

Riedel, Frank

Working Paper

Uncertain acts in games

Center for Mathematical Economics Working Papers, No. 571

Provided in Cooperation with:

Center for Mathematical Economics (IMW), Bielefeld University

Suggested Citation: Riedel, Frank (2017) : Uncertain acts in games, Center for Mathematical Economics Working Papers, No. 571, Bielefeld University, Center for Mathematical Economics (IMW), Bielefeld,
<https://nbn-resolving.de/urn:nbn:de:0070-pub-29091716>

This Version is available at:

<https://hdl.handle.net/10419/175872>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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


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

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

March 2017

Uncertain Acts in Games

Frank Riedel



Uncertain Acts in Games

Frank Riedel*

March 15, 2017

Abstract

This text reviews a recent approach to modeling “radically uncertain” behavior in strategic interactions. By rigorously rooting the approach in decision theory, we provide a foundation for applications of Knightian uncertainty in mechanism design, principal agent and moral hazard models. We discuss critical assessments and provide alternative interpretations of the new equilibria in terms of equilibrium in beliefs, and as a boundedly rational equilibrium in the sense of a population equilibrium. We also discuss the purification of equilibria in the spirit of Harsanyi.

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

JEL subject classification: C72, D81

1 Introduction

In games of conflict, players have an incentive to conceal their behavior. A player has to be unpredictable, in some sense, since otherwise their opponent(s) can exploit the player. To find a minmax-value for zero-sum games, v. Neumann (1928) introduces mixed strategies, i.e. objective random devices with a known probability distribution. While a probabilistic device is perfectly natural for a mathematically trained person, one might feel that other, more “radically uncertain” approaches to playing a game might appear interesting as well. In this paper, I review and extend a recent attempt (in

*Center for Mathematical Economics, Bielefeld University and Department of Economic and Financial Sciences, University of Johannesburg. The author thanks Quentin Couanau and Jean-Marc Tallon for helpful comments. Financial support through the German Research Foundation Grant Ri-1128-6-1 is gratefully acknowledged.

Riedel and Sass (2014)) to study the outcomes in games when such radical uncertainty is allowed.

This paper extends the basic approach in Riedel and Sass (2014) and provides the conceptual and decision-theoretic foundations. It shows that players might profit from playing in an uncertain way, rather than a probabilistic one. We work with one main example to illustrate the new findings because we want to focus on the conceptual foundations of the new approach rather than on the generality of results that are the subject of ongoing work.

When extending the set of possible strategies, it is necessary to extend the utility functions. We provide a sound decision-theoretic foundation for our approach by rooting the extended utility function in the setup of Gajdos, Hayashi, Tallon, and Vergnaud (2008) who provide an axiomatic analysis of preferences over acts with imprecise probabilistic information. We also detail how to use the smooth model of Klibanoff, Marinacci, and Mukerji (2005) by using the appropriate extension of this class of utility functionals to uncertain acts with imprecise probabilistic information provided by Giraud (2014).

Of course, the “objective” interpretation by von Neumann has been criticized. Schelling (1960) has pointed out that there is no incentive to conceal one’s own behavior in common interest games. Hence, mixed equilibria in such games require a different justification. Harsanyi (1973) has shown that it is possible to interpret mixed Nash equilibria as pure strategy equilibria in incomplete information games with private information in which the players’ payoffs are randomly disturbed. If the perturbation is sufficiently small, agents’ equilibrium behavior approximates the mixed strategy equilibrium for an outside observer.

In contrast to what some people might have expected, such a purification is also possible in our leading example. The new equilibria can be interpreted as pure strategy equilibria in an incomplete information game with private information where players’ payoffs are disturbed in a Knightian way, i.e. the probability distribution of the perturbation is not known.

We discuss some objections to the new approach, in particular how one should interpret these “new” strategies in real life situations and, from a rather dogmatic side, whether players “have to” be Bayesian, or maybe not. We also describe different interpretations of the new equilibria in terms of equilibrium in beliefs (Lo (1996)) or Nash’s boundedly rational population interpretation.

The last section reviews recent interesting applications of the new approach to mechanism design, signaling games, and other economically relevant models, and contains a selective literature review¹.

¹The literature on Knightian uncertainty in games has grown at a rapid speed in

2 Profiting from Uncertainty

I assume that you, dear reader, are an economist, or a game theorist (the latter not necessarily being a subset of the former); as such, you are probably so familiar with the idea of randomizing that it is necessary for you to free yourself from this familiarity at the beginning of our reflections. It would thus be helpful to remember a time when you did not know about the formal version of a mixed strategy in a game as introduced by John von Neumann in his first text on “Gesellschaftsspiele” in 1928. It is clear to anyone who has played parlor games that being unpredictable is important in some situations. If the opponents are able to anticipate your next move, you lose in poker and, in fact, in all competitive games. I had certain doubts, however, when I first studied the minmax theorem; while it is perfectly natural for a mathematically trained person to introduce a probabilistic device, a critical mind immediately doubts if real players, humans, mix *probabilistically*. We know, of course, that humans do have problems with these tasks; even a well trained person is barely able to produce a sequence of zeroes and ones in which the probability of a one is 71.2 %. However, if we examine the zero-sum game in Table 1, then this is exactly what is required in Nash equilibrium.

	L	R
U	$-18, 18$	$18, -18$
D	$71, -71$	$-18, 18$

Table 1: Payoff matrix for a game in which player 1 needs to play U with probability 71.2 % in equilibrium.

On the other hand, for most real world situations, it is not necessary to

recent years. We mention here the approaches that use uncertain actions or strategies. Bade (2011) studies two player games using uncertain strategies in an Anscombe–Aumann setting where players have subjective uncertainty aversion and do not use the imprecise probabilistic information contained in strategies. She focuses on the fact that in two player games, the support of equilibrium actions is the same as in Nash equilibrium. Grant, Meneghel, and Tourky (2016) introduce uncertain actions without allowing randomization in the classic sense. Then the issue of existence of an equilibrium arises as strategy sets are no longer convex. In a rich Savage-like setting, they are able to prove existence of an equilibrium. Stauber (2017) develops an interesting different interpretation of ambiguity about actions as coming from irrational behavior of players.

Lo (1996), Marinacci (2000), and Eichberger and Kelsey (2011) are examples of approaches to games in which players use pure strategies, but players hold ambiguous beliefs about the other players’ actions. We would like to refer to Riedel and Sass (2014) for a review on belief-based approaches to ambiguity.

randomize perfectly; some degree of uncertainty is sufficient. It is thus important to distinguish between the mathematical model and the real world application. In the mathematical model, agents are fully rational and compute perfectly; for the real world application, the model abstracts away from many details anyway, and we are pleased if the actual behavior matches the theoretical prediction to a sufficient degree. It has been documented that humans – or let us say more accurately, students in the laboratory – tend to have problems playing the mixed Nash equilibrium exactly in games like the one above (see Goeree and Holt (2001), for example).

While some skepticism towards perfect randomization for real players is well deserved, we should not go too far by considering only pure strategies in games. We do not want to call the necessity to be unpredictable in games of conflict into question. Every sportsman knows that you lose if the opponent is able to foresee your next action, be it in tennis, soccer, or baseball. In the words of the German soccer player Lukas Podolski (who might well be the modern counterpart to America’s Yogi Berra): “It is best when I do not know what I do.” Although Podolski was speaking about penalties in this quotation, the same idea holds true whenever there is an incentive to disguise a planned move in any strategic interaction.

Let us now return to the very beginnings of game theory, and in fact, to John von Neumann himself, and re-think why and how he introduced mixed strategies. Have a look again at the zero-sum game in Table 1. Whenever you play a pure strategy like U for the first player, your opponent has the strict incentive to play the best reply, and you will get a payoff of -18 causing the opponent to “win” the game. The game does not have a unique value in von Neumann’s sense; the lower value (maxmin) of -18 is strictly smaller than the upper value (minmax) of 71 . Things change, of course, if we introduce random strategies and expected utility. When both players randomize (with the probability 0.712 for U resp. 0.288 for L) and evaluate the payoffs by computing the expected payoff, the lower and the upper value coincide at -7.632 .

In their later book (1953), von Neumann and Morgenstern gave the following justification for the introduction of mixed strategies: suppose you have to write down and commit to your planned action before the game is carried out. There is now a risk that your opponent, in modern words, will hack your computer and finds out what you plan to do. If such a risk exists, it is best to commit to a random action because the spy will not learn anything useful in that case.

We now aim to radicalize this approach of being unpredictable². In some

²without being too radical, though, as you are going to see.

situations, it is advantageous to be even more unpredictable than a random device with well specified probabilities. Note that the random devices that von Neumann introduced provide very specific information about the probability distribution of the actions. This information is used by agents to determine their expected payoffs. In the following example (see Table 2), we reveal that being imprecise about probabilities when players are averse to such radical uncertainty can pay off.

	<i>L</i>	<i>R</i>	
<i>T</i>	0, 8, 0	3, 3, 3	
<i>B</i>	1, 1, 1	0, 0, 0	
	<i>l</i>		

	<i>L</i>	<i>R</i>
<i>T</i>	0, 0, 0	3, 3, 3
<i>B</i>	1, 1, 1	8, 0, 0
	<i>r</i>	

Table 2: Player 1 chooses the row, player 2 the column, and player 3 the matrix.

In this game, Player 3 is indifferent regarding the choice between his actions *l* and *r* since the payoff is always the same for him for the given strategy choices of Players 1 and 2. The best payoff for Player would be 3. Player 3 might want to induce this outcome by creating some uncertainty for the other players concerning which matrix game, left or right, is played. However, this does not work with a classic mixed strategy: in every Nash equilibrium, Players 1 and 2 play *B* and *L*, and each player obtains the payoff 1. To see this, note that Player 2 has the dominant action *L* in the left matrix game; in the right matrix game, Player 1 has the dominant action *B*. If Player 3 plays *l*, Player 2 thus plays *L*, and Player 1 best replies with *B*. If Player 3 plays *r*, Player 1 plays *B*, and Player 2 best replies with *L*. Randomization does not get us away from this outcome: if the probability of *l* is greater than $3/8$, Player 2 has the unique best reply *L*, to which Player 1 again replies *B*; if *p* is smaller than $5/8$, Player 1 has the unique best reply *B*, to which Player 2 replies *L*. The payoffs are always 1 for each player.

Now let us suppose that Player 3 can behave in a more radically uncertain way; his action is uncertain in the sense that Players 1 and 2 do not even have an idea of the probability for *l* or *r*. We explain how such radical uncertainty can be modeled in the next section; for the moment, though, it is sufficient to think of a completely uncertain action for which we do not know the probabilities (a horse race, in the language of decision theory, rather than a lottery, or, if you allow, Lukas Podolski).

Now suppose for the sake of the example that the players do not like such uncertainty; they assume that Player 3 is acting against them and thus consider the worst possible outcome.

We claim that the strategy profile (T, R, H) (here, 'H' stands for “horse race”) is an equilibrium. Indeed, by deviating to B, Player 1 can receive 8 or 0; as he considers the worst case, he assumes that Player 3 might play l, resulting in a payoff of 0. The same reasoning applies to Player 2 who might receive 0 or 8 by deviating to L; as she is completely uncertain, she considers the worst case, namely Player 3 playing r, to be relevant, and thus does not deviate. As 3 is the optimal payoff for Player 3, no one has an incentive to deviate, and we have a new equilibrium.

Radical uncertainty in strategies in a game where players are uncertainty-averse can lead to new equilibria. I think that this example is sufficiently interesting to start a program to discover the consequences of extending game theory to more general notions of mixed strategies. A first step was made in Riedel and Sass (2014). We will generalize the approach and root it firmly in decision theory in this paper.

3 Modeling Ambiguous Acts

We aim to formally model the idea of behaving in an uncertain, unpredictable way without relying on precise probabilistic devices. In decision theory, it is common to distinguish between *horse races* and *lotteries*. In the latter, probabilities are exactly known (as in a good roulette wheel), in the former, few or even no information about the probability distribution is available.

Let us consider a game in normal form $G = \langle N, (S_i, u_i)_{i \in N} \rangle$ where $N = \{1, \dots, n\}$ is the set of players, $S_i, i = 1, \dots, n$ are the finite strategy sets, $S = \prod_{i=1}^n S_i$ is the set of pure strategy profiles, and players' payoffs are given by functions

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

3.1 Ambiguous Acts

We model an uncertain strategy over the pure strategy set S_i of some player i . In principle, the device we are looking for is well known in decision theory: we use the urns of unknown composition discussed by Daniel Ellsberg (1961). In his famous experiments, Ellsberg had subjects choose between bets on different urns; in some urns, the composition of, e.g., red and blue balls was exactly known, whereas in other urns, only bounds for the number of red and blue balls were given.

We use this ingenious device to model our uncertain strategies. It is important to note an important feature of Ellsberg experiments that is often overlooked. An Ellsberg bet consists of an act $f(\omega)$ (in which ω is the outcome

of the draw and $f(\omega)$ is the associated payoff) *and* of some information about the composition of the urn. For example, the urn is known to contain one hundred blue and red balls at least thirty of which are red. An Ellsberg urn is thus described by a pair (f, \mathcal{P}) where f is the act and \mathcal{P} the information about the possible probability distributions.

Definition 1. *Let $G = \langle N, (S_i, u_i)_{i \in N} \rangle$ be a normal form game. Denote by ΔS_i the set of all probability vectors over S_i (mixed strategies in the sense of von Neumann).*

A strategy for player i in game G consists of a finite state space Ω , an act $f_i : \Omega \rightarrow \Delta S_i$, and a set of probability distributions \mathcal{P} over Ω .

Let us illustrate this definition by the behavior of Player 3 in the above example. If Player 3 aims to be completely ambiguous, he can take an Ellsberg urn fore which the composition is completely unknown. Let $\Omega = \{0, 1\}$, and set $\mathcal{P} = \Delta\Omega$, the set of all probabilities over Ω . The act is then modeled as $f_3(0) = l, f_3(1) = r$.

As stated above, we do not discard the usual mixed strategies. Note that the act f_i maps into the set ΔS_i of mixed strategies in the classic sense. If a player does not want to behave in an ambiguous way, he can simply choose to use a random device.

In the next step, we reduce the complexity that we have allowed for so far by assuming that players only care about the ambiguity that is induced over strategies $S_i = \{a_{i,1}, \dots, a_{i,n_i}\}$. They do not care about the specific properties of the Ellsberg experiment or how the act f depends on the outcome of the Ellsberg experiment; rather, they care about the game, and the Knightian uncertainty regarding the strategies played.

If Player i uses a state space Ω , an act f_i , and a set of probability distributions \mathcal{P} over Ω , we obtain an induced distribution q over S_i for every probability $p \in \mathcal{P}$ in the following way. Denote by $f_i(\omega)(a_{ik})$ the objective probability that a_{ik} is chosen in state ω . Then

$$q_k^{\Omega, p, f_i} = \sum_{\omega \in \Omega} f_i(\omega)(a_{ik})$$

is the induced probability of action a_{ik} . The uncertain strategy given by the triple $(\Omega, \mathcal{P}, f_i)$ thus leads to a set of probability distributions

$$\mathcal{Q} = \{q^{\Omega, p, f_i} : p \in \mathcal{P}\}$$

over S_i .

Definition 2. *A reduced uncertain act (or reduced strategy) for player i in the game $G = \langle N, (S_i, u_i)_{i \in N} \rangle$ consists of a nonempty subset $\mathcal{Q} \subset \Delta S_i$.*

3.2 Independence

We consider *noncooperative* games and thus maintain the standing assumption that players act independently. Now, Knightian uncertainty sets new challenges for some concepts from probability theory that we are used to taking for granted. Especially in dynamic games, when players need to update their beliefs after observing previous actions, independence and the related issue of dynamic consistency become difficult issues (see, e.g., the discussion of Kuhn’s theorem in Muraviev, Riedel, and Sass (2017)).

For normal form games, we model independent Ellsberg experiments by requiring independence uniformly over all priors. Let us illustrate this approach for a two-player game. Given the reduced uncertain acts \mathcal{Q}_1 and \mathcal{Q}_2 , and possible (mixed) actions $q_1 \in \mathcal{Q}_1$ and $q_2 \in \mathcal{Q}_2$, both players consider the profile of induced actions to be independent. Therefore, the only possible joint probability distribution for evaluating the payoffs is the product measure $q_1 \otimes q_2$ on $S_1 \times S_2$.

In general, given the reduced strategy profile $(\mathcal{Q}_1, \dots, \mathcal{Q}_n)$ in a game, the corresponding induced independent profile over the set of pure strategy profiles S is

$$\otimes_{i=1}^n \mathcal{Q}_i = \{q_1 \otimes \dots \otimes q_n : q_i \in \mathcal{Q}_i, i = 1, \dots, n\} .$$

3.3 Extending the Payoff Functions

Having extended the set of possible strategies, our next task is to extend the payoff functions to this new domain. We want to do so parsimoniously and are particularly interested in keeping the received theory of normal form games as long as players use pure strategies or random devices. We thus keep the independence axiom for objective probabilities. Moreover, there should be no incentive to introduce ambiguity unilaterally. We thus assume that players are averse to such Knightian uncertainty.

3.3.1 Imprecision Aversion

In decision theory, uncertainty-averse preferences have been discussed in great detail; however, the authors usually consider preferences over acts f only whereas we are interested here in pairs (f, \mathcal{P}) of acts and (imprecise) probabilistic information, or, in the reduced form, subsets \mathcal{Q} of ΔS_i .

Gajdos, Hayashi, Tallon, and Vergnaud (2008) represent a notable exception as these authors consider preferences over acts and imprecise probabilistic information. Their paper is certainly the best axiomatic treatment of the Ellsberg “paradox” in decision theory because it is the only one that

models the imprecise probabilistic information that is given objectively in the Ellsberg experiments.

Denote for a mixed strategy profile $q = (q_1, \dots, q_n) \in \prod_{j=1}^n \Delta S_j$ the expected payoff of player i by

$$E^q u_i = E^{q_1 \otimes \dots \otimes q_n} u_i,$$

where, with a slight abuse of notation, we identify q and the product of the single probabilities $q_1 \otimes \dots \otimes q_n$.

Confronted with the set of independent profiles $\mathcal{Q} = \prod_{i=1}^n \mathcal{Q}_i$, players need to evaluate their payoffs. As we want to extend the model such that classic game theory is included, we use the usual expected payoff $E^q u_i$ for a single strategy profile. If the set $\mathcal{Q} = \prod_{i=1}^n \mathcal{Q}_i$ is not a singleton, the players are confronted with Knightian uncertainty in which some probabilistic information is given. Gajdos, Hayashi, Tallon, and Vergnaud (2008) model uncertainty aversion by a function $\phi_i(\mathcal{Q})$ that selects a set of priors from the given set of priors; this selection is used to evaluate the payoffs, in conjunction with the uncertainty-averse axioms of Gilboa and Schmeidler (1989).

The natural extension of the utility function along the lines of Gajdos, Hayashi, Tallon, and Vergnaud (2008) thus reads as follows:

$$U_i(\mathcal{Q}_1, \dots, \mathcal{Q}_n) = \min_{q \in \phi_i(\mathcal{Q})} E^q u_i,$$

where ϕ_i is a correspondence that selects a set of subjective priors $\phi_i(\mathcal{Q})$ from the given imprecise probabilistic information \mathcal{Q} .

Given that the players are known to play independently, the selection ϕ_i should respect this independence, in general. One might thus want to assume that we can write

$$\phi_i(\mathcal{Q}) = \prod_{j=1}^n \phi_{ij}(\mathcal{Q}_j)$$

where now ϕ_{ij} selects a suitable subjective set of priors for Player i for a set of priors \mathcal{Q}_j among the strategies ΔS_j of Player j . The functions ϕ_{ij} reflect Player i 's ambiguity (or imprecision) aversion, see Gajdos, Hayashi, Tallon, and Vergnaud (2008). Moreover, the selected set of priors should not contradict the information contained in the set \mathcal{Q}_j , i.e. for the supports of the respective sets we have

$$\text{supp } \phi_{ij}(\mathcal{Q}_j) \subset \text{supp } \mathcal{Q}_j.$$

When the information \mathcal{Q}_j is given, a player should not select priors placing mass outside the states which are supported by \mathcal{Q}_j . In case a singleton

is played, $\mathcal{Q}_j = \{q_j\}$ for some $q_j \in \Delta S_j$, it makes sense to require that ϕ_{ij} selects exactly this singleton: $\phi_{ij}(\{q_j\}) = \{q_j\}$. We call such correspondences precise; they are characterized by Axiom 7 in Gajdos, Hayashi, Tallon, and Vergnaud (2008), Reduction under Precise Information.

Definition 3. Let $G = \langle N, (S_i, u_i)_{i \in N} \rangle$ be a game in normal form. Let D_i be the set of nonempty subsets of ΔS_i for $i = 1, \dots, n$. The extended game with ambiguity-averse players consists of G and precise correspondences

$$\phi_{ij} : D_j \rightarrow \rightarrow D_j$$

with

$$\text{supp } \phi_{ij}(\mathcal{Q}_j) \subset \text{supp } \mathcal{Q}_j$$

for all $\mathcal{Q}_j \in D_j$ and $i, j \in N$. The payoff of player i for a reduced strategy profile $\mathcal{Q} = (\mathcal{Q}_1, \dots, \mathcal{Q}_n)$ is

$$U_i(\mathcal{Q}) = \min_{\substack{q_1 \in \phi_{i1}(\mathcal{Q}_1) \\ \vdots \\ q_n \in \phi_{in}(\mathcal{Q}_n)}} E^{q_1 \otimes \dots \otimes q_n} u_i. \quad (1)$$

We call $(G, (\phi_{ij})_{i,j \in N})$ the extension of G .

Note that this is a parsimonious extension of classic game theory: first, we still include the possibility of ambiguity-neutral agents by always allowing a player to select a single prior when confronted objectively with a set of priors. In particular, Bayesian players i who use a second-order prior μ^{ij} over ΔS_j and who use the single prior $\phi_{ij}(\mathcal{Q}_j) = \int_{\mathcal{Q}_j} q \mu(dq | \mathcal{Q}_j)$ when confronted with the probabilistic information $\mathcal{P} \subset \Delta S_i$ are included. Second, we do not make it easy for players to use ambiguous acts; in fact, due to ambiguity aversion, a player has no incentive to play ambiguously unilaterally as this can only decrease his or her payoff. If ambiguous acts do play a role, it thus has to be an interaction effect³.

3.3.2 Smooth Ambiguity Model

The smooth model proposed by Klibanoff, Marinacci, and Mukerji (2005) is another popular approach to modeling ambiguity-averse behavior. It combines a second-order prior μ over the unknown probability distributions with a concave real function ψ which allows for modeling ambiguity-aversion in

³With the exception of cases where the players are just indifferent about any action they use.

the same way as the Bernoulli utility function allows for modeling risk aversion.

I propose a natural extension of the smooth model to our setting in which we do not have preferences over acts but rather over pairs (f, \mathcal{P}) , or, as we discussed above, over sets of probabilities \mathcal{Q} . Giraud (2014) provides an axiomatic foundation for smooth ambiguity-averse utility functions under imprecise probabilistic information; in contrast to his approach, I would like to add an updating procedure on the second-order beliefs which seems natural given that the players obtain new, yet imprecise information about the possible actions being played.

Fix a player i . Let μ_{ij} be a second order prior for player i over the mixed strategies ΔS_j of player j . Given the imprecise probabilistic information \mathcal{Q}_j for strategies chosen by player j , player i uses the updated second order prior $\mu_{ij}(\cdot|\mathcal{Q}_j)$. If ψ_i denotes the ambiguity index of player i , the natural version of the smooth ambiguity utility function reads as

$$U_i(\mathcal{Q}_1, \dots, \mathcal{Q}_n) = \int_{\Delta S_1} \dots \int_{\Delta S_n} \psi_i(E^{q_1 \otimes \dots \otimes q_n} u_i) \mu_{i1}(dq_1|\mathcal{Q}_1) \dots \mu_{in}(dq_n|\mathcal{Q}_n),$$

3.3.3 Examples

We are now able to formally describe the new type of equilibrium that we found in the three-player game in Table 2. In this game, having a sufficient degree of ambiguity aversion is enough to sustain the new equilibrium in which all players obtain the efficient payoff 3. As long as the subjective ambiguity of players 1 and 2 derived from the actual strategy of player 3 contains at least the probabilities $3/8$ and $5/8$, we can support the efficient outcome as an equilibrium.

We identify the reduced ambiguous act for player 3 with an interval $[a, b] \subset [0, 1]$ for the probability of playing l .

Theorem 1. *Let G be the three player game described in Table 2 and let $(G, (\phi_{ij})_{i,j \in \{1,2,3\}})$ be its extension.*

The outcome $(3, 3, 3)$ is an equilibrium outcome whenever players 1 and 2 play T and R resp. (i.e. \mathcal{Q}_1 and \mathcal{Q}_2 consist of the singleton pure strategy $(1, 0)$), player 3 plays $\mathcal{Q}_3 = [a, b] \subset [0, 1]$ and we have

$$\phi_{i3}([a, b]) \supset \left[\frac{3}{8}, \frac{5}{8} \right].$$

The above result reveals that the efficient equilibrium is very robust: as long as a critical degree of ambiguity aversion is known to exist, the

efficient equilibrium outcome is possible. This has an important consequence for applications: one would not necessarily trust an outcome that relies on players' being able to exactly produce a certain Ellsberg urn with given probability bounds. For the game above, however, it is sufficient that an adequate amount of ambiguity is created. In this sense, such equilibria are easier to play in practice than the classic mixed strategy equilibria for which we rely on players' ability to randomize exactly.

We also want to point out that our extension does not lead to arbitrary equilibria. As a first example, we reveal that strictly dominated strategies remain strictly dominated, and we thus do not obtain new equilibria in the extended game.

Theorem 2. *Let G be a normal form game that can be solved by iterated deletion of strictly dominated strategies. Let $q = (q_1, \dots, q_n)$ be the unique Nash equilibrium of G . The extended game $(G, (\phi_{ij})_{i,j \in N})$ has also the unique equilibrium q .*

As we pointed out in Riedel and Sass (2014), in zero-sum games, essentially no new equilibria emerge.

Theorem 3. *Let G be the two-player zero-sum game described in Table 1 and let $(G, (\phi_{ij})_{i,j \in \{1,2,3\}})$ be its extension.*

At least one of the players plays the singleton maxmin strategy in every equilibrium of the extended game.

Given that one of the players plays the maxmin strategy (which is equal to the Nash equilibrium strategy in this zero sum game), the other player is indifferent and can thus play ambiguous acts. However, the ambiguous acts have no role to play in equilibrium as they do not affect the best reply structure of the opponent.

4 Purification

One of the first objections to this extension of game theory I encountered came from Faruk Gul during my sabbatical at Princeton university, which is where von Neumann lived when he and Oskar Morgenstern wrote *Theory of Games and Economic Behavior*. Faruk said: "I do not like this approach because you cannot do 'Harsanyi'." In this section, I am going to show that it is indeed possible to purify the new equilibrium in our leading example in the spirit of Harsanyi (1973).

In his famous paper, John Harsanyi begins by discussing a peculiar property of mixed Nash equilibria. Players are indifferent to all actions they take

in equilibrium; in a two player game, for example, the probabilities by one player are chosen in such a way that the *other* player is indifferent. Harsanyi thus asks why a player should act like this given that he or she has no incentive to do so. Players have to take the payoffs of the other player into account when computing their Nash equilibrium strategy. As a consequence, Nash equilibrium is (at times) difficult to play. In fact, one reason why we observe deviations from Nash equilibrium in laboratory experiments might be due to this difficulty (compare the Constant Sum Experiment and its modifications in Goeree and Holt (2001) and our discussion in Riedel and Sass (2014)).

Harsanyi comes up with the following alternative interpretation of mixed Nash equilibria. Suppose that players' payoffs are randomly disturbed by an exogenous device, independently for each player, and that players' payoffs are private information. In this case, almost every player uses a pure strategy upon observing his payoffs. Players are thus no longer indifferent amongst actions. To an outside observer, the actions look random as he does not know the private perturbation of the payoffs. When the perturbation is sufficiently small, the probabilities in the associated Bayes–Nash equilibria are close to the probabilities in the mixed Nash equilibrium. The mixed equilibrium is thus purified.

In the same spirit, I am now going to show that one can purify the new equilibrium in our three-player game above. We introduce exogenous Knightian uncertainty over payoffs. For simplicity, we just perturb the payoffs of Player 3 who chooses the matrix game being played.

We change the payoffs as follows. When Player 3 plays l , a private per-

	L	R	
T	0, 8, s	3, 3, 3 + s	
B	1, 1, 1 + s	0, 0, s	
	l		

	L	R
T	0, 0, 0	3, 3, 3
B	1, 1, 1	8, 0, 0
	r	

Table 3: Player 1 chooses the row, player 2 the column, and player 3 the matrix.

turbation $s \in [-1, 1]$ is added to his former payoff. Let us assume, again for simplicity, that all distributions for s over $[-1, 1]$ are possible; players have thus complete Knightian uncertainty over the exogenous perturbation.

An equilibrium (in the classic sense of a private information game) consists then of strategies $p \in [0, 1]$ and $q \in [0, 1]$ for Player 1 to play T and for Player 2 to play L , and probabilities $r(s) \in [0, 1]$ for each type of Player 3 to play l . As usual, $r(s)$ has to be measurable in the private information s in order to have well-defined ex ante payoffs.

We claim that $p = 1$ (T is played), $q = 0$ (R is played), and $r(s) = 1$ if $s \geq 0$ and $r(s) = 0$ if $s < 0$ (the types with a positive perturbation play l , the others play r), is a classic equilibrium in the perturbed game that purifies the equilibrium in the extended game.

This is easy to see for Player 3 of type s . Given that Player 1 and 2 play the combination (T, R) , l is better if the signal is positive, and r is better otherwise. As usual, the marginal player is the one with no perturbation, $s = 0$. So the proposed "Bayesian" pure strategy for player 3 is thus a best reply.

Given that Player 3 plays in this manner, and Player 2 plays R , Player 1 faces exogenous Knightian uncertainty given by $\mathcal{Q}_3 = [0, 1]$ over possible probabilities for action l of Player 3. Indeed, since we have allowed for complete Knightian uncertainty regarding perturbation, we can easily write down a probability distribution over $[-1, 1]$ that assigns probability p to the interval $[-1, 0]$ for any arbitrary p . All probability distributions over the actions of Player 3 are possible from the point of view of Player 1. If Player 1's ambiguity aversion is sufficiently high in the sense that

$$\phi_{13}([0, 1]) \supset \left[\frac{3}{8}, \frac{5}{8} \right],$$

then his best reply is to play T since the minimal expected payoff by deviating to B is less or equal than $3/8 \cdot 8 \leq 3$. The analogous reasoning holds true for Player 2. We have thus purified the equilibrium.

The new equilibrium can alternatively be interpreted as a usual equilibrium in a Bayesian game when the players have insufficient information about the exogenous distribution of types.

A completely general theory of the purification of equilibria in the extended game is not yet available. In Decerf and Riedel (2016), we discuss two player games with two strategies. In these games, as we discussed in Riedel and Sass (2014), the maxmin strategy plays a crucial role as it allows for hedging Knightian uncertainty. When the maxmin strategy is a pure strategy, the new equilibria of the extended game can be purified. In general, some types use the maxmin strategy in the equilibria of the perturbed game. In this sense, we are able to disambiguate the (Ellsberg) equilibria of the extended game.

5 Objections and Interpretations

I have presented these ideas at many universities over the past five years. As the new theory deviates from the received one, a certain reservation or

skepticism is natural and well deserved. Next, I would like to discuss some of the objections.

How can one play such strategies? The first question is typically: “But how can such a new strategy be played?” This is a natural question, of course. Note, however, that the same question can be asked of a mixed strategy as well. How do you play a mixed strategy for which you are supposed to play U with a probability of 71.2 percent as in our first example above? We know that most humans are not able to randomize perfectly. So how can it be done? In some sense, you would try to come close to 70 percent, and if you have some experience with randomization, I trust you could come close to it. Nevertheless, the fact that it is difficult to play in this uncertain way applies both to mixed strategies and to the new uncertain strategies.

Note that when one plays a zero-sum game like Rock-Scissors-Paper, you do not really tell your opponent that you are going to throw a die. You just have to act, and you try to do so “in your mind”. In the same way, you can perform an uncertain action by behaving in an unpredictable way. What is important for the application of the theory, is not the exact device, but rather if the equilibrium prediction is robust enough such that it can be trusted to be applied to the real world.

In a sense, the new equilibria are more robust than classic mixed Nash equilibria. In our three-player example, a sufficient degree of uncertainty aversion is enough to generate the new equilibrium outcome. We do not have to rely on the exact possibility of humans to generate an Ellsberg urn with given probability bounds. It is sufficient that they behave ambiguously enough. This kind of uncertain behavior is much easier to achieve, in my opinion, than a randomized strategy. Podolski’s words “It is best when I do not know what I do” are thus much closer to an ambiguous behavior than to an exact random behavior. Some people attribute to Podolski the statement: “Today’s soccer is a lot like chess; just without dice.” Funny as it is, the statement contains a truth that we use in our approach: we use uncertainty without dice.

Bayesian Dogma From the more dogmatic side, the objection that in a “small” world, players “have to be Bayesian” is sometimes heard. Since no human “has to be” anything in a free world, this criticism can only apply to the “agent”, i.e. to the rational homo oeconomicus whom we model in game theory. The dogma goes back to Savage (1954) who derives subjective expected utility theory axiomatically. Savage claims that his axioms should hold true in small worlds in which the state space is not too complicated.

The simple games we looked at certainly qualify as small worlds. While I admire Savage’s work, it is nevertheless clear that his axioms (as well as Anscombe and Aumann’s axioms in their (1963) paper) impose very strong conditions on rational behavior. The Ellsberg experiments show convincingly, in my opinion, that rationality in small worlds does not imply Bayesian (subjective expected) utility. It is perfectly rational to apply some sort of sophisticated worst–case approach when the conditions of the environment are not perfectly known. Nothing more than this small requirement is needed for the extended games analyzed here. In fact, the vast decision–theoretic literature that has followed Schmeidler (1989) and Gilboa and Schmeidler (1993) has developed a sound conceptual foundation for rational behavior under Knightian uncertainty in a non-subjective expected utility fashion. The rational agent can thus exhibit uncertainty aversion, and it makes sense to study the consequences of such uncertainty aversion in games.

Belief-based interpretation Nash equilibria can be interpreted in different ways. In applied work, most people actually tend to take mixed strategies literally, as this paper did at the beginning. When we say “Player 1 mixes in equilibrium” we usually tend to mean it as a deliberately random behavior. On the other hand, such an interpretation is not always meaningful, as Schelling (1960) noted. We have already discussed Harsanyi’s purification approach above, but another approach interprets the mixed strategies as beliefs of the other players (Aumann and Brandenburger (1995)). In fact, to support a Nash equilibrium, it is sufficient that Player 1 believes that Player 2 is behaving in such and such a random way. Such a belief interpretation is feasible for the extended games described here as well. In this interpretation, the belief of Player i about other players’ behavior is the product of the reduced uncertain strategy profiles of $\otimes_{j \neq i} \mathcal{Q}_j$. If the beliefs are correct and if all players agree on the respective beliefs about the other players, we get another interpretation of the new equilibria (compare also Lo (1996)).

Population Interpretation The population interpretation of equilibrium can be traced back to Nash’s dissertation (1950), see also Hofbauer (2000) and Weibull (1995). Nash points out that one can view a mixed Nash equilibrium also as the result of boundedly rational players playing *in a statistical sense* against a certain population of *randomly* drawn opponents. The players are opposed to gather empirical information about the play of the game. If the average action observed converges to a stable number, players will best reply to the population, and the total population behavior will correspond to the mixed strategy equilibrium.

Note that the above reasoning assumes a stationary stochastic environment such that some form of the law of large numbers applies. As many games are played rather few than many times, it is quite natural to assume that we do not have such a stationary environment; there is substantial Knightian uncertainty about the actions in the population. In fact, the law of large numbers under Knightian uncertainty just states that the (cluster points of) long-run mean of a repeated number of independent experiments will be in some interval (see Peng (2007) and Maccheroni and Marinacci (2005)). There thus remains Knightian uncertainty even after many observations of past play. If players are averse to such statistical (model) uncertainty, we get another interpretation of the new equilibria in the extended game.

6 Vague Language, Uncertain Mechanisms, and Extensive-form Games

The idea of endogenous uncertainty in interactions is currently being explored in a variety of interesting studies.

Our approach has interesting applications for language (or signaling) games; in fact, ambiguous language can be quite useful in some situations, e.g. in diplomatic negotiations or in public announcements that might otherwise trigger harsh immediate reactions. In our leading example, we have found a new equilibrium which has a higher payoff for all players. Player 3 introduces Knightian uncertainty and thus leads the other players to see no reason to deviate from the efficient equilibrium. This example thus highlights already two things: 1) vagueness can be used as a threat and 2) Knightian uncertainty renders outcomes incentive-compatible that would not be under expected utility.

In a beautiful paper, Kellner and Quement (2017) show that players can achieve higher payoffs in cheap talk games when ambiguous communication is allowed. This approach provides an interesting rationale for the use of vague language in situations in which players share some interest in the outcome, but are not completely aligned.

Bose and Renou (2014) use the same idea in mechanism design to show that the designer can implement a larger class of social choice functions when he is allowed to introduce ambiguity by a cheap talk signaling game before the actual mechanism is carried out. Similarly, Di Tillio, Kos, and Messner (2017) show that a seller can increase his profit by using an ambiguous mechanism. Lang and Wambach (2013) show that uncertainty about fraud detection deters ambiguity-averse agents from reporting false insurance claims.

The above-mentioned papers have to deal with the important aspect of updating beliefs under ambiguity. For example, Bose and Renou explicitly assume that the agents behave in a dynamically inconsistent way by updating their beliefs prior by prior. For single optimization problems, the issue of dynamic consistency in situations with multiple priors is well understood. The set of priors has to be stable under pasting (or rectangular, as Epstein and Schneider (2003) have dubbed it; see also Riedel (2004) and Sarin and Wakker (1998)). It is not possible to directly apply this reasoning in games as players typically face different information flows. For example, Kuhn’s theorem does not hold true at the same level of generality as in the expected utility case as Muraviev, Riedel, and Sass (2017) show (compare also Aryal and Stauber (2014)). Nevertheless, in large classes of games, it is possible to identify reasonable restrictions to the set of ex ante uncertain strategies that allow to re-establish equivalence. In ongoing work, Hanany, Klibanoff, and Mukerji (2016) study extensive-form games with Knightian uncertainty regarding types and smooth utility. They develop the notion of sequential equilibrium for such games. The purified version of our extended game corresponds to such incomplete information games; it would thus be possible to develop a notion of sequential equilibrium from that point of view.

7 Conclusion

In this paper, we have taken up a new approach to uncertain actions in games which was originally proposed in Riedel and Sass (2014). We extend their approach by rooting it firmly in decision theory (Gajdos, Hayashi, Tallon, and Vergnaud (2008), Klibanoff, Marinacci, and Mukerji (2005)). Although the new approach is a parsimonious extension of classic game theory in the sense that no player has an incentive to introduce uncertain actions unilaterally, interesting new equilibria arise. We discuss purification in the spirit of Harsanyi (1973) and other possible interpretations like the original mass-action or population interpretation of Nash (1950) and the more modern belief-based interpretation of Aumann and Brandenburger (1995). Our model can be used as a rigorous modeling background for the successful applications of ambiguous strategies much as they have been used in the literature on mechanism design, principal agents, and moral hazard.

References

- ANSCOMBE, F., AND R. AUMANN (1963): “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199–205.
- ARYAL, G., AND R. STAUBER (2014): “A note on Kuhn’s Theorem with ambiguity averse players,” *Economics Letters*, 125(1), 110–114.
- AUMANN, R., AND A. BRANDENBURGER (1995): “Epistemic conditions for Nash equilibrium,” *Econometrica*, 63, 1161–1180.
- BADE, S. (2011): “Ambiguous act equilibria,” *Games and Economic Behavior*, 71(2), 246–260.
- BOSE, S., AND L. RENOU (2014): “Mechanism design with ambiguous communication devices,” *Econometrica*, 82(5), 1853–1872.
- DECERF, B., AND F. RIEDEL (2016): “Disambiguation of Ellsberg Equilibria in 2x2 Normal Form Games,” Center for Mathematical Economics Working Paper 554.
- DI TILLIO, A., N. KOS, AND M. MESSNER (2017): “The design of ambiguous mechanisms,” *The Review of Economic Studies*, pp. 237–276.
- EICHBERGER, J., AND D. KELSEY (2011): “Are the Treasures of Game Theory Ambiguous?,” *Economic Theory*, 48, 313–339.
- ELLSBERG, D. (1961): “Risk, ambiguity, and the Savage axioms,” *The Quarterly Journal of Economics*, 75(4), 643–669.
- EPSTEIN, L., AND M. SCHNEIDER (2003): “Recursive Multiple Priors,” *Journal of Economic Theory*, 113, 1–31.
- 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, 33–49.
- GIRAUD, R. (2014): “Second order beliefs models of choice under imprecise risk: Nonadditive second order beliefs versus nonlinear second order utility,” *Theoretical Economics*, 9(3), 779–816.

- GOEREE, J., AND C. HOLT (2001): “Ten little treasures of game theory and ten intuitive contradictions,” *American Economic Review*, 91(5), 1402–1422.
- GRANT, S., I. MENEGHEL, AND R. TOURKY (2016): “Savage games,” *Theoretical Economics*, 11(2), 641–682.
- HANANY, E., P. KLIBANOFF, AND S. MUKERJI (2016): “Incomplete information games with smooth ambiguity preferences,” .
- HARSANYI, J. (1973): “Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points,” *International Journal of Game Theory*, 2(1), 1–23.
- HOFBAUER, J. (2000): “From Nash and Brown to Maynard Smith:Equilibria, Dynamics, and ESS,” *Selection*, 1, 81–88.
- KELLNER, C., AND M. T.-L. QUEMENT (2017): “Endogenous AMbiguity in Cheap Talk,” Working Paper, Bonn University.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A smooth model of decision making under ambiguity,” *Econometrica*, 73(6), 1849–1892.
- LANG, M., AND A. WAMBACH (2013): “The fog of fraud—mitigating fraud by strategic ambiguity,” *Games and Economic Behavior*, 81, 255–275.
- LO, K. (1996): “Equilibrium in beliefs under uncertainty,” *Journal of Economic Theory*, 71(2), 443–484.
- MACCHERONI, F., AND M. MARINACCI (2005): “A strong law of large numbers for capacities,” *Annals of probability*, pp. 1171–1178.
- MARINACCI, M. (2000): “Ambiguous Games,” *Games and Economic Behavior*, 31, 191–219.
- MURAVIEV, I., F. RIEDEL, AND L. SASS (2017): “Kuhn’s theorem for extensive form Ellsberg games,” *Journal of Mathematical Economics*, 68, 26–41.
- NASH, J. (1950): “Non-cooperative Games,” Ph.D. thesis, Princeton University.
- PENG, S. (2007): “Law of large numbers and central limit theorem under nonlinear expectations,” *arXiv preprint math/0702358*.

- RIEDEL, F. (2004): “Dynamic Coherent Risk Measures,” *Stochastic Processes and Their Applications*, 112, 185–200.
- RIEDEL, F., AND L. SASS (2014): “Ellsberg Games,” *Theory and Decision*, 76(4), 469–509.
- SARIN, R., AND P. WAKKER (1998): “Dynamic Choice and Nonexpected Utility,” *Journal of Risk and Uncertainty*, 17, 87–119.
- SAVAGE, L. (1954): *The Foundations of Statistics*. Wiley, New York.
- SHELLING, T. C. (1960): *The Strategy of Conflict*. Harvard University Press.
- SCHMEIDLER, D. (1989): “Subjective probability and Expected Utility without Additivity,” *Econometrica*, 57, 571–587.
- STAUBER, R. (2017): “Irrationality and ambiguity in extensive games,” *Games and Economic Behavior*, 102, 409–432.
- V. NEUMANN, J. (1928): “Zur Theorie der Gesellschaftsspiele,” *Mathematische Annalen*, 100(1), 295–320.
- VON NEUMANN, J., AND O. MORGENSTERN (1953): *Theory of Games and Economic Behavior*. Princeton University Press.
- WEIBULL, J. (1995): *Evolutionary Game Theory*. MIT Press, Cambridge, MT.