MPRA
Munich Personal RePEc Archive

Mixed Strategies in Games with Ambiguity Averse Agents

Calford, Evan

Purdue University

31 October 2016

Online at https://mpra.ub.uni-muenchen.de/74909/
MPRA Paper No. 74909, posted 08 Nov 2016 14:42 UTC

Mixed strategies in games with ambiguity averse agents *

Evan M. Calford †

October 31, 2016

Abstract

In normal form games, when agents exhibit ambiguity aversion the exclusion of mixed strategies from agents' choice sets can enlarge the set of equilibria. While it is possible, in a game theoretic experiment, to enforce pure strategy reporting it is not possible to prevent subjects from mixing before reporting a pure strategy. This short paper establishes conditions under which the set of equilibrium in a game with ambiguity averse agents and pure strategy reporting is invariant to the existence of pre-play mixing devices. This result is crucial for the interpretation of recent experimental work on the role of ambiguity aversion in normal form games.

Keywords: Ambiguity Aversion, Mixed Strategies, Game Theory, Experimental Economics

JEL Codes: D81, C92, C72 and D03

1 Introduction

Consider a game between two agents that is mediated by a game theorist. The agents report their strategies to the game theorist, who then resolves the outcome of the game and pays the agents their winnings (or collects their losses). The game theorist may allow mixed strategy reports from the agents, and resolve the mixed strategy herself, or she may require that the agents report a pure strategy. If the game theorist requires pure strategy reports, as is typically the case in economic experiments, then the game theorist should be aware that the agents may still be using a mixed

*This paper was originally, in a much longer form, a part of my PhD thesis. I would like to thank Yoram Halevy and Wei Li for providing many helpful suggestions that have improved this paper, as well as Ryan Oprea and Mike Peters.

†Department of Economics, Krannert School of Management, Purdue University, West Lafayette, IN, 47907, ecalford@purdue.edu.

strategy that they are resolving privately before reporting. In the standard case, where agents have expected utility preferences, the set of Nash equilibrium under the two reporting requirements will be indistinguishable.

However, when agents have ambiguity averse preferences then the different reporting requirements may induce different games with different equilibria. Equilibrium concepts such as Lo (2009), Dow and Werlang (1994) and Eichberger and Kelsey (2000), which enforce pure strategy reporting, generate larger equilibrium sets than Lo (1996) which allows for mixed strategy reports. The difference arises because ambiguity averse preferences are non-linear and agents may have a strict preference for reporting mixed strategies. The equilibrium concepts that enforce pure strategy reporting assume that only pure strategies are available to agents: they implicitly rule out private pre-play mixing. By its very nature, however, private pre-play mixing will be unobservable to the game theorist and cannot readily be prevented.

How, then, does allowing for private pre-play mixing affect the pure strategy only equilibrium in papers such as Lo (2009)? The answer, provided in this paper, is that allowing for private pre-play mixing has no effect on the equilibrium set for agents with preferences that lie in the intersection of Choquet Expected Utility (Schmeidler, 1989) and Maxmin Expected Utility (Gilboa and Schmeidler, 1989). While this result is of independent interest, it is particularly relevant for experimental tests of ambiguity averse equilibrium concepts. Recent experiments have used equilibrium concepts that restrict agents to pure strategies, and thereby implicitly assume that their subjects are not engaging in pre-play mixing.[1] Given that it is not possible to actively prevent subjects from pre-play mixing, the results in this paper are essential for a direct interpretation of the data in the previous experimental literature using pure strategy solution concepts.

To establish the main result, we reinterpret a formal mathematical result from Gilboa and Schmeidler (1994) (that a non-additive measure can be spanned by an appropriately chosen additive measure over a larger state space) and translate it from an individual decision making environment to a non-cooperative game. Each agent then plays a “mental” game that is closely related to, but distinct from, the game presented by the game theorist. The agent resolves their mixed strategy with respect to the mental game and reports the pure strategy realization to the game theorist, who then implements the strategy in the original game. The mental game is interpreted as a fictional accounting device used by the agent. Under this interpretation of a game, it is possible to rewrite a pure strategy equilibrium in a fashion that is consistent with pre-play mixing (see Section 3.1).

[1]See Kelsey and le Roux (2015), Kelsey and le Roux (2016), Calford (2016), Ivanov (2011) and Eichberger et al. (2008) for examples.

The rest of this paper is organized as follows. Section 2 introduces some mathematical tools from Gilboa and Schmeidler (1994) that are used in the rest of the paper. Section 2.1 introduces the notion of a game as a set of interacting decision problems and presents the key idea of a "mental" state space. Section 3 applies the structure of a "mental" state space to the equilibrium concept contained in Lo (2009) and establishes the key result of this paper. Section 4 concludes.

2 Preliminaries

Suppose that there exists a finite set of *states* of the world, $\omega \in \Omega$, and that an *act*, f, maps each state to an *outcome* in $\mathbb{R}$; that is $f : \Omega \mapsto \mathbb{R}$. Choquet Expected Utility (CEU) generalizes Subjective Expected Utility by allowing a decision maker to hold non-additive beliefs which are represented by a capacity, ν, defined over the set of events $\Sigma = 2^{\Omega}$. Suppose, without loss of generality, that for a given act, f, the set of states can be ordered so that $f(\omega_1) \geq f(\omega_2) \geq \ldots \geq f(\omega_n)$. A CEU agent calculates her utility of an act by evaluating the Choquet integral:

$$\int f\, dx = \int_0^{\infty} \nu(\{\omega | f(\omega) \geq t\})\, dx + \int_{-\infty}^{0} [\nu(\{\omega | f(\omega) \geq t\}) - \nu(\Omega)]\ dx \tag{1}$$

We shall assume throughout that the capacities, ν, are *belief functions* and we use V to denote the space of all such capacities. That is, we assume that $\nu(\Sigma) = 1$ and that ν is totally monotone.[2] Under these assumptions it is also possible to represent the agent's preferences using Maxmin Expected Utility (Gilboa and Schmeidler, 1994).

Throughout this paper, we shall rely on two basic mathematical results that are demonstrated in Gilboa and Schmeidler (1994). Firstly, the non-additive measure ν can be spanned by an additive measure over an appropriately defined (larger) state space. Secondly, we can represent an agent with CEU preferences over Ω as, equivalently, having SEU preferences over the larger state space with an appropriately transformed set of acts.

[2] A capacity is totally monotone if $\nu(A) \geq 0$ for all $A \in \Sigma$ and, for every $A_1, A_2, \ldots, A_n \in \Sigma$, $\nu(\cup_{i=1}^{n} A_i) \geq \sum_{\{I|\emptyset \neq I \subseteq \{1,2,\ldots,n\}\}} (-1)^{|I|+1} \nu(\cap_{i\in I} A_i)$. Intuitively, a capacity is totally monotone if it is non-negative and the capacity of every event is larger than the sum of the capacities of all its sub-events (after accounting for the fact that each state is a member of multiple events).

Result 1 (Adapted from Gilboa and Schmeidler (1994)). *For $T, A \in \Sigma' = \Sigma \backslash \{\emptyset\}$, define*

$$e_T(A) = \begin{cases} 1 & T \subseteq A \\ 0 & otherwise \end{cases}$$

Then the set $\{e_T\}_{T \in \Sigma'}$ forms a linear basis for V. The unique coefficients $\{\alpha_T^\nu\}$ satisfying

$$\nu = \sum_{T \in \Sigma'} \alpha_T^\nu e_T$$

are given by

$$\alpha_T^\nu = \sum_{S \subseteq T} (-1)^{|T|-|S|} \nu(S) \tag{2}$$

Furthermore, if ν is totally monotone then $\alpha_T^\nu \geq 0$ for all $T \in \Sigma'$ and if ν is normalized then $\sum_{T \in \Sigma'} \alpha_T^\nu = 1$.

Result 1 provides the key building block for this paper: any non-additive measure over a state space can be spanned by an appropriately formed set of states constructed from the power set of the original state space. Furthermore, when the non-additive measure is a belief function then the spanning coefficients can be interpreted as probabilities over the newly constructed state space. Note the relationship between Result 1 and the proof of the representation theorem for CEU in Schmeidler (1989). In Result 1, we begin with a non-additive measure and 'restore' additivity by extending the state space. In Schmeidler (1989), the primitive is a SEU representation with respect to an additive measure, which is then extended to generate a CEU representation with respect to a non-additive measure. This tight relationship between Choquet Expected Utility and Subjective Expected Utility is formalized in the next result.

Recall that when ν is a belief function the core of ν, denoted by Core(ν), is simply the set of probability measures, p, such that $p(A) \geq \nu(A)$ for all $A \in \Sigma$.[3] We now state Result 2.

Result 2 (Corollary 4.4 from Gilboa and Schmeidler (1994)). *Suppose that ν is a belief function. Then for every $f \in F$*

$$\int f \, d\nu = \sum_{T \in \Sigma'} \alpha_T^\nu \left[\min_{\omega \in T} f(\omega) \right] \tag{3}$$

$$= \min_{p \in Core(\nu)} \sum_{\omega \in \Omega} p(\{\omega\}) f(\omega). \tag{4}$$

[3]Note that, because each p in the core is additive, the core can be equivalently defined as the restriction to an equivalent set of probability measures defined over Ω. In some places we use core (ν) to denote this restriction, but this should be clear from the context.

Result 2 demonstrates that an agent with CEU preferences with respect to a belief function can have their preferences represented via either MEU or SEU preferences. While this relation between MEU and CEU preferences is both straightforward and well known, the representation with SEU preferences requires the formation of a new set of acts over the set $\Sigma = 2^{\Omega}$, with the outcome associated with each new act defined by the min function in Equation 3. We shall call these acts "mental" acts, and will sometimes refer to the event space as the "mental" state space. This terminology reflects that the event space may not be observable to an external observer and may, therefore, represent the mental accounting of the agent.[4]

Definition 1. A mental act, f' , is an extension of an act, f, defined over the event space, $\Sigma' = 2^{\Omega}/\emptyset$, such that $f' : \Sigma' \mapsto \mathbb{R}$ with $f'(T) = \min_{\omega \in T} f(\omega)$ for all $T \in \Sigma'$.

It follows from Result 2 that the preferences of an agent who has CEU preferences with respect to a belief function can be written in the expected utility form with respect to mental acts over the mental state space.[5] The key feature of this preference representation is that preferences over mental acts are linear which, as discussed below, implies that there is no preference for randomization over mental acts.

2.1 Games as Interacting Decision Problems

It is possible to reformulate a standard game as a set of interacting decision problems.[6] Recall that a normal-form game is fully defined by the set of players, I, a set of (pure) strategies, $\{S_i\}$, for each player and a set of utility functions, $\{u_i\}$. For each agent we can formulate their decision problem as choosing from a set of acts over states where the states are given by $\Omega_i = \{\times S_{j\neq i}\}$ and $f_i(\omega) = s_i(\omega) = u_i(s_i, \omega) = u_i(s_i, s_{j\neq i})$.

The interpretation of mixed strategies in a game with ambiguity averse agents requires some

[4]The interpretation of the event space as a mental "state" space implies that the agent, at least in his mental accounting, is not well calibrated about the nature of the world. In a sense, this is precisely the trade-off that allows us to move from non-linear preferences in the observable state space to linear preferences in the mental state space.

[5]Billot and Vergopoulos (2014) generalize this result to a broader set of preferences than just CEU with respect to a belief function, and explore the implications of Result 2 for problems with a single decision maker in some detail.

[6]Aumann (1987) is the classic example. Note that one key feature of Aumann's approach is that mixed strategy equilibrium can be interpreted as an equilibrium in beliefs, without any need for explicit mixing by agents. In contrast, this paper considers an environment where agents may strictly prefer to use mixing devices; assuming away mixing devices when they may in fact be available is then a mistake.

care. There are two approaches that have been taken in the prior literature. The first approach, analagous to using a Savage (1954) interpretation of the state space in decision theory, postulates that mixed strategies are not available to agents. A mixed strategy equilibrium is then interpreted as an equilibrium in beliefs, and the resultant equilibrium concepts typically generalize Nash equilibrium even in two-player games (see Dow and Werlang (1994), Epstein (1997), Groes et al. (1998), Marinacci (2000), Lo (2009), Grant et al. (2016) and Eichberger and Kelsey (2014) for examples).

The second approach, analagous to using an Anscombe and Aumann (1963) interpretation of uncertainty in decision theory, gives agents access to mixed strategies. That is, the choice set of the agent now consists of $\sigma_i \in \Delta(S_i)$. Payoffs for mixed strategies are defined state-by-state so that mixed strategies may, in some cases, provide a hedge against strategic uncertainty (see Section A.2 for an example). One consequence of including mixed strategies in this fashion is that the equilibrium set in a two-player game with ambiguity averse agents has the same (pure strategy) support as the set of Nash equilibrium. See Lo (1996) for a detailed discussion, as well as Klibanoff (1996) and Azrieli and Teper (2011).

Recently, the decision theory and experimental design literature has drawn a distinction between the case where subjective uncertainty is resolved before objective risk and the case where risk is resolved before uncertainty (Seo (2009), Azrieli et al. (2016), Baillon et al. (2014) and Eichberger et al. (2016)). If subjective uncertainty is resolved before objective risk then a preference for randomization arises, while if risk is resolved before uncertainty then there is no preference for randomization. The approach found in Lo (1996), Klibanoff (1996), and Azrieli and Teper (2011) is consistent with the subjective before objective interpretation: agents report mixed strategies to a game theorist who then realizes the mixed strategies.

In contrast, when working in the mental state space it is natural that objective uncertainty is resolved before subjective uncertainty. The mental state space (and the mental acts) are a construct of the agents' mental accounting within which the agent may randomize as they please, but is detached from the environment that is identifiable by an outside observer. When applied to a game, this implies that we can interpret a game played in the mental state space as having agents that resolve their mixed strategy privately before reporting a pure strategy to the game theorist.

This provides an alternative *interpretation* of a game. The mental state space has two properties that make it particularly useful in experimental applications. First, as demonstrated below, the equilibrium set is unchanged when moving from a pure-strategies-only framework to the mental state space framework. Second, because experimental economists cannot readily prevent pre-play mixing by their subjects, this ensures that there is an equivalence, when the experiment is restricted to pure strategies, between experiments with subjects who utilize pre-play mixing and experiments where subjects do not utilize pre-play mixing. In contrast, experiments with explicitly mixed strategy sets are not equivalent.

In what follows, we more formally establish the equivalence between these two interpretations of equilibrium for the case of Lo-Nash equilibrium (Lo, 2009). The same proof technique can be applied to establish the equivalence for other existing pure-strategy-only equilibrium concepts.[7]

3 Lo-Nash Equilibrium

This section introduces Lo-Nash equilibrium, following Lo (2009) closely. Define a set of players $N = \{1, \ldots, n\}$, let each player $i \in N$ have a finite set of actions A_i, and define $A = \times_{i \in N} A_i$ and $A_{-i} = \times_{j \neq i \in N} A_j$. We shall endow each agent with a Von Neumann-Morgenstern utility function $u : A \mapsto \mathbb{R}$. Suppose that an agent has uncertainty regarding the strategy choices of their opponents, A_{-i}. Then we can regard a strategy, a_i, as an act over the state space A_{-i} generating a payoff $u_i(a_i, a_{-i})$ when the state a_{-i} is realized.

In a manner consistent with Gilboa and Schmeidler's (1989) MEU formulation, we suppose that an agent's beliefs regarding their opponents strategies are a closed and convex set of probability measures $\Phi_i \subseteq \Delta(A_{-i})$. Given Φ_i an agents preferences are represented by

$$\min_{\phi \in \Phi} \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \phi_i(a_{-i}).$$

Furthermore, we use σ to denote a probability measure on A. We define $\sigma^{A_i}(a_i) = \sum_{a_{-i} \in A_{-i}} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_i and $\sigma^{A_{-i}}(a_{-i}) = \sum_{a_i \in A_i} \sigma(a_i, a_{-i})$ as the marginal distribution of σ on A_{-i}. Then, in the usual fashion we write $\sigma(a_{-i}|a_i) = \frac{\sigma(a_i, a_{-i})}{\sigma^{A_i}(a_i)}$.

Finally, we write $\operatorname{supp} \sigma$ to denote the support of the probability distribution σ, and define $\operatorname{supp} \Phi_i$ to be the union of the supports of the elements of Φ_i, and write Φ to denote the profile $(\Phi_i)_{i \in N}$. We are now ready to define a Lo-Nash equilibrium.

Definition 2 (Lo-Nash equilibrium)**.** A pair $< \sigma, \Phi >$ forms a Lo-Nash equilibrium if it satisfies

$$\sigma(\cdot|a_i) \in \Phi_i \quad \forall a_i \in \operatorname{supp} \sigma^{A_i}, \forall i \in N \tag{5}$$

$$\operatorname{supp} \Phi_i = \times_{j \neq i} \operatorname{supp} \sigma^{A_j} \quad \forall i \in N \tag{6}$$

[7] An earlier version of this paper, which was included in the author's PhD thesis, also establishes the result for the equilibrium concept found in Dow and Werlang (1994).

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \min_{\phi_i \in \Phi_i} \sum_{a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i})\phi_i(a_{-i}) \quad \forall a_i \in \operatorname{supp} \sigma^{A_i}, \forall i \in N \tag{7}$$

Equation 7 requires that all strategies that are played in an equilibrium are best responses *among the set of pure strategies*, with preferences defined as MEU preferences with respect to the equilibrium conjectures Φ. Equations 5 and 6 are the consistency requirements: Equation 6 ensures that a strategy is played with a positive probability iff it is expected to be played with a positive probability, and Equation 5 forces actual strategies to be contained in the belief sets. Note that Equation 5 allows for conditioning of σ on a_i - this allows for strategies to be correlated, but the realized strategy must lie within player i's belief set for all a_i. Nash equilibrium is a special case of Lo-Nash equilibrium, thereby ensuring existence of Lo-Nash equilibrium for all finite normal form games.

Lo explicitly restricts attention to pure strategies to prevent agents from holding a strict preference for mixing. Mixed equilibrium are, therefore, interpreted as equilibrium in beliefs. In the following section we demonstrate that by reformulating the problem using the mental state space, we can allow for explicit pre-play mixing without altering the equilibrium set.

3.1 Formulating Lo-Nash equilibrium as a mental equilibrium

To introduce a mental version of Lo-Nash equilibrium we first need to extend the state space into the event space and define preferences over the new state space. In a Lo-Nash equilibrium, each agent i faces a state space A_{-i}: their opponent's strategy space is their state space. We denote the extended state space for agent i to be $\Sigma_i = \{2^{A_{-i}}/\emptyset\}$. Then we define an agents' extended utility function, for all $T \in \Sigma_i$, as:

$$u'(a_i, T) = \min_{a_{-i} \in T} u(a_i, a_{-i}).$$

Notice that when $T \in A_{-i}$ then $u'(a_i, T) = u(a_i, a_{-i})$ so that the extended utility function is consistent with the original utility function. We also define the natural extension of σ over Σ_i: $\sigma(a_i, T) = \sum_{a_{-i} \in T} \sigma(a_i, a_{-i})$ for all $T \in \Sigma_i$. Finally, we introduce the probability measure $\alpha_i \in \Delta(\Sigma_i)$ which can be interpreted as agent i's belief over the event space Σ_i. We write α to denote the profile $(\alpha_i)_{i \in N}$.

We are now ready to define our Mental Lo-Nash equilibrium. In essence Mental Lo-Nash equilibrium is simply an application of Result 1 to Lo-Nash equilibrium, as we establish in Theorem 1 below.

Definition 3 (Mental Lo-Nash equilibrium)**.** A pair $<\sigma, \alpha>$ form a Mental Lo-Nash equilibrium if

$$\sigma(T|a_i) \geq \sum_{\tau \subseteq T} \alpha_i(\tau) \quad \forall T \in \Sigma_i, \forall a_i \in \operatorname{supp} \sigma^{A_i}, \forall i \in N \tag{8}$$

$$\sum_{\{T: a_{-i} \notin T\}} \alpha_i(T) = 1 \Leftrightarrow \prod_{j \neq i} \sigma^{A_j}(a_j) = 0 \quad \forall a_{-i} \in A_{-i}, \forall i \in N \tag{9}$$

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \sum_{T \in \Sigma_i} u'_i(\hat{a}_i, T)\alpha_i(T) \quad \forall a_i \in \text{ supp } \sigma^{A_i}, \forall i \in N \tag{10}$$

Note that the existence of Mental Lo-Nash equilibrium is guaranteed whenever a Nash equilibrium exists. Furthermore, the formulation of the Mental Lo-Nash equilibrium is linear in all parameters, which implies that solving for a Lo-Nash equilibrium is a linear complementarity problem. In particular, if the game theorist is willing to assume a support for σ, then solving for equilibrium reduces to solving a linear program.

Theorem 1. *If $<\sigma, \alpha>$ is a Mental Lo-Nash equilibrium then there exists a Φ such that $<\sigma, \Phi>$ is a Lo-Nash equilibrium. Conversely, if $<\sigma, \Phi>$ is a Lo-Nash equilibrium and Φ_i is the core of a belief function for all i then there exists an α such that $<\sigma, \alpha>$ is a mental Lo-Nash equilibrium.*

The interpretation of a Mental Lo-Nash equilibrium is straightforward when σ is a product measure (i.e. strategies are uncorrelated). In this case, each agent uses an independent pre-play mixing device yet may place a positive weight on beliefs that imply correlation between their opponents, or even that their opponents strategies may be correlated with their own strategy. Such beliefs may be justified in cases where an agent believes that their opponents may be able to communicate, or where they fear that their opponents may be able to see the agent's strategy choice (or a signal correlated with their strategy choice) before making a decision. It is feasible that some subjects in a standard game theory experiment may not trust the experimenter and may hold such concerns.

When σ is not a product measure, the interpretation is more complicated. As an example, imagine a Professor of Business Strategy who suggests that their students use the positioning of the second hand on a clock face as a randomization device – when students of the same Professor play against each other in a simultaneous game, this will be enough to induce correlated randomizations.

4 Conclusions

This paper presents a methodology for extending pure strategy only ambiguity averse equilibrium concepts to allow for (pre-play) mixed strategies.

The theoretical interest in the structure and interpretation of mixed strategy equilibrium for agents with uncertainty averse preferences is readily apparent. Unlike SEU agents, ambiguity averse agents are not necessarily indifferent between a mixed strategy and the pure strategy supports of the mixed strategy. Considering this, some of the standard interpretations of mixed strategies (via purification arguments or population interpretations of mixed strategies) may not be appropriate in the context of ambiguity averse agents; indeed, a majority of ambiguity averse equilibrium concepts explicitly restrict their analysis to pure strategies. In this context, the methodology introduced here can be viewed as an equilibrium-preserving interpretation of mixed strategies: when agents are ambiguity averse and have access to a mixing device that resolves before the strategic interaction occurs then the agents will, in equilibrium, be indifferent between using the mixing device or not. In this fashion the population interpretation of mixing can be restored without the need to explicitly assume away the existence of mixing devices.

The ubiquitous rise of the internet and digital communications has lead to a proliferation of pre-programmed bots that are increasingly interacting in strategic situations.[8] It may be the case that, in some cases, bots that play pure strategies are exploitable thereby necessitating the use of randomization. In cases where the programmer has expected utility preferences over the outcomes produced by their bot there is no conceptual problem with evaluating the (ex-ante) expected performance of differently specified programs. However, in cases where the programmer has ambiguity averse preferences the evaluation of such programs is a conceptually challenging task. The theory presented in this paper presents one potential way forward.

References

F.J. Anscombe and R.J. Aumann. A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1):199–205, 1963. → pages 6

Robert J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55(1):1–18, 1987. → pages 5

[8]Online pricing algorithms for concerts and sporting events interacting with automated purchasing algorithms owned by scalpers are but one example.

Yaron Azrieli and Roee Teper. Uncertainty aversion and equilibrium existence in games with incomplete information. *Games and Economic Behavior*, 73:310–317, 2011. → pages 6

Yaron Azrieli, Christopher P. Chambers, and P. J. Healy. Incentives in experiments: A theoretical analysis. Working Paper, 2016. → pages 6

Aurelien Baillon, Yoram Halevy, and Li Chen. Experimental elicitation of ambiguity attitude using the random incentive system. Mimeo, 2014. → pages 6

Antoine Billot and Vassili Vergopoulos. Expected utility without parsimony. CES Working Paper, March 2014. → pages 5

Evan M. Calford. Ambiguity aversion in game theory: Experimental evidence. Mimeo, 2016. → pages 2

James Dow and Sergio Riberio da Costa Werlang. Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *Journal of Economic Theory*, 64:305–324, 1994. → pages 2, 6, 7

Jurgen Eichberger and David Kelsey. Non-additive beliefs and strategic equilibria. *Games and Economic Behavior*, 30:183–215, 2000. → pages 2

Jurgen Eichberger and David Kelsey. Optimism and pessimism in games. *International Economic Review*, 55(2):483–505, 2014. → pages 6

Jurgen Eichberger, David Kelsey, and Burkhard C. Schipper. Granny versus game theorist: Ambiguity in experimental games. *Theory and Decision*, 64:333–362, 2008. → pages 2

Jurgen Eichberger, Simon Grant, and David Kelsey. Randomization and dynamic consistency. *Economic Theory*, 62(3):547–566, 2016. → pages 6

Larry G. Epstein. Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73: 1–29, 1997. → pages 6

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989. → pages 2, 7

Itzhak Gilboa and David Schmeidler. Additive representations of non-additive measures and the Choquet integral. *Annals of Operations Research*, 52:43–65, 1994. → pages 2, 3, 4

Simon Grant, Idione Meneghel, and Rabee Tourky. Savage games. *Theoretical Economics*, 11: 641–682, 2016. → pages 6

Ebbe Groes, Hans Jorgen Jacobson, Birgitte Sloth, and Torben Tranaes. Nash equilibrium with lower probabilities. *Theory and Decision*, 44:37–66, 1998. → pages 6

Asen Ivanov. Attitudes to ambiguity in one-shot normal-form games: An experimental study. *Games and Economic Behavior*, 71:366–394, 2011. → pages 2

David Kelsey and Sara le Roux. An experimental study on the effect of ambiguity in a coordination game. *Theory and Decision*, 79(4):667–688, 2015. → pages 2

David Kelsey and Sara le Roux. Dragon slaying with ambiguity: Theory and experiments. *Journal of Public Economic Theory*, Forthcoming, 2016. → pages 2

Peter Klibanoff. Uncertainty, decision and normal form games. Mimeo, July 1996. → pages 6

Kin Chung Lo. Equilibrium in beliefs under uncertainty. *Journal of Economic Theory*, 71: 443–484, 1996. → pages 2, 6, 14

Kin Chung Lo. Correlated Nash Equilibrium. *Journal of Economic Theory*, 144:722–743, 2009. → pages 2, 3, 6, 7

Massimo Marinacci. Ambiguous games. *Games and Economic Behavior*, 31:191–219, 2000. → pages 6

Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954. → pages 6

David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989. → pages 2, 4

Kyoungwon Seo. Ambiguity and second-order belief. *Econometrica*, 77(5):1575–1605, 2009. → pages 6

A Appendix

A.1 Proof of Theorem 1

Lemma 1. Suppose that ν is a belief function and $\Phi = \text{core}\ (\nu)$. Then $A \in \text{ supp } \Phi$ if and only if $\nu(A^c) < 1$.

Proof of Lemma 1. Suppose that $A \in \text{supp } \Phi$. Therefore there exists a $\phi \in \Phi$ such that $\phi(A) > 0$ and $\phi(A^c) < 1$. Therefore there exists a $p \in \text{core}\ (\nu)$ such that $p(A^c) < 1$. Therefore $\nu(A^c) < 1$.

Suppose that $A \notin \text{supp } \Phi$. Therefore $\phi(A) = 0$ for all $\phi \in \Phi$. Therefore $p(A) = 0$ for all $p \in \text{core}\ (\nu)$. Therefore $p(A^c) = 1$ for all $p \in \text{core}\ (\nu)$. Note that there always exists a $p \in \text{core}\ (\nu)$ such that $p(A^c) = \nu(A^c)$. Therefore $\nu(A^c) = 1$. □

Lemma 2 (Lo-Nash equilibrium)**.** Consider a pair $< \sigma, \Phi >$ such that $\Phi_i = \text{ core } \nu_i$ for all i, where ν_i is a belief function. If σ and ν_i satisfy:

$$\sigma(\cdot|a_i) \in \text{core } \nu_i \quad \forall a_i \in \text{supp}\, \sigma^{A_i}, \forall i \in N \tag{A.1}$$

$$\nu_i(a_{-i}^c) = 1 \Leftrightarrow \prod_{j \neq i} \sigma^{A_j}(a_j) = 0 \quad \forall a_{-i} \in A_{-i}, \forall i \in N \tag{A.2}$$

and

$$a_i \in \arg\max_{\hat{a}_i \in A_i} \int \hat{a}_i d\nu_i \quad \forall a_i \in \text{supp}\, \sigma^{A_i}, \forall i \in N \tag{A.3}$$

then the pair $< \sigma, \Phi >$ form a Lo-Nash equilibrium. Conversely, if $< \sigma, \Phi >$ are a Lo-Nash equilibrium then σ and ν_i satisfy Equations A.1, A.2 and A.3.

Proof of Lemma 2. We need to demonstrate that $< \sigma, \Phi >$ is a Lo-Nash equilibrium.

Equation A.1 is equivalent to Equation 5.

The equivalence of Equation 7 and equation Equation A.3 follows directly from Result 2.

$\prod_{j \neq i} \sigma^{A_j}(a_j) = 0 \Leftrightarrow \exists j \neq i \text{ s.t. } \sigma^{A_j}(a_j) = 0 \Leftrightarrow a_{-i} \notin \times_{j \neq i} \text{ supp } \sigma^{A_j}$ and Lemma 1 establish that equation 6 and equation A.2 are contrapositives.

□

Proof of Theorem 1. Given the above Lemmas, it is sufficient to show an equivalence between Equations A.1 and 8, Equations A.2 and 9, and Equations A.3 and 10. We begin by noting that

Result 1 ensures that each α_i can be associated with a belief function ν_i (and vice-versa) and that Φ_i can be defined as the core of ν_i.

The equivalence of Equations A.3 and 10 follows from Results 1 and 2.

The equivalence of Equations A.1 and 8 is a consequence of the definition of the core of a capacity and the fact that $\nu_i(B) = \sum_{T \subseteq B} \alpha_i(T)$ (which follows directly from the definition of α_i).

The equivalence of Equations A.2 and 9 follows immediately from $\nu_i(a^c_{-i}) = \sum_{T: a_{-i} \notin T} \alpha_i(T)$ which again is an immediate consequence of the definition of α_i.

□

A.2 Example of mixing as a hedge against uncertainty

To illustrate the role of mixing in the mental state space we present a simple example. Consider a world with two states $\Omega = \{U, D\}$ and two acts l and r. The agent earns a payoff of 1 if they choose l when the state is U or choose r when the state is D and nothing otherwise.[9] Further, suppose that the agent has complete subjective uncertainty regarding the true state of the world, so that $\nu(U) = \nu(D) = 0$ and $\nu(\{U, D\}) = 1$. The agent then holds preferences such that $l \sim r \sim 0$ (where 0 should be understood as an act that pays 0 in every state).

In the Lo (1996) framework, which defines mixtures state-by-state, the mixture $\alpha r + (1 - \alpha)l$ will earn a payoff of α in state D and $1 - \alpha$ in state U. Therefore, mixing can provide a hedge against uncertainty: the strategy $\frac{1}{2}r + \frac{1}{2}l$ earns a payoff of $\frac{1}{2}$ in each state so that $\frac{1}{2}r + \frac{1}{2}l \sim \frac{1}{2} \succ l \sim r$.

Now, consider the mental state space constructed from Ω: $\Sigma = \{U, D, \Omega\}$. An application of Result 1 generates $\alpha^\nu_U = \alpha^\nu_D = 0$ and $\alpha^\nu_\Omega = 1$. The act l' pays 1 in state U and nothing in states D and Ω, and the act r' pays 1 in state D and nothing in states U and Ω. In this environment, mixing does not provide a hedge against uncertainty: the act $\frac{1}{2}r' + \frac{1}{2}l'$ earns $\frac{1}{2}$ in states U and D, but only earns 0 in state Ω. Given that $\alpha^\nu_\Omega = 1$ then $\frac{1}{2}r' + \frac{1}{2}l' \sim l' \sim r' \sim 0$. The existence of the additional state, Ω, in the mental state space implies that state-by-state mixing cannot provide a hedge against the uncertainty.

[9] Although this example is being presented solely as an individual decision maker problem, it should be clear that the setup could also describe a standard coordination game as viewed by a single player (the opponent's payoffs have been suppressed, but they are not material to the current discussion).