The Strategic Use of Ambiguity^{*}

Frank Riedel^{†1} and Linda Sass^{‡2}

¹Bielefeld University and Princeton University ²Bielefeld University and Université Paris 1 Panthéon–Sorbonne

First Version: August 2, 2011, This Version: May 28, 2012

^{*}We thank Andreas Blume, Stephen Morris, Daniela Puzzello, Ariel Rubinstein, Burkhard Schipper, and Marco Scarsini for discussions about the decision-theoretic foundations of game theory. We thank Jürgen Eichberger, Joseph Greenberg, Itzhak Gilboa, Kin Chung Lo, and Bill Sandholm for valuable comments on the exposition of the paper. The paper was presented at Indiana University, Minnesota University, Institut Henri Poincaré Paris, University of Exeter, and University of Texas at Austin as well as at Games Toulouse 2011, Mathematical Aspects of Game Theory and Applications. We thank seminar participants for their unambiguous attention.

[†]Financial Support through the German Research Foundation, International Graduate College "Stochastics and Real World Models", Research Training Group EBIM, "Economic Behavior and Interaction Models", and Grant Ri–1128–4–1 is gratefully acknowledged. Parts of the paper were written during a sabbatical at Princeton University. I thank Patrick Cheridito and the ORFE department for their hospitality.

[‡]Financial support through the German Research Foundation, International Research Training Group EBIM "Economic Behavior and Interaction Models", and through the DFH-UFA, French-German University, is gratefully acknowledged. I thank University Paris 1 and Jean-Marc Tallon for hospitality during my research visit 2011–2012.

Abstract

We propose a framework for normal form games where players can use Knightian uncertainty strategically. In such *Ellsberg games*, ambiguity– averse players may render their actions objectively ambiguous by using devices such as Ellsberg urns, in addition to the standard mixed strategies. While Nash equilibria remain equilibria in the extended game, there arise new *Ellsberg equilibria* with distinct outcomes, as we illustrate by negotiation games with three players. We characterize Ellsberg equilibria in two-person games with conflicting interests. These equilibria turn out to be consistent with experimental deviations from Nash equilibrium play.

Key words and phrases: Knightian Uncertainty in Games, Strategic Ambiguity, Ellsberg Games

JEL subject classification: C72, D81

1 Introduction

Common sense suggests that a certain strategic ambiguity can be useful in conflicts. "Many different strategies are used to orient toward conflicting interactional goals; some examples include avoiding interaction altogether, remaining silent, or changing the topic." says Eric Eisenberg in his famous article (Eisenberg (1984)), and he points out that applying one's resources of ambiguity is key in successful communication when conflicts of interest are present.

This paper introduces such strategic use of ambiguity into games. Although game theory was invented to model conflicts of interest, so far the theory does not allow players to intentionally choose ambiguity as a strategy. To this end we are going back to the beginnings of game theory. Von Neumann and Morgenstern (1953) introduced mixed strategies as random devices that are used to conceal one's behavior. We take up this interpretation and propose a generalization. In short, we allow players to use Ellsberg urns in addition to probabilistic devices like a roulette wheel or a die. For example, a player can base his action on the draw from an urn that contains hundred red and blue balls and it is only known that the number of red balls is between thirty and fifty. Such urns are objectively ambiguous, by design; players can thus create ambiguity. The recent advances in decision theory that were motivated by Ellsberg's famous experiments and Knight's distinction between risk and uncertainty, allow to model such strategic behavior formally.

Objective ambiguity means that players cannot have a look into their own urn. They intentionally choose a device that leaves themselves uncertain about the pure strategy eventually played once uncertainty is resolved. Players choose to know less about their own play than they could know. This may sound surprising or implausible at first. It is, however, the natural generalization of the usual mixed strategy when it is interpreted as an objective device to conceal one's behavior, as in the classical justification by von Neumann and Morgenstern. The founders of game theory justify the use of random devices in zero sum games by a thought experiment. If your opponent might find out your strategy, then it is optimal to conceal your behavior by using a random device. We just go one step further and allow players to use Ellsberg urns to conceal their plans.

Another interpretation of ambiguous actions in game theory is closer to the behavioral and psychological literature. Gigerenzer (2007), e.g., claims that not assessing all possibilities and information about a choice is often better, more efficient or more satisfying for a human decision maker. Using an Ellsberg strategy where the exact probability of choosing an action is not specified might be viewed as one way to model such mental efficiency.

In this paper we propose a model of *Ellsberg games* that differs from the classical model merely in the one aspect that players may use (objectively ambiguous) Ellsberg urns in addition to mixed strategies. We first discuss the conceptual foundations of our approach; to this end, we compare the two interpretations of mixed strategies as objective random devices (von Neumann and Morgenstern) or as beliefs about other players' pure actions (the Bayesian view). The latter, Bayesian view has been generalized to ambiguity before in the literature starting with Dow and Werlang (1994), Lo (1996), and Marinacci (2000). There, the interpretation of mixed strategies as beliefs about others' actions is generalized to uncertain beliefs in the sense of the decisiontheoretic literature on Knightian uncertainty (Schmeidler (1989), Gilboa and Schmeidler (1989)). Bade (2011) allows for ambiguous Anscombe–Aumann acts, and is closer to our approach in that sense; however, she assumes that ambiguity is a subjective part of players' preferences, thus generalizing Aumann's subjective equilibria (Aumann (1974)). A detailed discussion of the literature is provided in Section 6. Our approach takes the von Neumann and Morgenstern point of view where mixed strategies are used to conceal

one's behavior. We generalize this classical model to incorporate Knightian uncertainty in an objective sense as it was recently axiomatized by Gajdos, Hayashi, Tallon, and Vergnaud (2008).

A remarkable consequence is that players, once offered the possibility to use Ellsberg urns, actually use them, even though they are ambiguity– averse¹. New equilibria emerge that are not Nash equilibria in the original game, with outcomes that are not in the support of the original Nash equilibrium. We explain this with an example of a peace negotiation taken from Greenberg (2000). This game has a unique Nash equilibrium in which war is the outcome; there is another, as we call it, Ellsberg equilibrium, in which peace is the outcome.

Our approach to games has its most natural and fruitful applications to conflicts where players are at least to some degree in opposition to each other. We consequently perform a detailed study of two-person 2×2 games with conflicting interests, as Matching Pennies, or similar competitive situations.

We discuss first two new phenomena, *immunization against ambiguity* and *nonlinearity of payoffs* that arise in Ellsberg games. We then derive all Ellsberg equilibria of 2×2 -conflict games. While our predictions are broader than the classical unique Nash equilibrium, they remain restrictive, and, at least in principle, testable. Our results do allow to explain the experimental findings of Goeree and Holt (2001) who show that humans tend to deviate from Matching Pennies in asymmetric modified Matching Pennies games, but tend to play Nash equilibrium in symmetric Matching Pennies. This corresponds and is consistent with our Ellsberg equilibria.

The paper is organized as follows. In Section 2 we explain how Ellsberg urns are understood as concealment device, in the line of von Neumann and Morgenstern's interpretation of mixed strategies. In Section 3 we develop the theoretical framework of Ellsberg games. The concept is applied to the negotiation example in Section 4. Section 5 analyses the use of strategic ambiguity in two-person conflicts. We compare Ellsberg games to existing equilibrium concepts with ambiguity-aversion in Section 6, and we conclude in Section 7.

¹In this sense, our assumption of ambiguity–aversion is parsimonious as it makes it harder for players to introduce Knightian uncertainty. They do not use Ellsberg urns for love of uncertainty.

2 Ellsberg Urns and Mixed Strategies: Concealment Device versus Beliefs about Opponents' Behavior

Let us go back to the very foundations of game theory. A game consists of a finite set N of players, a finite set of (pure) strategies $S_i, i \in N$ for each player, as well as a collection of payoff functions $u_i : S \to \mathbb{R}$ defined over strategy profiles $S = \times_{i \in N} S_i$. Von Neumann and Morgenstern (1953) introduce mixed strategies as probability vectors P_i over pure strategies S_i . The question then emerges how players evaluate profiles of such mixed strategies $P = (P_1, \ldots, P_n)$; as the reader knows, von Neumann and Morgenstern adopt expected utility (and axiomatize their choice).

In this paper, we are going back to these foundations and propose a generalization. We allow players to use Ellsberg urns in addition to probabilistic devices like a roulette wheel or a die. So we imagine that a player can credibly commit his behavior on the outcome of an Ellsberg urn whose parameters he has chosen. To give an example, in a Matching Pennies game he would play HEAD if the draw from an urn with 100 red and blue balls yields a red ball; he himself and the other players only know that the proportion of red balls lies between 30 and 50 percent. The Ellsberg urn thus displays objective, common knowledge of ambiguity. All players know the possible probability distributions of outcomes, but no player has an informational advantage over others. We want to find out what the consequences for game theory are if we change the foundations in such a way.

Before we justify conceptually our new approach, let us go back again to classical game theory and ask how mixed strategies are justified and interpreted there. Our discussion follows closely the excellent account delivered in Reny and Robson (2004).

From a mathematical point of view, mixed strategies lead to convex strategy sets; and if one wants to assign a unique value to zero sum games, e.g., such convexity is needed. Convexity and linearity of the payoff functions are useful in many other respects as well, of course. Just think about the indifference principle by which we usually find Nash equilibria.

The purely mathematical aspect would not be very compelling, of course, had it not a plausible interpretation. In the words of Von Neumann and Morgenstern (1953) (p. 144): "In playing Matching Pennies against an at least moderately intelligent opponent, the player will not attempt to find out the opponent's intentions, but will concentrate on avoiding having his own intentions found out, by playing irregularly 'heads' and 'tails' in successive games. Since we wish to describe the strategy in one play — indeed we must discuss the course in one play and not that of a sequence of plays — it is preferable to express this as follows: The player's strategy consists neither of playing 'tails' and 'heads', but of playing 'tails' with the probability of $\frac{1}{2}$ and 'heads' with the probability of $\frac{1}{2}$." They then point out that this strategy protects the player against losses as his expected gain is always zero regardless how the opponent plays.

Von Neumann and Morgenstern always interpret these strategies as *objective* random devices like a fair coin or die. In particular, all players assign the same probabilities to the device's outcomes. In zero sum games, the use of such mixing can be justified as an attempt to conceal your behavior from your opponents. Indeed, von Neumann and Morgenstern also offer a Stackelberg game–like argument. Suppose that you have to write down your strategy on a sheet of paper before you play. If your opponent sends a spy able to find out what you have written down, then, in a zero sum game, it is strictly better for you to have concealed your behavior by writing "I will use a fair coin to determine my behavior." A mixed strategy, in von Neumann and Morgenstern's interpretation, is thus deliberate, objective randomization.

These arguments run into problems in common interest games as one would prefer one's own strategy to be found out by the opponent in simple coordination games (as has been pointed out by Schelling (1960) and Lewis (1969) already). Harsanyi (1967)'s construction allows to resolve this plausibility problem; he shows that mixed strategy equilibria can be interpreted as pure strategy equilibria in nearby incomplete information games where the payoffs are suitably perturbed, and players have private information. A common interpretation of mixed strategies nowadays goes even a step further. Several authors, including Aumann (1987), Armbruster and Boege (1979), Tan and Werlang (1988), Aumann and Brandenburger (1995), propose to forgo Harsanyis's construction and to interpret mixed strategies directly as the *belief* about which pure strategies the other players are going to use. Players are assumed to choose a definite action; as other players do not know exactly which one, the mixed strategy represents their uncertainty. For more than two players, this requires some consistency among beliefs in equilibrium, of course. We will come back to this below when we discuss beliefs equilibria.

Summing up, we have here two opposing interpretations of mixed strategies: on the one hand, the "objective" interpretation by von Neumann and Morgenstern where players deliberately use random devices with known probabilities to conceal their behavior; on the other hand, the "subjective" beliefs interpretation where the probability distributions represent players' uncertainty about other players' pure strategy choice.²

The literature on ambiguity in games has mainly focused on the beliefs interpretation of mixed strategies. So far, it has usually been assumed that the players choose pure strategies and that the opponents are uncertain in the Knightian sense about their choice. Please refer to Section 6 for a detailed discussion of the literature and the relation of the different approaches to our model.

Our paper differs from that literature as we take up the approach by von Neumann and Morgenstern and allow players to use objective devices that create Knightian uncertainty. We let them play Ellsberg urns to conceal their behavior. As in the classical case, our approach will have the most fruitful applications in games of conflict where it is in players own interest to conceal their behavior. This will also be highlighted by our examples. The approach might seem less plausible in common interest games; however, the theory in its abstract form applies there as well, of course, and we think it is useful to study the consequences in such games, too.

$$(1-t_i)u_i(m_i, m_{-i}) + t_iu_i(m_i, \text{bestreplyto}(m_i)).$$

²Both approaches are merged in the framework of Reny and Robson (2004). Reny and Robson unify both views of mixed equilibria with the help of another construction. In their perturbed game, every player *i* is characterized by a privately known subjective probability $t_i \in [a, b]$, 0 < a < b < 1 according to which he believes that a spy finds out his strategy. The payoff of the perturbed game is then

So with probability $1 - t_i$ the normal static payoff is obtained, whereas with probability t_i , the other player finds out one's strategy and is allowed to play a best reply to it. In such a situation, for a generic class of games, one can approximate all Nash equilibria by suitable pure strategy equilibria of the perturbed game, with a and b close to zero. In contrast to Harsanyi, players do use mixed strategies in the perturbed game sometimes, e.g. in zero sum games. We refer to their excellent paper for a more detailed discussion which is beyond the scope of our aims here.

3 Ellsberg Games

3.1 Creating ambiguity: objective Ellsberg urns

Let us formalize the intuitive idea that players can create ambiguity with the help of Ellsberg urns. An Ellsberg urn is, for us, a triple $(\Omega, \mathcal{F}, \mathcal{P})$ of a nonempty set Ω of states of the world, a σ -field \mathcal{F} on Ω (where one can take the power set in case of a finite Ω), and a set of probability measures \mathcal{P} on the measurable space (Ω, \mathcal{F}) . This *set* of probability measures represents the Knightian uncertainty of the strategy.

A typical example is the classical Ellsberg urn that contains 30 red balls, and 60 balls that are either black or yellow. One ball is drawn from that urn. The state space consists of three elements $\{R, B, Y\}$, \mathcal{F} is the power set, and \mathcal{P} the set of probability vectors (P_1, P_2, P_3) such³ that $P_1 = 1/3$, $P_2 = k/60, P_3 = (60 - k)/60$ for any $k = 0, \ldots, 60$.

We assume that the players of our game have access to and can design the parameters of such Ellsberg urns; imagine that there is an independent, trustworthy laboratory that sets up such urns and reports the outcome truthfully.

Note that we allow the player to choose the degree of ambiguity of his urn. He tells the experimentalists of his laboratory to set up such and such an Ellsberg experiment that generates exactly the set of distributions \mathcal{P}_i . In this sense, the ambiguity in our formulation of the game is "objective"; it is not a matter of agents' *beliefs* about the actions of other players, but rather a property of the device used to determine his action.

3.2 Ellsberg Games

We come now to the game where players can use such urns in addition to the usual mixed strategies (that correspond to roulette wheels or dice). Let $N = \{1, ..., n\}$ be the set of players. Each player *i* has a finite strategy set S_i . Let $S = \prod_{i=1}^n S_i$ be the set of pure strategy profiles. Players' payoffs are given by functions

 $u_i: S \to \mathbb{R}$ $(i \in N)$.

The normal form game is denoted $G = \langle N, (S_i), (u_i) \rangle$.

³We are always going to work with convex sets of probability measures. In this case, this means that we would allow for any $P_2, P_3 \ge 0$ with $P_2 + P_3 = 2/3$ here. In our framework, this is without loss of generality, of course.

Players can now use different devices. On the one hand, we assume that they have "roulette wheels" or "dices" at their disposal, i.e. randomizing devices with objectively known probabilities. The set of these probabilities over S_i is denoted ΔS_i . The players evaluate such devices according to expected utility, as in von Neumann and Morgenstern's formulation of game theory.

Moreover, and this is the new part, players can use Ellsberg urns. As we said above, we imagine that the players can credibly commit to base their actions on ambiguous outcomes. Technically, we model the Ellsberg urn of player i as a triple $(\Omega_i, \mathcal{F}_i, \mathcal{P}_i)$ as explained above.

Player i acts in the game by choosing a measurable function (or Anscombe-Aumann act)

$$f_i: (\Omega_i, \mathcal{F}_i) \to \Delta S_i$$

which specifies the classical mixed strategy played once the outcome of the Ellsberg urn is revealed. An *Ellsberg strategy* for player i is then a pair

$$((\Omega_i, \mathcal{F}_i, \mathcal{P}_i), f_i)$$

of an Ellsberg urn and an act.

To finish the description of our Ellsberg game, we have to determine players' payoffs. We suppose that all players are ambiguity-averse: in the face of ambiguous events (as opposed to simply random events) they evaluate their utility in a cautious and pessimistic way. This behavior in response to ambiguity has been observed in the famous experiments of Ellsberg (1961) and confirmed in further experiments, for example of Pulford (2009) and Camerer and Weber (1992), see also Etner, Jeleva, and Tallon (2012) for references. For our purpose we follow the axiomatization of attitude towards objective but imprecise information in Gajdos, Hayashi, Tallon, and Vergnaud (2008). In the case of extreme pessimism the utility is evaluated as a maxmin expected utility similar to the axiomatization of Gilboa and Schmeidler (1989), but with the difference of the decision maker facing objective instead of subjective ambiguity. Starting with Jaffray (1989), Giraud (2006) and Giraud and Tallon (2011) are other papers that make a case for objective ambiguity, as well as Stinchcombe (2007) and Olszewski (2007).

The payoff of player $i \in N$ at an Ellsberg strategy profile $((\Omega, \mathcal{F}, \mathcal{P}), f)$ is thus the minimal expected utility with respect to all different probability

distributions in the closed and convex set \mathcal{P} ,

$$U_i(((\Omega, \mathcal{F}, \mathcal{P}), f)) := \min_{P_1 \in \mathcal{P}_1, \dots, P_n \in \mathcal{P}_n} \int_{\Omega_1} \cdots \int_{\Omega_n} u_i(f(\omega)) \, dP_n \dots dP_1$$

We call the described larger game an *Ellsberg game*.

An *Ellsberg equilibrium* is, in the same spirit as Nash equilibrium, a profile of Ellsberg strategies

$$(((\Omega_1^*, \mathcal{F}_1^*, \mathcal{P}_1^*), f_1^*), \dots, ((\Omega_n^*, \mathcal{F}_n^*, \mathcal{P}_n^*), f_n^*))$$

where no player has an incentive to deviate, i.e. for all players $i \in N$, all Ellsberg urns $(\Omega_i, \mathcal{F}_i, \mathcal{P}_i)$, and all acts f_i for player i we have

$$U_{i}\left(\left((\Omega^{*}, \mathcal{F}^{*}, \mathcal{P}^{*}), f^{*}\right)\right) \geq U_{i}\left(\left((\Omega_{i}, \mathcal{F}_{i}, \mathcal{P}_{i}), f_{i}\right), \left((\Omega^{*}_{-i}, \mathcal{F}^{*}_{-i}, \mathcal{P}^{*}_{-i}), f^{*}_{-i}\right)\right) .$$

3.3 Reduced Form Strategies

This definition of an Ellsberg game depends on the particular Ellsberg urn used by each player i. As there are arbitrarily many possible state spaces⁵, the definition of Ellsberg equilibrium might not seem very tractable. Fortunately, there is a more concise way to define Ellsberg equilibrium. The procedure is similar to the reduced form of a correlated equilibrium, see Aumann (1974) or Fudenberg and Tirole (1991). Instead of working with arbitrary Ellsberg urns, we note that the players' payoff depends, in the end, on the set of distributions that the Ellsberg urns and the associated acts induce on the set of strategies. One can then work with that set of distributions directly.

Definition 1. Let $G = \langle N, (S_i), (u_i) \rangle$ be a normal form game. A reduced form Ellsberg equilibrium of the game G is a profile of sets of probability measures $\mathcal{Q}_i^* \subseteq \Delta S_i$, such that for all players $i \in N$ and all sets of probability measures $\mathcal{Q}_i \subseteq \Delta S_i$ we have

$$\min_{P_i \in \mathcal{Q}_i^*, P_{-i} \in \mathcal{Q}_{-i}^*} \int_{S_i} \int_{S_{-i}} u_i(s_i, s_{-i}) dP_{-i} dP_i$$

$$\geq \min_{P_i \in \mathcal{Q}_i, P_{-i} \in \mathcal{Q}_{-i}^*} \int_{S_i} \int_{S_{-i}} u_i(s_i, s_{-i}) dP_{-i} dP_i$$

⁴Throughout the paper we follow the notational convention that $(f_i, f_{-i}^*) := (f_1^*, ..., f_{i-1}^*, f_i, f_{i+1}^*, ..., f_n^*)$. The same convention is used for profiles of pure strategies (s_i, s_{-i}) and probability distributions (P_i, P_{-i}) .

⁵In fact, the class of all state spaces is too large to be a well–defined set according to set theory.

The two definitions of Ellsberg equilibrium are equivalent in the sense that every Ellsberg equilibrium induces a payoff–equivalent reduced form Ellsberg equilibrium; and every reduced form Ellsberg equilibrium is an Ellsberg equilibrium with state space $\Omega_i = S_i$ and f_i^* the constant act.

Proposition 1. Ellsberg equilibrium and reduced form Ellsberg equilibrium are equivalent in the sense that every Ellsberg equilibrium $((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)$ induces a payoff-equivalent reduced form Ellsberg equilibrium on $\Omega^* = S$; and every reduced form Ellsberg equilibrium \mathcal{Q}^* is an Ellsberg equilibrium $((S, \mathcal{F}, \mathcal{Q}^*), f^*)$ with f^* the constant act.

This is shown formally in the appendix.

We henceforth call a set $Q_i \subseteq \Delta S_i$ an Ellsberg strategy whenever it is clear that we are in the reduced form context.

3.4 Ellsberg Equilibria Generalize Nash Equilibria

Note that the classical game is contained in our formulation: players just choose a singleton $\mathcal{P}_i = \{\delta_{\pi_i}\}$ that puts all weight on a particular (classical) mixed strategy π_i .

Player 2

$$L$$
 R
Player 1 T $3,3$ $0,0$
 B $0,0$ $1,1$

Figure 1: Strategic Ambiguity does not unilaterally make a player better off.

Now let (π_1, \ldots, π_n) be a Nash equilibrium of the game G. Can any player unilaterally gain by creating ambiguity in such a situation? The answer is no. Take the game in Figure 1 and look at the pure strategy Nash equilibrium (B, R) with equilibrium payoff 1 for both players. If player 1 introduces ambiguity, he will play T in some states of the world (without knowing the exact probability of those states). But this does not help here because player 2 sticks to his strategy R, so playing T just leads to a payoff of zero. Unilateral introduction of ambiguity does not increase one's own payoff. We think that this is an important property of our formulation.

Proposition 2. Let $G = \langle N, (S_i), (u_i) \rangle$ be a normal form game. Then a mixed strategy profile (π_1, \ldots, π_n) of G is a Nash equilibrium of G if and

only if the corresponding profile of singletons $(\mathcal{P}_1, \ldots, \mathcal{P}_n)$ with $\mathcal{P}_i = \{\delta_{\pi_i}\}$ is an Ellsberg equilibrium.

In particular, Ellsberg equilibria exist when the strategy sets S_i are finite.

By including and generalizing Nash equilibria, our formulation avoids the existence pitfalls that one encounters when players are assumed to play pure strategies and beliefs are uncertain about those pure actions.

Let us next turn to the question whether non–Nash behavior can arise in Ellsberg games.

4 Non–Nash Outcomes: Strategic Ambiguity in Negotiation Games

Strategic ambiguity can lead to new phenomena that lie outside the scope of classical game theory. As our first example, we consider the following peace negotiation game taken from Greenberg (2000). There are two small countries who can either opt for peace, or war. If both countries opt for peace, all three players obtain a payoff of 4. If one of the countries does not opt for peace, war breaks out, but the superpower cannot decide whose action started the war. The superpower can punish one country and support the other. The game tree is in Figure 2 below.⁶



Figure 2: Peace Negotiation

⁶We take the payoffs as in Greenberg's paper. In case the reader is puzzled by the slight asymmetry between country A and B in payoffs: it does not play a role for our argument. One could replace the payoffs 3 and 6 for country A by 0 and 9.

As we deal here only with static equilibrium concepts, we also present the normal form, where country A chooses rows, country B columns, and the superpower chooses the matrix.



Figure 3: Peace Negotiation in normal form

This game possesses a unique Nash equilibrium where country A mixes with equal probabilities, and country B opts for war; the superpower has no clue who started the war given these strategies. It is thus indifferent about whom to punish and mixes with equal probabilities as well. War occurs with probability 1. The resulting equilibrium payoff vector is (4.5, 4.5, 0.5).

If the superpower can create ambiguity (and if the countries A and B are ambiguity-averse), the picture changes. Suppose for simplicity, that the superpower creates maximal ambiguity by using a device that allows for any probability between 0 and 1 for its strategy *punish* A. The pessimistic players A and B are ambiguity-averse and thus maximize against the worst case. For both of them, the worst case is to be punished by the superpower, with a payoff of 0. Hence, both prefer to opt for peace given that the superpower creates ambiguity. As this leads to a very desirable outcome for the superpower, it has no incentive to deviate from this strategy. We have thus found an equilibrium where the strategic use of ambiguity leads to an equilibrium outcome.

Let us formalize the above considerations. We claim that there is the following type of Ellsberg equilibria. The superpower creates ambiguity about its decision; if this ambiguity is sufficiently large, both players fear to be punished by the superpower in case of war. As a consequence, they opt for peace.

In our game with just two actions for the superpower, we can identify an Ellsberg strategy with an interval $[P_0, P_1]$ where $P \in [P_0, P_1]$ is the probability that the superpower punishes country A. Suppose the superpower plays so with $P_0 < \frac{4}{9}$ and $P_1 > \frac{5}{9}$. Assume also that country B opts for *peace*. If A goes for war, it uses that prior in $[P_0, P_1]$ which minimizes its expected payoff, which is P_1 . This yields $U_A(war, war, [P_0, P_1]) = P_1 \cdot 0 + (1 - P_1) \cdot 9 < 4$.

Hence, opting for peace is country A's best reply. The reasoning for country B is similar, but with the opposite probability P_0 . If both countries A and B go for peace, the superpower gets 4 regardless of what it does; in particular, the ambiguous strategy described above is optimal. We conclude that (*peace*, *peace*, $[P_0, P_1]$) is a (reduced form) Ellsberg equilibrium.

Proposition 3. In Greenberg's game, the strategies (peace, peace, $[P_0, P_1]$) with $P_0 < \frac{4}{9}$ and $P_1 > \frac{5}{9}$ form an Ellsberg equilibrium.

Note that this Ellsberg equilibrium is very different from the game's unique Nash equilibrium. In Nash equilibrium, war occurs in every play of the game; in our Ellsberg equilibrium, peace is the unique outcome.⁷ By using the strategy $[P_0, P_1]$ which is a set of probability distributions, the superpower *creates* ambiguity. This supports an Ellsberg equilibrium where players' strategies do not lie in the support of the unique Nash equilibrium. We also point out that the countries A and B use different *worst-case priors* in equilibrium; this is a typical phenomenon in Ellsberg equilibria that are supported by strategies which are not in the support of any Nash equilibrium of the game.

Greenberg refers to historic peace negotiations between Israel and Egypt (countries A and B in the negotiation example) mediated by the USA (superpower C) after the 1973 war. As explained by Kissinger (1982)⁸, the fact that both Egypt and Israel were too afraid to be punished if negotiations broke down partly contributed to the success of the peace negotiations. This story is supported by our Ellsberg equilibrium, a first evidence that Ellsberg equilibria might capture some real world phenomena better than Nash equilibria.

⁷Other equilibrium concepts for extensive form games (without Knightian uncertainty) such as conjectural equilibrium Battigalli and Guaitoli (1988), self-confirming equilibrium Fudenberg and Levine (1993), and subjective equilibrium Kalai and Lehrer (1995) can also assure the peace equilibrium outcome in the example by Greenberg. Other equillibrium concepts for extensive form games with Knightian uncertainty are e.g. Battigalli, Cerreia-Vioglio, Maccheroni, and Marinacci (2011) and Lo (1999). Postponing the analysis of the relation of these equilibrium concepts to Ellsberg equilibrium to a later paper, we only want to stress here that in difference to the existing concepts the driving factor in Ellsberg equilibrium is that ambiguity is employed strategically and objectively.

⁸See p. 802 therein, in particular.

5 Strategic Ambiguity in Two–Person Conflicts

Our approach to games has its most natural and fruitful applications to conflicts where players are at least to some degree in opposition to each other. We start this section by discussing a modified version of Matching Pennies to illustrate the phenomena of *immunization against ambiguity* and *nonlinearity of payoffs* that arise in Ellsberg games. We then provide a general analysis of 2×2 -conflict games. While our predictions are broader than the classical Nash equilibrium, they remain restrictive, and, at least in principle, testable. Our results do allow to explain the experimental findings of Goeree and Holt (2001) who show that humans tend to deviate from Matching Pennies in asymmetric modified Matching Pennies, but tend to play Nash equilibrium in symmetric Matching Pennies. This corresponds and is consistent with our Ellsberg equilibria.

5.1 A Matching Pennies Example

Let us now consider a modified version of Matching Pennies. The payoff matrix for this game is in Figure 4.

		Player 2		
		HEAD	TAIL	
Player 1	HEAD	3, -1	-1, 1	
	TAIL	-1, 1	1, -1	

Figure 4: Modified Matching Pennies I

We point out two effects that arise due to strategic ambiguity in this class of games. On the one hand, the Ellsberg equilibria are different from what one might expect first; in a game like the one above, one might intuitively guess that "full ambiguity" would be an Ellsberg equilibrium, as the natural generalization of "full randomness" (completely mixed Nash equilibrium). This is not the case.

On the other hand, we emphasize an important property of Ellsberg games (or ambiguity aversion in general): the best reply functions are no longer linear in the probabilities. As a consequence, the indifference principle of classical game theory – when two pure strategies yield the same payoff, then the player is indifferent about mixing in any arbitrary way between the two strategies – does not carry over to Ellsberg games. When a player is indifferent between two Anscombe–Aumann acts, this does not imply that he is indifferent between all mixtures over these two acts. This is due to the hedging or diversification effect provided by a (classical) mixed strategy when players are ambiguity–averse. We call this effect *immunization against strategic ambiguity*.

Immunization against Strategic Ambiguity

In our modified version of Matching Pennies, the unique Nash equilibrium is that player 1 mixes uniformly over his strategies, and player 2 mixes with (1/3, 2/3). This yields the equilibrium payoffs 1/3 and 0. One might guess that one can get an Ellsberg equilibrium where both players use a set of probability measures around the Nash equilibrium distribution as their strategy. This is not true.

The crucial point to understand here is the following. Players can immunize themselves against ambiguity; in the modified Matching Pennies example, player 1 can use the mixed strategy (1/3, 2/3) to make himself independent of any ambiguity used by the opponent. Indeed, with this strategy, his expected payoff is 1/3 against any mixed strategy of the opponent, and a fortiori against Ellsberg strategies as well. This strategy is also the unique best reply of player 1 to Ellsberg strategies with ambiguity around the Nash equilibrium; in particular, such strategic ambiguity is not part of an Ellsberg equilibrium.

Let us explain this somewhat more formally. An Ellsberg strategy for player 2 can be identified with an interval $[Q_0, Q_1] \subseteq [0, 1]$ where $Q \in [Q_0, Q_1]$ is the probability to play *HEAD*. Suppose player 2 uses many probabilities around 1/3, so $Q_0 < 1/3 < Q_1$. The (minimal) expected payoff for player 1 when he uses the mixed strategy with probability *P* for *HEAD* is then

$$\begin{split} \min_{Q_0 \le Q \le Q_1} 3PQ - P(1-Q) &- (1-P)Q + (1-P)(1-Q) \\ &= \min \left\{ Q_0(6P-2), Q_1(6P-2) \right\} + 1 - 2P \\ &= \begin{cases} Q_1(6P-2) + 1 - 2P & \text{if } P < 1/3 \\ 1/3 & \text{if } P = 1/3 \\ Q_0(6P-2) + 1 - 2P & \text{else} . \end{cases} \end{split}$$

We plot the payoff function in Figure 5.



Figure 5: Player 1's (minimal expected) payoff as a function of the probability P of playing HEAD when player 2 uses the Ellsberg strategy [1/4, 1/2].

By choosing the mixed strategy P = 1/3, player 1 becomes immune against any ambiguity and ensures the (Nash) equilibrium payoff of 1/3. If there was an Ellsberg equilibrium with $P_0 < 1/2 < P_1$ and $Q_0 < 1/3 < Q_1$, then the minimal expected payoff would be below 1/3. Hence, such Ellsberg equilibria do not exist.

Such immunization plays frequently a role in two-person games, and it need not always be the Nash equilibrium strategy that is used to render oneself immune. Consider, e.g., the slightly changed payoff matrix

		Player 2		
		HEAD	TAIL	
Player 1	HEAD	1, -1	-1, 1	
	TAIL	-2, 1	1, -1	

Figure 6: Modified Matching Pennies II

In the unique Nash equilibrium, player 1 still plays both strategies with probability 1/2 (to render player 2 indifferent); however, in order to be im-

mune against Ellsberg strategies, he has to play HEAD with probability 3/5. Then his payoff is -1/5 regardless of what player 2 does. This strategy does not play any role in either Nash or Ellsberg equilibrium. It is only important in so far as it excludes possible Ellsberg equilibria by being the unique best reply to some Ellsberg strategies.

Ellsberg Equilibria

The question thus arises if there are any Ellsberg equilibria different from the Nash equilibrium at all. There are, and they take the following form for our first version of modified Matching Pennies (Figure 4). Player 1 plays HEAD with probability $P \in [1/2, P_1]$ for some $1/2 \leq P_1 \leq 1$ and player 2 plays HEAD with probability $Q \in [1/3, Q_1]$ for some $1/3 \leq Q_1 \leq 1/2$. This Ellsberg equilibrium yields the same payoffs 1/3 and 0 as in Nash equilibrium. We prove a more general theorem covering this case in the next section.

Proposition 4. In Modified Matching Pennies I, the Ellsberg equilibria are of the form $([1/2, P_1], [1/3, Q_1])$ for $1/2 \le P_1 \le 1$ and $1/3 \le Q_1 \le 1/2$.

The typical Ellsberg equilibrium strategy takes the following form. Player 1 says : "I will play HEAD with a probability of at least 50%, but not less." And Player 2 replies: "I will play HEAD with at least 33%, but not more than 50%."

5.2 General Conflict Games

Let us now generalize the above results. Consider the competitive two–person 2×2 game with payoff matrix

Player 2

$$L R$$

Player 1 $U a, d b, e$
 $D b, e c, f$

such that

$$a, c > b$$
 and $d, f < e$.

Due to these assumptions, our game has conflicting interests and no pure strategy Nash equilibria. In the unique mixed strategy Nash equilibrium, player 1 plays U with probability

$$P^* = \frac{f - e}{d - 2e + f},$$

and player 2 plays L with probability

$$Q^* = \frac{c-b}{a-2b+c}.$$

Proposition 5. The Ellsberg equilibria of the above game are the following:

For $P^* > Q^*$ all Ellsberg equilibria are of the form

 $([P^*, P_1], [Q^*, Q_1]) \text{ for } P^* \leq P_1 \leq 1, Q^* \leq Q_1 \leq P^*;$

for $P^* < Q^*$ all Ellsberg equilibria are of the form

$$([P_0, P^*], [Q_0, Q^*]) \text{ for } 0 \le P_0 \le P^*, P^* \le Q_0 \le Q^*;$$

and for $P^* = Q^*$ all Ellsberg equilibria are of the form

 $(Q^*, [Q_0, Q_1])$ where $Q_0 \le Q^* \le Q_1$

and

$$([P_0, P_1], P^*)$$
 where $P_0 \leq P^* \leq P_1$.

There are several noteworthy properties here. First of all, in symmetric Matching Pennies games, including the 2×2 zero sum games, essentially no new equilibria arise. Symmetry implies that the immunization strategy and the Nash equilibrium strategy coincide, leaving no room for nontrivial Ellsberg equilibria beyond the artificial fact that one player can use some ambiguity given that the other player plays the immunization strategy that renders him immune to any ambiguity. But we do interpret this kind of equilibria as meaning that one can only play Nash in symmetric conflicts.

The picture changes for asymmetric conflicts. In that case, some ambiguity is possible; however, the ambiguity is not *around* the Nash equilibrium, but rather the Nash equilibrium is a bound for the set of probabilities used in the Ellsberg strategies. For one player, the probability bounds are exactly given by the Nash equilibrium probabilities.⁹

⁹This is, however, partly due to the special case treated in Proposition 5. When (D, L) and (U, R) do not yield the same payoff, the probability bounds are the Nash equilibrium strategy on one side and the immunization strategy on the other. See also Remark 1 in the appendix following the proof of Proposition 5.

We also see that our theory leads to a broader set of equilibria as the classical theory, but it is not arbitrary. The probabilities used in Ellsberg equilibrium do have to satisfy certain nontrivial bounds. In particular, one can test these bounds in the lab. We will come back to testability below.

5.3 Human Behavior in Matching Pennies Games and Ellsberg Equilibria

Whereas the support of the Ellsberg and Nash equilibria is obviously the same, we do think that the Ellsberg equilibria reveal a new class of behavior not encountered in game theory before. It might be very difficult for humans to play exactly a randomizing strategy with equal probabilities; indeed, the ability to do so has been a debate since the early days of game theory, and some claim that humans cannot randomize, see Dang (2009) for a recent account and references therein. Our result shows that it is not necessary to randomize exactly to support a similar equilibrium outcome (with the same expected payoff). It is just enough that your opponent knows that you are randomizing with some probability, and that it could be that this probability is one half, but not less. It is thus sufficient that the player is able to control the lower bound of his device. This might be easier to implement than the perfectly random behavior required in classical game theory.

In fact, there are experimental findings which suggest that the Ellsberg equilibrium strategy in the modified Matching Pennies game is closer to real behavior than the Nash equilibrium prediction. To illustrate this, let us consider the interesting results by Goeree and Holt (2001) who ran experiments on three different versions of Matching Pennies; the three payoff matrices can be seen in Table 1.

In the first game, we have a typical symmetric conflict game with a unique mixed Nash equilibrium in which both players randomize uniformly over both pure strategies. The aggregate play of humans in the experiment is closely consistent with the Nash equilibrium prediction, 48% of players choosing "Top" or "Left", resp.

Remember that the probabilities in a mixed strategy equilibrium are chosen in such a way as to render the opponent indifferent between her two pure strategies. As a consequence, if we change the payoffs of player 1 only (while keeping the ordering of payoffs), his Nash equilibrium strategy does not change because he has to make player 2 indifferent between her two pure

		Left (48)	Right (52)
Symmetric	Top (48)	80,40	40,80
	Bottom (52)	40,80	80,40
		Left (16)	$\operatorname{Right}(84)$
Asymmetric	Top (96)	320 , 40	40,80
	Bottom (4)	40,80	80,40
		Left (80)	Right (20)
Reversed	Top (8)	44 ,40	40,80
	Bottom (92)	40,80	80, 40

Table 1: The Goeree–Holt Results on three different versions of Matching Pennies.

actions, and her payoffs have not been modified.

In the second game, called the asymmetric Matching Pennies game, player 1 gets 320 instead of 80 in the upper left outcome. All other payoffs remain the same. Many humans now deviate from Nash, as is reported in brackets, 96% of the players taking the action "Top". Interestingly, also the humans playing the role of player 2 change their behavior, and most of them play "Right", the best reply to "Top".

In the third case, player 1's payoff in the upper left outcome is decreased to a lowly 44. Then only 8% of players choose "Top"; 80% of humans in the role of player 2 choose "Left".

While aggregate behavior by humans is certainly inconsistent with the predictions of Nash equilibrium, it is consistent with Ellsberg equilibria. We summarize the results in Table 2.

In the symmetric game, our Proposition 5 essentially predicts only Nash equilibrium behavior, and this is what we observe in the experiment as well.

In the asymmetric Matching Pennies game, the Nash equilibrium strategies are $P^* = 1/2$ for player 1 and $Q^* = 1/8$ for player 2. According to our proposition, the Ellsberg equilibria allow for probabilities in the interval [1/2, 1] for player 1 choosing "Top", and for the interval [1/8, 1/2] for player 2 choosing "Left". The observed percentages of 96% and 16% do lie in these intervals.

And in the "reversed" version of the game, the Nash equilibrium strategies are $P^* = 1/2$ and $Q^* = 10/11$. So we have the reversed relation $Q^* > P^*$. The Ellsberg equilibria allow for probabilities for "Top" in the interval [0, 1/2]

Game	Nash Equilibrium		Ellsberg Equilibrium		Observations	
	Player 1	Player 2	Player 1	Player 2	Player 1	Player 2
symmetric	0.5	0.5	0.5	0.5	0.48	0.48
asymmetric	0.5	0.125	[0.5,1]	[0.125, 0.5]	0.96	0.16
reversed	0.5	$0.\overline{90}$	[0, 0.5]	$[0.5, 0.\overline{90}]$	0.08	0.8

Table 2: Comparison of Nash and Ellsberg Predicitions with the Experimantal Observations. We record the probabilities (or intervals of probabilities) for each player to play the first pure strategy ("Top" resp. "Left") and the observed aggregate frequency of these actions in the Goeree–Holt experiments.

for player 1, and for probabilities in [1/2, 10/11] for player 2. The aggregate observed quantities of 8% and 80% do lie in these intervals.

5.4 Observational Implications of Game Theory and Human Behavior

Game Theory studies equilibrium outcomes of social conflicts when rational agents interact. Human beings are quite different from rational agents in general, so one can only expect to see a consistency with Nash equilibrium predictions and human behavior when the situation is controlled in such a way as to bring out the rational part of humans.

Nevertheless, it does make sense to ask what the observational implications of our theory are. For three player games, this is quite clear, as our theory predicts new equilibria outside the support of Nash equilibria; this is a testable implication, and we shall proceed one day to carry out such a test.

For two player games, the situation is more subtle. Both the Nash equilibrium and the Ellsberg equilibria have full support, so the only thing that we can learn from our theory seems to be that either action is fine in a one shot game. This is indeed the stance of Bade (2011), in line with a number of predecessors.

There is, however, a way to distinguish the predictions of Ellsberg equilibria and Nash equilibria even in two player games. To understand this, we need first explain how the law of large numbers looks like under ambiguity. The classical law states that the frequency of HEAD in an infinite sequence of independent coin tosses will converge to the probability of HEAD. Now let us look at a typical Ellsberg urn that contains 100 balls, red and black, and we only know that the number of red balls is between 30 and 60. What can we say about the average frequency drawn from independent repetitions of the Ellsberg experiment? The natural guess would be that the average lies in the interval between 30% and 60% in the long run. This is indeed correct, and mathematical versions of that theorem have recently been proven, see Maccheroni and Marinacci (2005) and Epstein and Schneider (2003), e.g. Peng (2007) has obtained the result that the average frequency will indeed fluctuate between both bounds, and *every* point in the interval [0.3, 0.6] is an accumulation point of the sequence.

What is then the empirical content of such laws of large numbers? If we adhere to the point of view that our observed humans play independently one shot games, and that they should play equilibrium strategies, then the average frequency will converge to the Nash equilibrium strategy according to the classical theory, and will fluctuate between two bounds according to the new Ellsberg theory.

We thus do get observational differences between the two theories; and we interpret the Goeree–Holt results as a first evidence that our theory can accommodate deviations from Nash equilibrium observed in laboratories.

6 Related Literature and Discussion

Several authors introduce Knightian uncertainty into complete–information normal form games. We discuss their concepts and compare them to our approach.

Dow and Werlang (1994), Lo (1996), Marinacci (2000), Eichberger and Kelsey (2000) and Eichberger, Kelsey, and Schipper (2009) all extend the interpretation of Nash equilibrium as an equilibrium in beliefs. For example, Dow and Werlang (1994) interpret their non-additive (Choquet) probabilities as uncertain beliefs about the other player's action. A pair (P_1, P_2) of nonadditive probabilities is then a Nash equilibrium under Knightian uncertainty if each action in a support of player 1's belief P_1 is optimal given that he uses P_2 to evaluate his expected payoff, and similarly for player 2. We thus have here a first version of an equilibrium in beliefs. This approach is refined by Marinacci (2000) and extended to *n*-person games by Eichberger and Kelsey (2000).

Lo (1996) introduces the concept of equilibrium in beliefs under uncertainty where the beliefs are represented by multiple priors over other players' mixed strategies. Each player i has a set of beliefs B_i over what the other players do, so over ΔS_{-i} . The profile (B_i) then forms a beliefs equilibrium if player j puts positive weight only on strategies of player i that maximize i's minimal expected payoff given the belief set B_i . This concept allows for disagreement of players' beliefs, and for correlation. Lo therefore introduces the refinement of a beliefs equilibrium with agreement in which player j and k agree about player i's actions and the beliefs of i over j and k are independent. Lo proves the nice result that every beliefs equilibrium contains a Bayesian beliefs equilibrium (where the belief sets are singletons). As a corollary, he obtains a precursor of Bade (2011)'s main theorem (which we discuss in a later paragraph): in two player games, every beliefs equilibrium contains a Nash equilibrium.

Note that all the equilibrium concepts discussed above do not specify which action will actually be played in equilibrium. In Lo (1996) players can play any pure or mixed strategy that is a best response to their belief set, in the other equilibrium notions mentioned, players only have access to pure strategies in the support of the capacities. This stands in contrast to Ellsberg equilibrium, where the equilibrium strategy is fixed by the Ellsberg urn chosen. The strategy is a best response to the belief, and the belief coincides with the strategy played.

Klibanoff (1996), Lehrer (2008) and Lo (2009) propose an approach similar to beliefs equilibrium. Uncertainty is present in players' beliefs that are represented by sets of distributions. Equilibrium is defined as a profile of beliefs and an objectively mixed (or pure) strategy for each player, which is the strategy that he plays in equilibrium. These strategies need to be contained in the belief sets. Accordingly, players have to anticipate their opponents' strategy correctly in the sense that the truth is part of their belief. This consistency requirement is weaker than in Nash equilibrium (and weaker than in Ellsberg equilibrium!) and typically the strategies in equilibrium are not best responses to the actual strategies played. Klibanoff (1996) proposes a refinement where only correlated rationalizable beliefs are allowed.¹⁰ Lehrer (2008) develops a model of decision making under uncertainty with partially–specified probabilities, these are used to represent the players' uncertain beliefs about their opponents. Lo (2009) establishes formal epistemic

 $^{^{10}}$ Lo (1996) requires every probability distribution in the belief sets to be a best response, therefore every beliefs equilibrium with agreement is a refinement of equilibrium with uncertainty aversion and rationalizable beliefs (this is shown in Lo (1996), Proposition 9).

foundations for an equilibrium concept with ambiguity-averse preferences. He finds that epistemically stochastic independence is not necessary for a generalized Nash equilibrium concept. A correlated Nash equilibrium is a pair $\langle \sigma, \Phi \rangle$ consisting of a profile of beliefs Φ_i and a profile of mixed strategies σ_i where, for consistency, each strategy a_i in the support of σ_i is a best response to the belief Φ_i .

Bade (2011) goes a first step in another direction, away from the beliefs interpretation of Nash equilibrium. She allows players to use acts in the sense of Anscombe–Aumann and players are uncertainty–averse over such acts.¹¹ In an ambiguous act equilibrium, players play best responses as in Nash equilibrium, but under the generalized framework. A large class of ambiguity-averse preferences are covered. The possible priors for an ambiguous act are part of the players' preferences in her setup. Bade then adds some appropriate consistency properties (agreement on null events) to exclude unreasonably divergent beliefs, and she imposes the rather strong assumption that preferences are *strictly* monotone, following Klibanoff (1996) here. This excludes beliefs on the boundary of strategy sets; such degenerated beliefs are sometimes important, though. For example, it excludes Ellsberg urns with full ambiguity where it is only known that the probability for a red ball is between 0 and 1. Bade's main theorem establishes that under her assumptions, in two-person games the support of ambiguous act equilibria and the support of Nash equilibria coincide.

In difference to her setup we let ambiguity be an objective instrument that is not derived from subjective preferences. Players can credibly commit to play an Ellsberg urn with a given and known degree of ambiguity. In Ellsberg games players use devices that create ambiguity, thus we extend the objective random devices interpretation of Nash equilibrium. The articles cited above impose non-expected utility representations derived from subjective preferences, like maxmin expected utility by Gilboa and Schmeidler (1989), Choquet expected utility by Schmeidler (1989), or they fix only certain axioms to allow for a large class of ambiguity-averse preferences. To

¹¹Aumann (1974), Epstein (1997) and Azrieli and Teper (2011) (amongst others) have also defined games that have Anscombe–Aumann acts as strategies, but to different ends. Aumann (1974) defines such a general game, then imposes Savage expected utility and analyses properties of correlated and subjective equilibrium. Epstein (1997) analyses games very similar to Bade's, but is mainly interested in rationalizability and iterated deletion of strictly dominated strategies in the generalized framework. Azrieli and Teper (2011) define an extension of an incomplete–information game.

model the preferences in Ellsberg games we use the representation results by Gajdos, Hayashi, Tallon, and Vergnaud (2008) on attitude towards imprecise information which capture the objective ambiguity we have in mind.

7 Conclusion

This article demonstrates that the strategic use of ambiguity is a relevant concept in game theory. Employing objective ambiguity as a strategic instrument leads to a new class of equilibria not encountered in classic game theory. We point out that in many games players choose to be deliberately ambiguous to gain a strategic advantage.

In some games this results in equilibrium outcomes which cannot be obtained as Nash equilibria. The peace negotiation game provides an example of such Ellsberg equilibria. Games with more than two players offer a strategic possibility that is not available in two-person games, because a third player is able to induce the use of different probability distributions. Although countries A and B observe the same Ellsberg strategy played by the superpower C, due to their ambiguity aversion the countries use different probability distributions to assess their utility. We plan to say more on this power of the third player, as well as on immunization against strategic ambiguity in games with more than two players, in a companion paper.

However, also two-person 2×2 games with conflicting interests have Ellsberg equilibria which are different from classic mixed strategy Nash equilibria. There are equilibria in which both players create ambiguity. They use an Ellsberg strategy where they only need to control the lower (or upper) bound of their set of probability distributions. We argue that this device is easier to use for a player than playing one precise probability distribution like in mixed strategy Nash equilibrium. What makes this argument attractive is that the payoffs in these Ellsberg equilibria are the same as in the unique mixed Nash equilibrium and thus the use of ambiguous strategies in competitive games is indeed an option. Our argument is strengthened by experimental results. Without any further assumptions besides ambiguity aversion, Ellsberg equilibria can explain human non-Nash behavior in modified Matching Pennies games. In symmetric Matching Pennies, humans tend to play the Nash equilibrium which is also in line with our result that essentially no new equilibria emerge in such symmetric games.

A Equivalence of the Different Formulations of Ellsberg Equilibrium

We provide here the proof of Proposition 1. First we recap the definition of an Ellsberg equilibrium, which was stated in the text of Section 3.

Definition 2. Let $G = \langle N, (S_i), (u_i) \rangle$ be a normal form game. A profile

$$(((\Omega_1^*, \mathcal{F}_1^*, \mathcal{P}_1^*), f_1^*), ..., ((\Omega_n^*, \mathcal{F}_n^*, \mathcal{P}_n^*), f_n^*))$$

of Ellsberg strategies is an Ellsberg equilibrium of G if no player has an incentive to deviate from $((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)$, i.e. for all players $i \in N$, all Ellsberg urns $(\Omega_i, \mathcal{F}_i, \mathcal{P}_i)$ and all acts f_i for player i we have

$$U_i(((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)) \ge U_i(((\Omega_i, \mathcal{F}_i, \mathcal{P}_i), f_i), ((\Omega^*_{-i}, \mathcal{F}^*_{-i}, \mathcal{P}^*_{-i}), f^*_{-i})), \text{ that is}$$

$$\min_{P_i \in \mathcal{P}_i^*, P_{-i} \in \mathcal{P}_{-i}^*} \int_{\Omega_i^*} \int_{\Omega_{-i}^*} u_i(f_i^*(\omega_i), f_{-i}^*(\omega_{-i})) dP_{-i} dP_i$$

$$\geq \min_{P_i \in \mathcal{P}_i, P_{-i} \in \mathcal{P}_{-i}^*} \int_{\Omega_i} \int_{\Omega_{-i}^*} u_i(f_i(\omega_i), f_{-i}^*(\omega_{-i})) dP_{-i} dP_i$$

The definition of the reduced form Ellsberg equilibrium was given in Definition 1.

Proof. " \Leftarrow " Let \mathcal{Q}^* be an Ellsberg equilibrium according to Definition 1. We choose the states of the world $\Omega = S$ to be the set of pure strategy profiles, thereby we see that player *i* uses the Ellsberg urn $(S_i, \mathcal{F}_i, \mathcal{Q}_i^*)$. We define the act $f_i^* : (S_i, \mathcal{F}_i) \to \Delta S_i$ to be the constant act that maps $f_i^*(s_i) = \{\delta_{s_i}\}$. $\{\delta_{s_i}\} \in \Delta S_i$ is the degenerate mixed strategy which puts all weight on the pure strategy s_i . Each measure $Q_i \in \mathcal{Q}_i^*$ has an image measure under f_i^* ,

$$Q_i \circ f_i^{*^{-1}} : \{\delta_{s_i}\} \mapsto Q_i(f_i^{*^{-1}}(\{\delta_{s_i}\})).$$

 $Q_i \circ f_i^{*^{-1}}$ can be identified with $Q_i \in \mathcal{Q}_i^*$. Thus the reduced form Ellsberg equilibrium strategy \mathcal{Q}_i^* can be written as the Ellsberg strategy $((S, \mathcal{F}, \mathcal{Q}^*), f^*)$. This strategy is an Ellsberg equilibrium according to Definition 2.

" \Rightarrow " Let now $((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)$ be an Ellsberg equilibrium according to Definition 2. Every $P_i \in \mathcal{P}_i^*$ induces an image measure $P_i \circ f_i^{*^{-1}}$ on ΔS_i that assigns a probability to a distribution $f_i^*(\omega_i) \in \Delta S_i$ to occur.

To describe the probability that a pure strategy s_i is played, given a distribution P_i and an Ellsberg strategy $((\Omega_i^*, \mathcal{F}_i^*, \mathcal{P}_i^*), f_i^*)$, we integrate $f_i^*(\omega_i)(s_i)$ over all states $\omega_i \in \Omega_i$. Thus we can define Q_i to be:

$$Q_i(s_i) := \int_{\Omega_i^*} f_i^*(\omega_i)(s_i) \, dP_i. \tag{1}$$

Recall that \mathcal{P}_i is a closed and convex set of probability distributions. We get a measure Q_i on S_i for each $P_i \in \mathcal{P}_i \subseteq \Delta \Omega_i$. We call the resulting set of probability measures \mathcal{Q}_i^* .

$$\mathcal{Q}_i^*(s_i) := \left\{ Q_i(s_i) = \int_{\Omega_i^*} f_i^*(\omega_i)(s_i) \, dP_i \mid P_i \in \mathcal{P}_i^* \right\}.$$

 \mathcal{Q}_i^* is closed and convex, since \mathcal{P}_i^* is.

Now suppose \mathcal{Q}^* was not a reduced form Ellsberg equilibrium. Then for some player $i \in N$ there existed a set \mathcal{Q}_i of probability measures on S_i that yields a higher minimal expected utility. This means we would have

$$\min_{Q_{i}\in\mathcal{Q}_{i},Q_{-i}\in\mathcal{Q}_{-i}^{*}} \int_{S_{i}} \int_{S_{-i}} u_{i}(s_{i},s_{-i}) \, dQ_{-i} dQ_{i}
> \min_{Q_{i}\in\mathcal{Q}_{i}^{*},Q_{-i}\in\mathcal{Q}_{-i}^{*}} \int_{S_{i}} \int_{S_{-i}} u_{i}(s_{i},s_{-i}) \, dQ_{-i} dQ_{i}$$
(2)

for some $Q_i \neq Q_i^*$. Let Q'_i be the minimizer of the first expression, then it must be that $Q'_i \notin Q_i^*$. We know that Q'_i is derived from some some P'_i under the equilibrium act,

$$Q'_{i}(s_{i}) = \int_{\Omega_{i}^{*}} f_{i}^{*}(\omega_{i})(s_{i}) \, dP'_{i}.$$
(3)

It follows that P'_i is not element of the equilibrium Ellsberg urn $(\Omega^*_i, \mathcal{F}^*_i, \mathcal{P}^*_i)$, that is $P'_i \notin \mathcal{P}^*_i$. Now it remains to show that in the original game P'_i yields a higher minimal expected utility than using \mathcal{P}^*_i . In that case $((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)$ is not an Ellsberg equilibrium and the proof is complete.

Let player i use P'_i in his maxmin expected utility evaluation in the original game. This yields

$$\min_{P_{-i}\in\mathcal{P}_{-i}^{*}} \int_{\Omega_{i}^{*}} \int_{\Omega_{-i}^{*}} u_{i}(f_{i}^{*}(\omega_{i}), f_{-i}^{*}(\omega_{-i})) dP_{-i}dP_{i}' \tag{4}$$

$$= \min_{P_{-i}\in\mathcal{P}_{-i}^{*}} \int_{\Omega_{i}^{*}} \int_{\Omega_{-i}^{*}} \int_{S_{i}} \int_{S_{-i}} u_{i}(s_{i}, s_{-i}) df_{-i}^{*}(\omega_{-i}) df_{i}^{*}(\omega_{i}) dP_{-i}dP_{i}'.$$

Recall that we use u_i to be the utility function on S_i as well as on ΔS_i . We use equations (1) and (3) to rewrite the expression and get

$$\min_{Q_{-i}\in\mathcal{Q}_{-i}^{*}} \int_{S_{i}} \int_{S_{-i}} u_{i}(s_{i}, s_{-i}) \, dQ_{-i} dQ'_{i}.$$
(5)

We know by equation (2) that this is larger than the minimal expected utility over Q_i^* and this gives

$$(5) > \min_{Q_i \in \mathcal{Q}_i^*, Q_{-i} \in \mathcal{Q}_{-i}^*} \int_{S_i} \int_{S_{-i}} u_i(s_i, s_{-i}) \, dQ_{-i} dQ_i$$
$$= \min_{P_i \in \mathcal{P}_i^*, P_{-i} \in \mathcal{P}_{-i}^*} \int_{\Omega_i^*} \int_{\Omega_{-i}^*} u_i(f_i^*(\omega_i), f_{-i}^*(\omega_{-i})) \, dP_{-i} dP_i$$

Going back to equation (4) we see that this contradicts the assumption that $((\Omega^*, \mathcal{F}^*, \mathcal{P}^*), f^*)$ was an Ellsberg equilibrium. Thus \mathcal{Q}^* is a reduced form Ellsberg equilibrium.

B Strategic Ambiguity in Two–Person Conflicts

We provide here the proof of Proposition 5.

Proof. The game has a unique Nash equilibrium in which player 1 plays U with probability

$$P^* = \frac{f - e}{d - 2e + f},$$

and player 2 plays L with probability

$$Q^* = \frac{c-b}{a-2b+c} \,.$$

The Nash equilibrium strategies follow from the usual analysis. The conditions on the payoffs assure that the Nash equilibrium is completely mixed, i.e. $0 < P^* < 1$ and $0 < Q^* < 1$.

Let now $[P_0, P_1]$ and $[Q_0, Q_1]$ be Ellsberg strategies of player 1 and 2, where $P \in [P_0, P_1]$ is the probability of player 1 to play U, and $Q \in [Q_0, Q_1]$ is the probability of player 2 to play L. Let us compute the minimal expected payoff.

The minimal expected payoff of player 1 when he plays the mixed strategy P is

$$\begin{split} \min_{Q_0 \leq Q \leq Q_1} u_1(P,Q) &= \min_{Q_0 \leq Q \leq Q_1} aPQ + bP(1-Q) + b(1-P)Q + c(1-P)(1-Q) \\ &= \min_{Q_0 \leq Q \leq Q_1} Q(b-c+P(a-2b+c)) + bP + c - cP \\ &= \begin{cases} Q_1(b-c+P(a-2b+c)) + bP + c - cP & \text{if } P < \frac{c-b}{a-2b+c} \\ \frac{ac-b^2}{a-2b+c} & \text{if } P = \frac{c-b}{a-2b+c} \\ Q_0(b-c+P(a-2b+c)) + bP + c - cP & \text{else} . \end{cases} \end{split}$$

Note that the payoff function has a fixed value at $P = \frac{c-b}{a-2b+c}$, which is player 2's Nash equilibrium strategy. Depending on Q_0 and Q_1 the minimal payoff function can have six different forms. It can be strictly increasing, strictly decreasing, have flat parts or be completely constant. To determine how player 1 maximizes his minimal payoff for different Q_0 and Q_1 we look at the borders of the minimal payoff function, where P = 0 and P = 1. This gives us two functions

$$\min_{Q_0 \le Q \le Q_1} u_1(0,Q) = Q_1(b-c) + c \quad \text{and} \quad \min_{Q_0 \le Q \le Q_1} u_1(1,Q) = Q_0(a-b) + b.$$

Note that b - c < 0 and a - b > 0, that is, the minimal payoff function is decreasing with Q_1 at P = 0 and increasing with Q_0 at P = 1. When

$$Q_1 = \frac{c-b}{a-2b+c} = Q^*$$
, then $\min_{Q_0 \le Q \le Q_1} u_1(0,Q) = \frac{ac-b^2}{a-2b+c}$,

that is, the minimal payoff function is constant for $0 \le P \le Q^*$. The same is true for the other boundary. When

$$Q_0 = \frac{c-b}{a-2b+c} = Q^*$$
, then $\min_{Q_0 \le Q \le Q_1} u_1(1,Q) = \frac{ac-b^2}{a-2b+c}$,

that is, the minimal payoff function is constant for $Q^* \leq P \leq 1$.

With this analysis one can see immediately that when $Q_0 = Q^* = Q_1$, the minimal payoff function is constant for all $P \in [0, 1]$ thus any Ellsberg strategy $[P_0, P_1] \subseteq [0, 1]$ is a best response for player 1.

Assume that $Q_0 > Q^*$, then (since a - b > 0) the minimal payoff function will be strictly increasing and the best response of player 1 is $P_0 = P_1 = 1$. The opposite is true for $Q_1 < Q^*$ and thus the best response is $P_0 = P_1 = 0$. Observe that when $Q_0 < Q^* < Q_1$, the values of both boundary functions drop below Q^* and the function takes its maximum at the kink $P = Q^*$. Therefore player 1's best response in this case is $P_0 = P_1 = Q^*$.

Two cases are still missing. The minimal expected payoff function can be flat exclusively to the left or to the right of Q^* . For all $P \in \left[0, \frac{c-b}{a-2b+c}\right] =$ $[0, Q^*]$, Player 1's utility is constant at $\frac{ac-b^2}{a-2b+c}$ when $Q_1 = \frac{c-b}{a-2b+c} = Q^*$, and it is strictly decreasing for $P > Q^*$. Hence, all $P \le Q^*$ are optimal for player 1. He can thus use any Ellsberg strategy $[P_0, P_1]$ with $P_1 \le Q^*$ as a best reply. Similarly, the payoff is constant for all $P \ge Q^*$ when $Q_0 = Q^*$ (and strictly increasing for $P < Q^*$). This means that player 1's best response to a strategy $[Q_0, Q^*]$ is any strategy $[P_0, P_1] \subseteq [0, Q^*]$, and player 1's best response to a strategy $[Q^*, Q_1]$ where $Q^* \le Q_1 \le 1$ is any strategy $[P_0, P_1] \subseteq [Q^*, 1]$.

We repeat the same analysis for player 2. His minimal expected utility when he plays the mixed strategy Q is

$$\begin{split} \min_{P_0 \leq P \leq P_1} u_2(P,Q) &= \min_{P_0 \leq P \leq P_1} dPQ + eP(1-Q) + e(1-P)Q + f(1-P)(1-Q) \\ &= \min_{P_0 \leq P \leq P_1} P(e-f+Q(d-2e+f)) + eQ + f - fQ \\ &= \begin{cases} P_0(e-f+Q(d-2e+f)) + eQ + f - fQ & \text{if } Q < \frac{f-e}{d-2e+f} \\ \frac{df-e^2}{d-2e+f} & \text{if } Q = \frac{f-e}{d-2e+f} \\ P_1(e-f+Q(d-2e+f)) + eQ + f - fQ & \text{else} \,. \end{cases} \end{split}$$

Note that the payoff function has a fixed value at $Q = \frac{f-e}{d-2e+f}$, which is player 1's Nash equilibrium strategy. Again, as for player 1, depending on P_0 and P_1 the minimal payoff function can have six different forms. We note the two functions that describe the minimal payoff function at the borders Q = 0 and Q = 1:

$$\min_{P_0 \le P \le P_1} u_1(P,0) = P_0(e-f) + f \quad \text{and} \quad \min_{P_0 \le P \le P_1} u_1(P,1) = P_1(d-e) + e.$$

Note that e - f > 0 and d - e < 0, that is, the minimal payoff function is increasing with P_0 in P = 0 and decreasing with P_1 in P = 1. When

$$P_0 = \frac{f-e}{d-2e+f} = P^*$$
, then $\min_{P_0 \le P \le P_1} u_2(P,0) = \frac{df-e^2}{d-2e+f}$,

that is, the minimal payoff function is constant for $0 \le Q \le P^*$. The same is true for the other boundary: When

$$P_1 = \frac{f-e}{d-2e+f} = P^*$$
, then $\min_{P_0 \le P \le P_1} u_2(P,1) = \frac{df-e^2}{d-2e+f}$,

that is, the minimal payoff function is constant for $P^* \leq Q \leq 1$.

Similar to the analysis of player 1 we now get the following best responses of player 2. When $P_0 = P^* = P_1$ player 2 can use any strategy $[Q_0, Q_1] \subseteq$ [0, 1], when $P_0 > P^*$ the best response is $Q_0 = Q_1 = 0$ and when $P_1 < P^*$ then $Q_0 = Q_1 = 1$. When $P_0 < P^* < P_1$ the minimal payoff function takes its maximum at the kink $Q = P^*$ and accordingly player 2's best response is $Q_0 = Q_1 = P^*$.

Finally note that Player 2's utility is constant at $\frac{df-e^2}{d-2e+f}$ for all $Q \in \left[0, \frac{f-e}{d-2e+f}\right] = [0, P^*]$ when $P_0 = \frac{f-e}{d-2e+f} = P^*$, and it is strictly decreasing for $Q > P^*$. Hence, all $Q \leq P^*$ are optimal for player 2. He can thus use any Ellsberg strategy $[Q_0, Q_1]$ with $Q_1 \leq P^*$ as a best reply. Similarly, the payoff is constant for all $Q \geq P^*$ when $P_0 = P^*$ (and strictly increasing for $Q < P^*$). This means that player 2's best response to a strategy $[P^*, P_1]$ is any strategy $[Q_0, Q_1] \subseteq [0, P^*]$, and player 2's best response to a strategy $[P^*, P_1]$ where $P^* \leq P_1 \leq 1$ is any strategy $[Q_0, Q_1] \subseteq [P^*, 1]$.

In Ellsberg equilibrium no player wants to unilaterally deviate from his equilibrium strategy. We analyze in the following which Ellsberg strategies have best responses such that no player wants to deviate.

We assume first that $Q^* < P^*$.

Three Ellsberg strategies can quickly be excluded to be part of an Ellsberg equilibrium. Suppose player 2 plays $[Q_0, Q_1]$ with $Q_0 > Q^*$, then player 1's best response is $P_0 = P_1 = 1$ and since we are looking at a strictly competitive game, player 2 would want to deviate from his original strategy to $Q_0 = Q_1 = 0$. A similar reasoning leads to the result that an Ellsberg strategy $[Q_0, Q_1]$ with $Q_1 < Q^*$ cannot be an equilibrium strategy. Thirdly, suppose player 2 plays $[Q_0, Q_1]$ with $Q_0 < Q < Q_1$, then player 1 would respond with $P_0 = P_1 = Q^*$. Since $Q^* < P^*$ player 2 would deviate from his original strategy to $Q_0 = Q_1 = 1$.

Now suppose player 2 plays $Q_0 = Q_1 = Q^*$, then player 1 can respond with any $[P_0, P_1] \subseteq [0, 1]$. Any choice with $P_0 \ge P^*$, $P_1 < P^*$ or $P_0 < P^* < P_1$ lead to contradictions similar to the cases above. The possibilities that $P_0 < P_1 = P^*$ or $P_0 = P_1 = P^*$ (which are Ellsberg equilibria) are contained in the Ellsberg equilibria that arise in the two remaining cases below.

Suppose player 2 plays $[Q^*, Q_1]$ with $Q^* \leq Q_1 \leq 1$, then if player 1 responds with $[P_0, P_1] = [P^*, P_1]$ with $P^* \leq P_1 \leq 1$ player 2 would play any strategy $[Q_0, Q_1] \subseteq [0, P^*]$ as a best response. Because $Q^* < P^*$, player 2 can choose $[Q_0, Q_1] = [Q^*, Q_1]$ with $Q^* \leq Q_1 \leq P^*$. These strategies are Ellsberg equilibria

$$([P^*, P_1], [Q^*, Q_1])$$
 where $P^* \le P_1 \le 1$ and $Q^* \le Q_1 \le P^*$.

In the case $Q^* < P^*$ this is the only type of Ellsberg equilibrium. Note that the Nash equilibrium is contained in these equilibrium strategies.

When we assume that $P^* < Q^*$ the analysis is very similar. We skip the first four cases and only look at the cases where the minimal payoff function has flat parts. Suppose player 2 plays $[Q_0, Q^*]$ with $0 \le Q_0 \le Q^*$, then if we let player 1 pick $[P_0, P_1] \subseteq [P_0, P^*]$ with $0 \le P_0 \le P^*$, player 2's best response is any subset $[Q_0, Q_1] \subseteq [P^*, 1]$. Again, because $P^* < Q^*$, he can choose $[Q_0, Q_1] = [Q_0, Q^*]$ with $P^* \le Q_0 \le Q^*$ as a best response. Player 1 would not want to deviate and thus these strategies are Ellsberg equilibria

$$([P_0, P^*], [Q_0, Q^*])$$
 where $0 \le P_0 \le P^*$ and $P^* \le Q_0 \le Q^*$.

As before, this is the only type of Ellsberg equilibrium in case $P^* < Q^*$.

Finally let $P^* = Q^*$. Repeat the considerations above having in mind the equality of the Nash equilibrium strategies. Since it was precisely the difference between P^* and Q^* that led to the Ellsberg equilibria in the above cases, we see that no Ellsberg equilibria exist where both players create ambiguity. But, in difference to the above analysis, two types of Ellsberg equilibria with unilateral ambiguity arise that could not be sustained above.

Remember that when player 2 plays $[Q_0, Q_1]$ with $Q_0 < Q^* < Q_1$ then it is optimal for player 1 to respond with $P_0 = P_1 = Q^*$. Since $P^* = Q^*$ these strategies are in equilibrium, even for $Q_0 \leq Q^* \leq Q_1$. One observes that, as long as player 2 makes sure that the mixed Nash equilibrium strategy Q^* is strictly contained in his Ellsberg strategy, player 1 will respond with Q^* and we have Ellsberg equilibria in which player 1 immunizes against the ambiguity of player 2. An analogous type of Ellsberg equilibrium exists for player 2 immunizing against the ambiguity of player 1 by playing $Q_0 = Q_1 = P^*$. Thus we have the following Ellsberg equilibria

$$(Q^*, [Q_0, Q_1])$$
 where $Q_0 \le Q^* \le Q_1$

and

$$([P_0, P_1], P^*)$$
 where $P_0 \le P^* \le P_1$.

These equilibria do not exist in the non-symmetric case $P^* \neq Q^*$ since the immunization strategies are in general not equilibrium strategies. Due to the assumptions on the payoffs in this proposition, the immunization strategy equals the Nash equilibrium strategy of the opponent. Thereby the immunization strategy is an equilibrium strategy only when $P^* = Q^*$.

- **Remark 1.** 1. In Proposition 5 we restrict to the case with (U, D) and (L, R) giving the same payoffs (b, e) for both players. Of course the Ellsberg equilibria of competitive games with more general payoffs can easily be calculated, the calculations are available upon demand. The more general competitive game yields two more types of Ellsberg equilibria. The nice feature of our restriction is that players use the mixed Nash equilibrium strategy of their respective opponent as their immunization strategy.
 - 2. Observe the asymmetry in the Ellsberg equilibria in the preceding proposition: no matter if $P^* < Q^*$ or $Q^* < P^*$, always it is player 2 who creates ambiguity between the Nash equilibrium strategies, player 1 never does so. This is due to the assumptions on the payoffs. If we assume that a, c < b and d, f > e player 1 will play between P^* and Q^* .
 - 3. Note that Proposition 5 holds likewise for zero sum games, but due to the assumptions on the payoffs the proposition restricts to zero sum games where $P^* = Q^*$. This can be seen easily by setting d := -a, f =: -c and e := -b, then $P^* = \frac{f-e}{d-2e+f} = \frac{c-b}{a-2b+c} = Q^*$. Therefore in the Ellsberg equilibria of two-person 2×2 zero sum games under the assumptions of Proposition 5 only one player creates ambiguity. This also holds for general two-person 2×2 zero sum games, although in those games in general $P^* \neq Q^*$. The result then hinges on the fact that the immunization strategy of player i is always equal to his own Nash equilibrium strategy.

References

- ARMBRUSTER, W., AND W. BOEGE (1979): "Bayesian Game Theory," in *Game Theory and Related Topics*, ed. by O. Moeschin, and D. Pallaschke. North-Holland, Amsterdam.
- AUMANN, R. (1974): "Subjectivity and correlation in randomized strategies," *Journal of Mathematical Economics*, 1(1), 67–96.

(1987): "Correlated equilibrium as an expression of Bayesian rationality," *Econometrica: Journal of the Econometric Society*, 55(1), 1–18.

- AUMANN, R., AND A. BRANDENBURGER (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161–1180.
- AZRIELI, Y., AND R. TEPER (2011): "Uncertainty aversion and equilibrium existence in games with incomplete information," *Games and Economic Behavior*.
- BADE, S. (2011): "Ambiguous act equilibria," Games and Economic Behavior, 71(2), 246–260.
- BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARI-NACCI (2011): "Selfconfirming Equilibrium and Uncertainty," Working Paper, Università Bocconi.
- BATTIGALLI, P., AND D. GUAITOLI (1988): "Conjectural Equilibria and Rationalizability in a Macroeconomic Game with Incomplete Information," (Extended Abstract) Università Bocconi.
- CAMERER, C., AND M. WEBER (1992): "Recent developments in modeling preferences: Uncertainty and ambiguity," *Journal of risk and uncertainty*, 5(4), 325–370.
- DANG, T. (2009): "Gaming or Guessing: Mixing and best-responding in Matching Pennies," *mimeo*.
- DOW, J., AND C. WERLANG (1994): "Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction," *Journal of Economic Theory*, 64, 305–324.

- EICHBERGER, J., AND D. KELSEY (2000): "Non-Additive Beliefs and Strategic Equilibria^{*} 1," *Games and Economic Behavior*, 30(2), 183–215.
- EICHBERGER, J., D. KELSEY, AND B. SCHIPPER (2009): "Ambiguity and social interaction," Oxford Economic Papers, 61, 355–379.
- EISENBERG, E. (1984): "Ambiguity as strategy in organizational communication," *Communication monographs*, 51(3), 227–242.
- ELLSBERG, D. (1961): "Risk, ambiguity, and the Savage axioms," *The Quarterly Journal of Economics*, pp. 643–669.
- EPSTEIN, L. (1997): "Preference, rationalizability and equilibrium," *Journal* of Economic Theory, 73(1), 1–29.
- EPSTEIN, L., AND M. SCHNEIDER (2003): "IID: independently and indistinguishably distributed," *Journal of Economic Theory*, 113(1), 32–50.
- ETNER, J., M. JELEVA, AND J. TALLON (2012): "Decision theory under ambiguity," *Journal of Economic Surveys*, 26(2), 234–270.
- FUDENBERG, D., AND D. LEVINE (1993): "Self-confirming equilibrium," Econometrica, 61(3), 523–545.
- FUDENBERG, D., AND J. TIROLE (1991): Game theory. MIT Press.
- GAJDOS, T., T. HAYASHI, J. TALLON, AND J. VERGNAUD (2008): "Attitude toward imprecise information," *Journal of Economic Theory*, 140(1), 27–65.
- GIGERENZER, G. (2007): Gut feelings: The intelligence of the unconscious. Viking Pr.
- GILBOA, I., AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.
- GIRAUD, R. (2006): "Objective imprecise probabilistic information, second order beliefs and ambiguity aversion: an axiomatization," *mimeo*.
- GIRAUD, R., AND J. TALLON (2011): "Are beliefs a matter of taste? A case for objective imprecise information," *Theory and decision*, 71(1), 23–31.

- GOEREE, J., AND C. HOLT (2001): "Ten little treasures of game theory and ten intuitive contradictions," *American Economic Review*, 91(5), 1402– 1422.
- GREENBERG, J. (2000): "The right to remain silent," *Theory and Decision*, 48(2), 193–204.
- HARSANYI, J. (1967): "Games with incomplete information played by "Bayesian" players, I-III," Mana, 14(3), 159–182.
- JAFFRAY, J.-Y. (1989): "Généralisation du critère de l'utilité à l'incertain régulier," Operations Research/Recherche Opérationnelle, 23(3), 237267.
- KALAI, E., AND E. LEHRER (1995): "Subjective games and equilibria," *Games and Economic Behavior*, 8(1), 123–163.
- KISSINGER, H. (1982): Years of upheaval. Weidenfeld & Nicolson.
- KLIBANOFF, P. (1996): "Uncertainty, decision, and normal form games," Manuscript, MIT.
- LEHRER, E. (2008): "Partially-specified probabilities: decisions and games," *mimeo*.
- LEWIS, D. (1969): Convention: A Philosophical Study. Harvard University Press.
- LO, K. (1996): "Equilibrium in beliefs under uncertainty," Journal of Economic Theory, 71(2), 443–484.
- (1999): "Extensive Form Games with Uncertainty Averse Players" 1," Games and Economic Behavior, 28(2), 256–270.
- (2009): "Correlated Nash equilibrium," Journal of Economic Theory, 144(2), 722–743.
- MACCHERONI, F., AND M. MARINACCI (2005): "A Strong Law of Large Numbers for Capacities," Annals of Applied Probability, 33, 1171–1178.
- MARINACCI, M. (2000): "Ambiguous games," Games and Economic Behavior, 31(2), 191–219.

- OLSZEWSKI, W. (2007): "Preferences over sets of lotteries," The Review of Economic Studies, 74(2), 567.
- PENG, S. (2007): "Law of large numbers and central limit theorem under nonlinear expectations," Arxiv preprint math/0702358.
- PULFORD, B. (2009): "Is luck on my side? Optimism, pessimism, and ambiguity aversion," The Quarterly Journal of Experimental Psychology, 62(6), 1079–1087.
- RENY, P., AND A. ROBSON (2004): "Reinterpreting mixed strategy equilibria: a unification of the classical and Bayesian views," *Games and Economic Behavior*, 48(2), 355–384.
- SCHELLING, T. C. (1960): *The Strategy of Conflict.* Harvard University Press.
- SCHMEIDLER, D. (1989): "Subjective probability and expected utility without additivity," *Econometrica*, pp. 571–587.
- STINCHCOMBE, M. (2007): "Consequentialist Ambiguity: Choice under Ambiguity as Choice Between Sets of Probabilities," *mimeo*.
- TAN, T., AND S. WERLANG (1988): "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory*, 45, 370–391.
- VON NEUMANN, J., AND O. MORGENSTERN (1953): Theory of Games and Economic Behavior. Princeton University Press (Princeton).