}(w) &= \left\{ \sum_{w' \in W} \pi^{t+1}(w') \cdot \pi_{+1}^t(w) \mid \right. \\
&\quad \left. \pi^{t+1}(w') \in \Pi_{\Theta}^{t+1}(w') \text{ for all } w' \in W, \pi_{+1}^t(w) \in \Pi_{\Theta,+1}^t(w) \right\}.
\end{aligned} \tag{12}$$

The set Π_{Θ}^R excludes exactly all those distributions in Π_{Θ} that yield conditionals outside Θ_1^1 . The rectangular set Π_{Θ}^R is depicted as the blue rectangle in Figure 4.

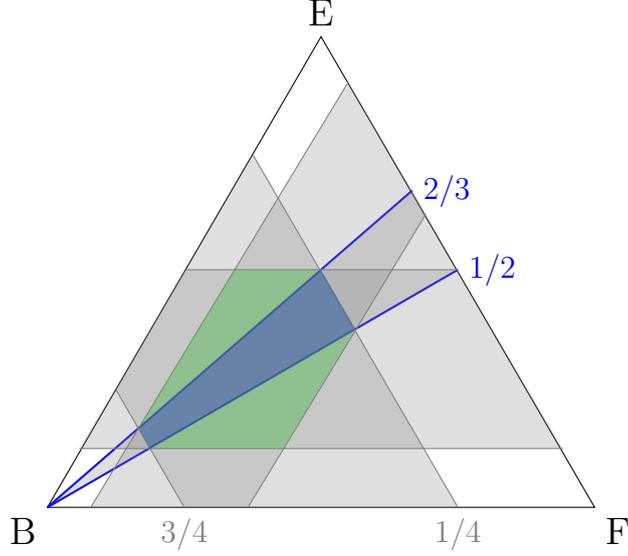


Figure 4: Set of realization probabilities Π_{Θ} (green hexagon), rectangular set of probabilities Π_{Θ}^R (blue rectangle). The grey numbers indicate the set of probabilities for B , the blue numbers the set of conditional one-step-ahead probabilities for E .

From the construction of $\Pi_{\Theta}^{R,t}$, it follows that

$$\begin{aligned}\Pi_{\Theta,+1}^{R,0} &= \Theta_1^0, \\ \Pi_{\Theta,+1}^{R,1} &= \Theta_1^1,\end{aligned}$$

that is, the induced conditional one-step-ahead measures are exactly the components of the Ellsberg behavior strategy which we started with. The set Π_{Θ}^R can also be constructed recursively as the set of all distributions $\pi \in \Delta(W, \mathcal{W}_T)$ for which $\pi_{+1}^t \in \Pi_{\Theta,+1}^{R,t}(w)$ for all t and all w , see Epstein and Schneider (2003) p. 8 for details. Hence, to find the induced Ellsberg strategy \mathcal{P}^{Θ} we first define

$$\Pi_{\mathcal{P}}^R := \Pi_{\Theta}^R$$

as in Figure 2. Then the rectangular set of realization probabilities $\Pi_{\mathcal{P}}^R$ is

$$\begin{aligned}\Pi_{\mathcal{P}}^R &= \left\{ \pi = (\pi_1, \pi_2, \pi_3) \in \Delta(W, \mathcal{W}_T) \mid \pi_1/(1 - \pi_3) \in [1/2, 2/3] = \Pi_{\mathcal{P},+1}^{R,1}(E), \right. \\ &\quad \left. \pi_2/(1 - \pi_3) \in [1/3, 1/2] = \Pi_{\mathcal{P},+1}^{R,1}(F), \pi_3 \in [1/4, 3/4] = \Pi_{\mathcal{P},+1}^{R,0}(B) \right\} \\ &= \left\{ \pi \in \Delta(W, \mathcal{W}_T) \mid \pi_3 \in [1/4, 3/4], \pi_1 \in [-\pi_3/2 + 1/2, -2\pi_3/3 + 2/3] \right\}.\end{aligned}\quad (13)$$

The notation in (13) is derived as follows. Every point in the simplex is fully described by

probabilities for B and E . The two lines passing through $B = 1$ and $E = 2/3$, $B = 1$ and $E = 1/2$ are functions of B and E with equations

$$\begin{aligned}\pi_1 &= \pi_3/2 + 1/2, \\ \pi_1 &= 2\pi_3/3 + 2/3.\end{aligned}$$

Now, finally, we construct \mathcal{P}^Θ from the rectangular set of realization probabilities $\Pi_{\mathcal{P}}^R$. We find

$$\mathcal{P}_1^\Theta = \left\{ P \in \Delta S_1 \mid P(A, E) = \pi_1, P(A, F) = \pi_2, P(B, E) = \frac{\pi_1}{\pi_1 + \pi_2} \pi_3, \right. \\ \left. P(B, F) = \frac{\pi_2}{\pi_1 + \pi_2} \pi_3, \text{ for all } (\pi_1, \pi_2, \pi_3) \in \Pi_{\mathcal{P}}^R \right\}.$$

We calculate the maxmin expected utility of player 1 of the Ellsberg strategy profile \mathcal{P}^Θ .

$$\begin{aligned}U_1(\mathcal{P}^\Theta) &= \min_{P \in \mathcal{P}} \sum_{w \in W} \Pi_P(w) u_i(w) \\ &= \min_{P \in \mathcal{P}} P(A, E) \cdot a + P(A, F) \cdot b + (P(B, E) + P(B, F)) \cdot d \\ &= 1/8 \cdot a + 1/8 \cdot b + (3/8 + 3/8) \cdot d \\ &= 1/8 \cdot a + 1/8 \cdot b + 3/4 \cdot d.\end{aligned}\tag{14}$$

Player 1 uses $P = (1/8, 1/8, 3/8, 3/8)$ as his worst case probability distribution, because it puts the greatest available probability to the worst outcome d . We thus find that the Ellsberg strategy profile \mathcal{P}^Θ induced by Θ yields the same maxmin expected utility as Θ :

$$U_1(\mathcal{P}^\Theta) \stackrel{(14)}{=} 1/8 \cdot a + 1/8 \cdot b + 3/4 \cdot d \stackrel{(11)}{=} U_1(\Theta).$$

Furthermore, doing the reasoning backwards, one sees directly that with the rectangular construction of $\Pi_{\mathcal{P}}^R = \Pi_{\Theta}^R$ in (12), the Ellsberg strategy yields the correct one-step-ahead conditionals for A, B and E, F .

Now, we start out from the rectangular Ellsberg strategy profile \mathcal{P} (the blue rectangle in Figure 4) and construct the induced Ellsberg behavior strategy profile $\Theta^{\mathcal{P}}$. We proceed as sketched in the diagram in Figure 1. The Ellsberg strategy profile has a set of realization probabilities $\Pi_{\mathcal{P}}$ which are by definition also rectangular. Then from Definition 5, $\Pi_{\mathcal{P}}$ possesses a decomposition of $\Pi_{\mathcal{P}}^t$ for all t into marginals and one-step-ahead conditionals. The rectangular set $\Pi_{\mathcal{P}}$ is constructed in the way to get exactly $\Pi_{\mathcal{P},+1}^0 = \Theta_1^{\mathcal{P},0}$ and $\Pi_{\mathcal{P},+1}^1 =$

$\Theta_1^{\mathcal{P},1}$ as one-step-ahead conditionals, and hence

$$\begin{aligned}\Theta_1^{\mathcal{P},0} &= \{(\theta_1^0, 1 - \theta_1^0) \mid \theta_1^0 \in [0, 1], 1/4 \leq \theta_1^0 \leq 3/4\}, \\ \Theta_1^{\mathcal{P},1} &= \{(\theta_1^1, 1 - \theta_1^1) \mid \theta_1^1 \in [0, 1], 1/2 \leq \theta_1^1 \leq 2/3\}.\end{aligned}$$

Obviously we have

$$U_1(\mathcal{P}) = U_1(\mathcal{P}^\Theta) \stackrel{(14)}{=} 1/8 \cdot a + 1/8 \cdot b + 3/4 \cdot d \stackrel{(11)}{=} U_1(\Theta) = U_1(\Theta^\mathcal{P}).$$

5 Ellsberg Equilibrium in Extensive Form Ellsberg Games

An Ellsberg equilibrium in an extensive form Ellsberg game is defined straightforwardly as in the static case: a profile of Ellsberg strategies is an Ellsberg equilibrium profile, if no player finds it profitable to deviate unilaterally.

Definition 10. *Let $(N, H, W, l, (\mathcal{I}_i), (U_i))$ be an extensive form Ellsberg game. An Ellsberg equilibrium is a profile $(\mathcal{P}_1^*, \dots, \mathcal{P}_n^*)$ of Ellsberg strategies such that for all $i = 1, \dots, n$ and every $\mathcal{P}_i \subseteq \Delta S_i$*

$$U_i(\mathcal{P}_i^*, \mathcal{P}_{-i}^*) \geq U_i(\mathcal{P}_i, \mathcal{P}_{-i}^*).$$

An *Ellsberg equilibrium in Ellsberg behavior strategies* is defined analogously. Given Theorem 7, the two definitions are equivalent for extensive form Ellsberg games with perfect recall.

Corollary 11 (of Theorem 7). *If a rectangular Ellsberg strategy profile \mathcal{P} is an Ellsberg equilibrium of an extensive form Ellsberg game, then the Ellsberg behavior strategy profile $\Theta^\mathcal{P}$ is an Ellsberg equilibrium in Ellsberg behavior strategies. If an Ellsberg behavior strategy profile Θ is an Ellsberg equilibrium in Ellsberg behavior strategies, then the rectangular Ellsberg strategy profile $\Theta^\mathcal{P}$ is an Ellsberg equilibrium.*

Every Nash equilibrium is an Ellsberg equilibrium. Conversely, every Ellsberg equilibrium in which \mathcal{P} is a single probability distribution is a Nash equilibrium.

Note that Corollary 11 now allows to adapt notions such as subgame perfect equilibrium and sequential equilibrium to extensive form Ellsberg games. However, in this paper we do not aim at further characterizing Ellsberg equilibria and the strategic implications of Theorem 7 for extensive form Ellsberg games.

6 Relation to Other Extensive Form Solution Concepts

Concepts

We already discussed the peace negotiation example by Greenberg (2000) in Riedel and Sass (2013): it is an example of an extensive form game in which one Ellsberg equilibrium (of the corresponding normal form) lies outside the set of Nash equilibria of the game. We will now analyze another example which has a similar structure. The example (see Figure 5), which Fudenberg and Kreps (1988) provided to show that mistakes about play off the equilibrium path can lead to non-Nash outcomes, is used to discuss the relation of Ellsberg equilibrium to other extensive form solution concepts.

Following Fudenberg and Kreps (1988), the “horse”-like game² has been used as an example in a number of papers (Battigalli, Gilli, and Molinari (1992), Fudenberg and Levine (1993), Rubinstein and Wolinsky (1994), Groes, Jacobsen, Sloth, and Tranaes (1998), Lo (1999)). The information structure is the same as in Greenberg (2000)’s peace negotiation example, and also the incentive structure is very similar. We are in a situation, where ambiguity can be used as a threat against deviation.³

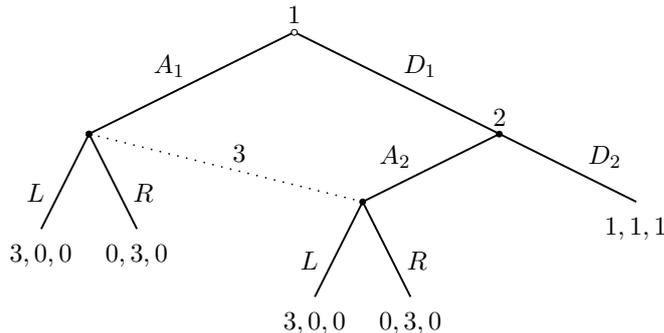


Figure 5: Fudenberg and Kreps’ three-player game.

Proposition 12. *The proper Ellsberg equilibria of Fudenberg and Kreps’ Three-Player game are of the form $(A_1, A_2, [R_0, R_1])$ where $R \in [R_0, R_1]$ is the probability that player 3 plays L, and $R_0 < 1/3$ and $R_1 > 2/3$.*

Proof. Let P, Q denote the probability that player 1 plays A_1 , player 2 plays A_2 , respec-

² Selten (1975) is first to analyze this “horse”-like information structure. In his numerical example, however, the possibility to use Ellsberg strategies does not lead to Ellsberg equilibria outside the Nash equilibrium support, because the incentive structure is such that players 1 and 2 are not in opposition to each other.

³ Kelsey and Spanjers (2004) observe that “Ambiguity can make threats more effective.” In the incentive structure of the games considered here, this observation is confirmed.

tively. Then we can calculate the minimal expected utility for each player.

$$\begin{aligned}
U_1(P, Q, [R_0, R_1]) &= \min_{R_0 \leq R \leq R_1} R(-3PQ + 3) + PQ \\
&= R_0(-3PQ + 3) + PQ \text{ for all } P, Q \in [0, 1]. \\
U_2(P, Q, [R_0, R_1]) &= \min_{R_0 \leq R \leq R_1} R(3PQ - 3) - 2PQ + 3 \\
&= R_1(3PQ - 3) - 2PQ + 3 \text{ for all } P, Q \in [0, 1]. \\
U_3(P, Q, [R_0, R_1]) &= \min_{R_0 \leq R \leq R_1} PQR + PQ(1 - R) = PQ.
\end{aligned}$$

Player 1 gets a payoff less than 1 if player 1 or 2 defect, when $R_0(-3PQ + 3) + PQ < 1 \Leftrightarrow R_0 < 1/3$. Player 2 gets a payoff less than 1 if player 1 or 2 defect, when $R_1(3PQ - 3) - 2PQ + 3 < 1 \Leftrightarrow R_1 > 2/3$. This yields the equilibrium $(A_1, A_2, [R_0, R_1])$ where $R_0 < 1/3$ and $R_1 > 2/3$. \square

The information structure of the game has been characterized by Fudenberg and Levine (1993) in the analysis of their concept of Self-confirming equilibrium. They define a property of *unobserved deviators* which captures the fact that a player does not observe who of his opponents leaves the equilibrium path at some point in the game. These games have a natural property to allow for “wrong” beliefs on plays *off* the equilibrium path. Therefore games with unobserved deviators can have Self-confirming or Conjectural (Battigalli (1987), Battigalli and Guaitoli (1988), Battigalli and Guaitoli (1997)) equilibria which are not Nash equilibria. Fudenberg and Levine (1993) show that games that have only *observed* deviators do not have non-Nash equilibria. We suppose that a similar characterization holds for Ellsberg equilibria as well. This is an interesting question for further research.

We briefly present other extensive form solution concepts with Knightian uncertainty. Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2012) define Self-confirming equilibrium with model uncertainty. Their concept also incorporates Knightian uncertainty into extensive form games, but is quite different from our extensive form Ellsberg games. The differences lie especially in the fact that players may only play pure strategies in Self-confirming equilibrium with model uncertainty, and that uncertainty is present in the environment and not in the strategies. In Aryal and Stauber (2013) players, anticipating possible small mistakes (trembles), can have ambiguous beliefs about their opponents’ strategies. A further extensive form solution concept with Knightian uncertainty is Multiple Priors Nash equilibrium by Lo (1999). Lo allows beliefs to be represented by multiple priors and demands every distribution in the sets of beliefs to be a best response to the beliefs on the other players. Players only have pure strategies at their disposition. Lo

(1999) also looks at the example in Figure 5 and shows that the path (D_1, D_2) can be supported by a Multiple Priors Nash equilibrium, that is, that there exist sets of beliefs such that playing (D_1, D_2) is optimal. However, this equilibrium in beliefs is entirely subjective and in that aspect differs fundamentally from Ellsberg equilibrium. In the same way, Nash equilibrium with lower probabilities in Groes, Jacobsen, Sloth, and Tranaes (1998) is distinguished from Ellsberg equilibrium. The Ellsberg equilibrium outcomes in Fudenberg and Kreps' game can also be achieved with subjective equilibrium by Kalai and Lehrer (1995), but, again, the fundamentals of the two concepts differ.

σ^* - equilibrium by Ma (2000) is closest to our approach. In his model, players use ambiguous plans modeled by sets of probability distributions over acts. The author also mentions the possibility to *create* Knightian uncertainty. In difference to our approach he relies entirely on subjective preference representations and focuses on belief systems and thus on equilibrium in beliefs. An interesting aspect of his paper is that it models the possibility to create ambiguity with pre-play communication.

7 Conclusion

We define extensive form Ellsberg games in which players can use imprecise probabilistic devices; either over their pure strategies, or over the available actions at every decision point. This is a generalization of classical mixed and behavioral strategies in extensive form games. We show that when Ellsberg strategies (the generalized mixed strategies) are rectangular in the sense of Epstein and Schneider (2003), they are equivalent to Ellsberg behavior strategies (the generalized behavioral strategies). We thus prove a generalization of Kuhn (1953)'s theorem for extensive form Ellsberg games.

When we apply the representation result by Epstein and Schneider (2003) we temporarily ignore the strategic element of the extensive form game. The game is represented as an event tree: this implies that which player decides when does not play a role. Thereby we establish, that it is equivalent to calculate Ellsberg equilibrium in Ellsberg strategies or Ellsberg behavior strategies. In classical extensive form theory this equivalence proved crucial for equilibrium analysis. Naturally, the behavioral and strategic implications of this equivalence in extensive form Ellsberg games are topic for further research.

References

ARYAL, G., AND R. STAUBER (2013): "Trembles in extensive games with ambiguity averse players," <http://cbe.anu.edu.au/media/2664963/wp606.pdf>, Working paper No. 606, Australian National University.

- BADE (2011a): “Independent randomization devices and the elicitation of ambiguity-averse preferences,” <http://www.coll.mpg.de/Download/Bade/PreferenceElicitation10.pdf>, Preprint MPI Research on Collective Goods.
- BADE, S. (2011b): “Ambiguous act equilibria,” *Games and Economic Behavior*, 71(2), 246–260.
- BATTIGALLI, P. (1987): “Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali,” *Unpublished undergraduate dissertation, Università Bocconi*.
- BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2012): “Selfconfirming equilibrium and model uncertainty,” *Working Paper n. 428, Università Bocconi*.
- BATTIGALLI, P., M. GILLI, AND M. MOLINARI (1992): “Learning and convergence to equilibrium in repeated strategic interactions: an introductory survey,” *Ricerche Economiche*, 46, 335–378.
- BATTIGALLI, P., AND D. GUAITOLI (1988): “Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information,” *Extended Abstract, Università Bocconi*.
- (1997): “Conjectural equilibria and rationalizability in a game with incomplete information,” *Studies in Risk and Uncertainty*, 8, 97–124.
- EPSTEIN, L., AND M. SCHNEIDER (2003): “Recursive multiple-priors,” *Journal of Economic Theory*, 113(1), 1–31.
- FUDENBERG, D., AND D. KREPS (1988): “Learning, experimentation, and equilibrium in games,” *mimeo, Department of Economics, Stanford University*.
- FUDENBERG, D., AND D. LEVINE (1993): “Self-confirming equilibrium,” *Econometrica*, 61(3), 523–545.
- GAJDOS, T., T. HAYASHI, J. TALLON, AND J. VERGNAUD (2008): “Attitude toward imprecise information,” *Journal of Economic Theory*, 140(1), 27–65.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- (1993): “Updating ambiguous beliefs,” *Journal of Economic Theory*, 59(1), 33–49.

- GREENBERG, J. (2000): “The right to remain silent,” *Theory and Decision*, 48(2), 193–204.
- GROES, E., H. JACOBSEN, B. SLOTH, AND T. TRANAES (1998): “Nash equilibrium with lower probabilities,” *Theory and Decision*, 44, 37–66.
- HANANY, E., AND P. KLIBANOFF (2007): “Updating preferences with multiple priors,” *Theoretical Economics*, 2, 261–298.
- KALAI, E., AND E. LEHRER (1995): “Subjective games and equilibria,” *Games and Economic Behavior*, 8(1), 123–163.
- KELSEY, D., AND W. SPANJERS (2004): “Ambiguity in partnerships,” *The Economic Journal*, 114, 528–546.
- KLIBANOFF, P. (2001): “Stochastically independent randomization and uncertainty aversion,” *Economic Theory*, 18(3), 605–620.
- KUHN, H. (1953): “Extensive games and the problem of information,” *Contributions to the Theory of Games*, 2(28), 193–216.
- LO, K. (1999): “Extensive form games with uncertainty averse players,” *Games and Economic Behavior*, 28(2), 256–270.
- MA, C. (2000): “Uncertainty aversion and rationality in games of perfect information,” *Journal of Economic Dynamics and Control*, 24(3), 451–482.
- OSBORNE, M., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT press.
- RIEDEL, F. (2004): “Dynamic coherent risk measures,” *Stochastic processes and their applications*, 112(2), 185–200.
- RIEDEL, F., AND L. SASS (2013): “Ellsberg games,” *Theory and Decision*, forthcoming.
- RUBINSTEIN, A., AND A. WOLINSKY (1994): “Rationalizable conjectural equilibrium: between Nash and rationalizability,” *Games and Economic Behavior*, 6(2), 299–311.
- SELTEN, R. (1975): “Reexamination of the perfectness concept for equilibrium points in extensive games,” *International Journal of Game Theory*, 4(1), 25–55.