Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 197 (2021) 105326

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Mixed strategies and preference for randomization in games with ambiguity averse agents ☆

Evan M. Calford

John Mitchell Fellow, Research School of Economics, Australian National University, Canberra, Australia

Received 4 October 2020; final version received 28 June 2021; accepted 26 July 2021
Available online 3 August 2021

Abstract

We study the use of mixed strategies in games by ambiguity averse agents with a preference for randomization. Applying the decision theoretic model of Saito (2015) to games, we establish that the set of rationalizable strategies grows larger as preference for randomization weakens. An agent's preference for randomization is partially observable: given the behavior of an agent in a game, we can determine an upper bound on the strength of randomization preference for that agent. Notably, data in previous experiments on ambiguity aversion in games is not consistent with a maximal preference for randomization for approximately 30% of subjects.

JEL classification: D81; C70; C72

Keywords: Ambiguity aversion; Mixed strategies; Game theory

☆ I would like to acknowledge Pierpaolo Battigalli, Anujit Chakraborty, Adam Dominiak, Yoram Halevy and Terri Kneeland for helpful comments and suggestions, and Leeat Yariv for a particularly helpful suggestion. Seminar and conference audiences at the Wabash-Jordan workshop and the University of Queensland provided useful feedback. I also wish to thank three referees for very helpful comments and, particularly, Marciano Siniscalchi for identifying an error in an earlier draft.

E-mail address: Evan.Calford@anu.edu.au.

https://doi.org/10.1016/j.jet.2021.105326

	R	P	C
R	$6,3$	$2,x$	$0,0$
P	$x,2$	x,x	$x,2$
C	$0,0$	$2,x$	$3,6$

Fig. 1. Modified Battle of the Sexes game. $0 < x < 2$.

1. Introduction

For an agent in a normal form game the behavior of an opponent is, even in the case where the game form and payoffs are common knowledge, a source of ambiguity. Given the well documented prevalence of ambiguity aversion (see the review by Machina and Siniscalchi (2013)), understanding the role of, and response to, strategic ambiguity in decision making is a first-order problem in understanding strategic behavior. This paper studies the role of mixed strategies, and preferences for randomization, in determining the set of rationalizable outcomes for ambiguity averse agents in normal form games.

Mixed strategies are particularly important for ambiguity averse agents. Suppose that an agent has access to one action that performs well when her opponent plays Up but poorly when her opponent plays Down, and a second action that performs well against Down but poorly against Up. If the agent can mix between these two actions, the agent can generate a strategy that *on average* performs reasonably against either Up or Down. A mixed strategy can, therefore, form an ex-ante hedge against strategic ambiguity and, potentially, eliminate the effects of strategic ambiguity.

There are two stylized facts from the recent experimental economics literature that are particularly relevant: [1] that ambiguity averse subjects play games differently than ambiguity neutral subjects, with ambiguity averse subjects choosing 'safe' strategies more frequently, and [2] that subjects often display a persistent preference for randomization.[1] The previous literature on rationalizability with ambiguity averse agents can be split into two categories: The first is consistent with [1] but not [2] and the second is consistent with [2] but not [1]. In contrast, this paper presents a model that is consistent with both [1] and [2].

We assume that ambiguity preferences can be represented by Maxmin Expected Utility (Gilboa and Schmeidler, 1989).[2] That is, each agent holds a set valued belief containing all strategies that are (subjectively) considered feasible for the opponent, and then seeks to maximize the minimum utility against strategies in the belief set. Maxmin Expected Utility (MEU) has been previously applied to games in both theoretical (Lo, 2009) and experimental (Calford, 2020) settings and nests standard Expected Utility as a special case.

1.1. An illustrative example

Consider the modernized Battle of the Sexes game in Fig. 1. Rowena and Colin have agreed to go on a date to either location R or C. Rowena prefers the date occurs at location R, and Colin

[1] For example, Calford (2020) and Li et al. (2019) provide evidence that ambiguity averse and ambiguity neutral subjects behave differently in games, while Ivanov (2011) uses behavior in games to estimate ambiguity preferences. Agranov and Ortoleva (2017) document a persistent preference for randomization in lottery tasks, and Agranov et al. (2020) also find a preference for randomization in games. See also the literature discussed in Section 7.

[2] In Section 6 we briefly discuss extensions to other models of ambiguity preferences, including the smooth ambiguity model (Klibanoff et al., 2005).

prefers location C. Unfortunately, neither person can remember the location at which they agreed to meet. However, both Colin and Rowena own mobile phones and each may call the other party to confirm their plans (action P). Calling the other person ensures that the date goes ahead but also involves a cost as it reveals, embarrassingly, that the caller was not able to remember the plans for the date. The outcome of the phone call has some fixed utility, $0 < x < 2$, for the caller and a slightly higher utility, 2, for the receiver. If neither Rowena or Colin chooses to call the other party, then they risk attending differing locations and earning a payoff of 0. If both players have standard Expected Utility preferences then the set of rationalizable outcomes for this game is $\{(R, R), (R, C), (C, R), (C, C)\}$.[3]

Consider the literature that is consistent with statement [1] but not statement [2] above. This literature assumes that agents may only select pure actions, and that mixing is not feasible (see Epstein (1997), for example).[4] Clearly, R and C are justifiable for each player given they are best responses to R and C, respectively. The strategy P can be justified by the case where Rowena faces complete uncertainty about the strategy of Colin. Given Rowena's MEU preferences, she considers the worst case scenario for each of her strategies. If she plays R, she fears that Colin will play C and therefore assigns a utility of 0 to R. If she plays C, she fears that Colin will play R and therefore assigns a utility of 0 to C. P earns a safe payoff of $x > 0$ and is therefore the best response to complete uncertainty. All strategies are justifiable, and the entire game is therefore rationalizable.

Next, consider the literature that is consistent with statement [2] but not statement [1] above. This literature allows agents to play explicitly mixed strategies and assumes that agents have a maximal preference for mixing (see Chen and Luo (2012), for example).[5] Intuitively, the preference for mixing arises because the agent believes mixing eliminates the effects of ambiguity and, being ambiguity averse, the agent values the elimination of ambiguity. A maximal preference for mixing describes the case where the agent considers it feasible for mixing to fully eliminate the effects of ambiguity.

Consider again Rowena's strategy $\frac{1}{3}R + \frac{2}{3}C$. Given the assumption of a maximal preference for mixing, this strategy provides Rowena with a complete hedge against the uncertainty generated by Colin's behavior because it provides a constant expected payoff. If Colin plays R, this strategy earns, on average, a utility of 2. If Colin plays C, this strategy earns, on average, a utility of 2. And, of course, if Colin plays P, this strategy earns a utility of 2. $\frac{1}{3}R + \frac{2}{3}C$ therefore strictly dominates P, which pays only $x < 2$, and the set of rationalizable strategies is the same as for an Expected Utility agent.

[3] For Rowena the strategy $\frac{1}{3}R + \frac{2}{3}C$ strictly dominates P, and for Colin $\frac{2}{3}R + \frac{1}{3}C$ strictly dominates P, where $\frac{1}{3}R + \frac{2}{3}C$ denotes the mixed strategy that plays R with probability $\frac{1}{3}$ and C with probability $\frac{2}{3}$.

[4] Mixed strategies are given a beliefs interpretation in this literature: while each individual agent must use a pure strategy, the agent may still form beliefs about the relative proportions each pure strategy will be used by others.

[5] Formally, Chen and Luo (2012) consider a more general class of games where each player has a compact Hausdorff strategy space. In the special case where the strategy space is a simplex in a finite-dimensional Euclidian space, then the strategy space can be interpreted as the set of mixed strategies associated with a finite game. In this special case, Chen and Lou's assumption that preferences satisfy a "concave-like" condition implies a preference for mixing in the underlying finite game.

1.2. When does mixing provide a hedge against ambiguity?

Is it, however, reasonable to assume that $\frac{1}{3}R + \frac{2}{3}C$ provides Rowena with a complete hedge against ambiguity? Suppose that, for the sake of argument, a date at location R requires Rowena to wear a cocktail dress while location C requires casual attire. In order to implement her mixed strategy Rowena must resolve her randomization, then get dressed, and then attend either location R or C. Rowena is not sure, however, whether she will be making her decision before, or after, Collin makes his decision. If she believes that Collin makes his decision first then Rowena's randomization can provide a hedge against Collin's predetermined choice (i.e. at the time Rowena makes her decision, her action fully determines the payoff conditional on all past events). However, if Rowena believes that she makes her decision first, then her randomization does not provide a hedge against Collin's future decision: once she is dressed, Rowena is again exposed to ambiguity.[6]

Recent advances in decision theory (Saito, 2015; Ke and Zhang, 2020) argue that the degree to which $\frac{1}{3}R + \frac{2}{3}C$ provides Rowena with a hedge against ambiguity is subjective and equivalent to Rowena's strength of preference for randomization. The stronger Rowena's preference for randomization, the greater her subjective sense that $\frac{1}{3}R + \frac{2}{3}C$ generates a hedge against the uncertainty of Colin's strategy choice. Using the interpretation of Ke and Zhang (2020), Rowena's subjective beliefs regarding the timing of decisions determines her preference for randomization.[7] Here, we study the effects of the strength of preference for randomization on the set of rationalizable strategies.

Following Saito (2015), we denote Rowena's preference for randomization by $0 \leq \delta \leq 1$ where $\delta = 1$ implies that mixing provides a complete hedge against ambiguity and $\delta = 0$ implies that mixing provides no hedge against ambiguity. The utility of a mixed strategy is taken to be a linear combination of the utility provided under the case where mixing provides no hedge and the case where mixing provides a full hedge, with the weights afforded to each case determined by δ.[8] In our example, the utility of $\frac{1}{3}R + \frac{2}{3}C$, when Rowena faces complete uncertainty about Colin's behavior, is given by 2δ. It is, therefore, the case that $\frac{1}{3}R + \frac{2}{3}C$ strictly dominates P whenever $2\delta > x$. The set of rationalizable strategies is $\{(R, R), (R, C), (C, R), (C, C)\}$ when $\delta > \frac{x}{2}$ and the entire game when $\delta \leq \frac{x}{2}$.[9]

We, therefore, identify behavior that is consistent with stylized facts [1] and [2]. The safe strategy, P, is rationalizable when δ is small, agents have a preference for randomization whenever $\delta > 0$, and agents may use mixed strategies to hedge against strategic uncertainty when δ

6 Similar arguments regarding the timing of the resolution of uncertainty can be found in the decision theory literature. See Bade (2015), Baillon et al. (2019), Eichberger and Kelsey (1996) and Epstein et al. (2007) for examples.

7 To further illustrate the intuition of Ke and Zhang (2020) applied to games, consider the often used example of a soccer player shooting a penalty kick at a goalkeeper. The shooter may shoot to the left or right, and the goalkeeper must simultaneously decide whether to jump to the left or right. The shooter might believe that the goalkeeper is going to follow a pre-determined action that has been assigned by the goalkeeper's coach. Or, the shooter might believe that the goalkeeper will determine their own action at the instant of implementation. In the former situation, the shooter believes that randomization will hedge against strategic uncertainty, and in the latter situation the shooter believes randomization will not hedge. If the shooter believes that each scenario is possible, then the shooter will think that randomization will provide a hedge with some probability between 0 and 1.

8 See Section 2 for definitions and details, and Saito (2015) for an axiomatic foundation.

9 While in this example the dominance of P implies that P is not justifiable, we demonstrate in Example 1 below that dominance does not, in general, guarantee that an action is not justifiable.

is large. Further, δ is partially recoverable from observable behavior: if, for example, Rowena is observed to play P, then we must conclude that $\delta \leq \frac{x}{2}$.

1.3. The main result

In the main result of this paper we establish, for finite games, that the set of rationalizable strategies weakly increases as preference for randomization decreases. More formally, denote the set of rationalizable outcomes when agents have preference parameter δ as Z^δ. Corollary 1 states that $Z^\delta \subseteq Z^{\delta'}$ for $\delta \geq \delta'$. In the extreme case of $\delta = 1$, Z^1 coincides with the standard notion of correlated rationalizability (and Chen and Luo (2012)). On the other extreme, $\delta = 0$, Z^0 is consistent with Epstein (1997).

The intuition underlying this result is straightforward, with one complication. Because δ captures preference for mixing, the utility of a pure strategy is independent of δ and the utility of a mixed strategy is (weakly) increasing in δ. Therefore, as δ increases, fewer pure strategies remain as best responses. A complication arises, however, that it is possible for a strategy to be simultaneously dominated as a pure strategy but remain a component of a mixed best response. Nevertheless, it remains true that for any strategy that is dominated as a pure strategy at $\delta < 1$ there exists some δ', with $\delta \leq \delta' < 1$, such that, for all $\delta'' \geq \delta'$ every mix using that strategy is dominated. Hence, as δ increases, more strategies are dominated and less strategies are justifiable.[10]

The two most closely related papers to this one are Battigalli et al. (2016) and Chen and Luo (2012). Battigalli et al. (2016) show that the set of rationalizable strategies increases with ambiguity aversion in the smooth ambiguity aversion model (Klibanoff et al., 2005). Our result is distinct: here, the underlying ambiguity aversion is constant but the degree to which mixed strategies provide a hedge against ambiguity varies. Chen and Luo (2012) study rationalizability with a general class of preferences, but impose a maximal preference for mixing. This paper considers a more restricted class of preferences, but allows for weaker preferences for mixing. Other related literature is discussed in Section 7.

The paper proceeds as follows. Section 2 introduces Saito (2015) preferences in the context of games. Section 3 presents the main result. Section 4 addresses the observability of δ, and Section 5 considers the case of heterogeneous δ. Section 6 considers extensions of the main result, and Section 7 discusses the results within the context of the broader literature. Section 8 concludes.

2. Ambiguity aversion and preferences over randomization in games

We define a game by $G = \langle N, (A_i)_{i \in N}, (v_i)_{i \in N}, \delta \rangle$ where $N = \{1, \dots, n\}$ is the set of players, A_i is a finite set of actions for player i, $v_i : \times_{j \in N} A_j \to \mathbb{R}$ is a von-Neumann Morgenstern utility function over outcomes, and δ is a preference parameter defined below.[11] Let $A = \times_{j \in N} A_j$ be the set of action profiles and $A_{-i} = \times_{j \neq i \in N} A_j$ be the set of action profiles of players other than i.

[10] Lemma 3 from Pearce (1984), which states the equivalence of dominated and non-justifiable strategies, holds in a slightly modified form in our environment.

[11] For convenience, we assume that δ is homogeneous across players. We discuss the case of heterogeneous δ in Section 5.

We seek to apply the decision theoretic model of Saito (2015), which builds on the canonical Maxmin Expected Utility model of Gilboa and Schmeidler (1989). The model of Gilboa and Schmeidler (1989) posits that agents entertain a set of possible beliefs and calculate the utility of an action with respect to the belief that minimizes the utility of that action. Saito (2015) defines utility to be a weighted average of the Maxmin Expected Utility when randomization is assumed to not provide a hedge against uncertainty and the Maxmin Expected Utility when randomization does provide a hedge against uncertainty. While the model of Saito (2015) has not previously been applied in a game theoretic context, the application of decision theoretic frameworks to games has a rich and fruitful history that extends back to, at least, Tan and Werlang (1988).

To apply Saito (2015) to games, we suppose that A_{-i} forms the state space for player i, whose set of feasible acts is A_i. The (set of) beliefs for player i is given by $\Phi_i \subseteq \Delta(A_{-i})$, and we restrict Φ_i to be closed and convex. We denote the set of all closed and convex subsets of $\Delta(A_{-i})$ by $\mathcal{P}_i$ and note that $\mathcal{P}_i$ is a compact metric space with respect to the Hausdorff metric. A (mixed) strategy for player i is denoted by $\sigma_i \in \Delta(A_i)$, and we define $\sigma = \times_{i\in N}\sigma_i$ and $\Phi = \times_{i\in N}\Phi_i$.[12]

The utility of player i with preference parameter δ, beliefs Φ_i, and playing mixed strategy σ_i is given by:

$$U(\sigma_i, \Phi_i, \delta) = \delta u(\sigma_i) + (1-\delta)\sum_{a_i\in A_i}\sigma_i(a_i)u(a_i) \tag{1}$$

where

$$u(\sigma_i) = \min_{\phi\in\Phi_i}\sum_{a_{-i}\in A_{-i}}\phi(a_{-i})\sum_{a_i\in A_i}\sigma_i(a_i)v(a_i, a_{-i}) \tag{2}$$

and, with an abuse of notation,

$$u(a_i) = \min_{\phi\in\Phi_i}\sum_{a_{-i}\in A_{-i}}\phi(a_{-i})v(a_i, a_{-i}). \tag{3}$$

δ captures the agent's subjective beliefs regarding the extent to which mixed strategies provide a hedge against ambiguity: if $\delta = 1$ then randomization fully eliminates ambiguity, and if $\delta = 0$ then randomization provides no hedge against ambiguity.[13]

Intuitively, when $\delta = 0$ the agent has Maxmin expected utility preferences over the set of pure strategies, and mixed strategies are evaluated as a linear combination of the payoffs of the underlying pure strategies. Consequently, the agent will only play mixed strategies when they are indifferent between two (or more) pure strategies. When $\delta = 1$, the agent has convex preferences over mixed strategies and hence may strictly prefer a mixed strategy over any of the pure strategies in the support of that mixed strategy. The intermediate cases are taken to be a weighted average of the $\delta = 0$ and $\delta = 1$ cases, with the weight determined by δ.

We shall write $\Sigma_i^*(\Phi_i, \delta) = \{\sigma_i : \sigma_i \in \arg\max_{\sigma_i'\in\Delta(A_i)} U(\sigma_i', \Phi_i, \delta)\}$ to be the set of mixed strategies that are a best response to beliefs Φ_i. The set of mixed strategies that are a best response to some belief are then given by $\Sigma_i^*(\delta) = \bigcup_{\Phi_i\in\mathcal{P}_i}\Sigma_i^*(\Phi_i, \delta)$.

[12] Saito (2015) considers the primitives of behavior to be a set of feasible actions, rather than an explicit randomization, under the contention that only the outcomes of a randomization, and not the randomization itself, can be observed. Here we take mixed strategies to be the primitive of choice, but note that we do not at any stage require the mixed strategy to be directly observable. As is standard, the definition of rationalizability used here specifies the set of pure strategies that are rationalizable.

[13] The case of $\delta = 1$ is analogous to interpreting mixing as an ex-post randomization in the decision theoretic literature, and $\delta = 0$ is analogous to interpreting mixing as an ex-ante randomization.

	X	Y	Z
A	10	10	19
B	40	0	20
C	0	40	20
D	80	80	0

Fig. 2. An example game where column player payoffs have been suppressed for readability. Strategy A is both dominated and justifiable.

3. Rationalizability

In standard game theory with rational expected utility agents the set of pure strategy realizations that could by seen by an outside observer is the same as the set of pure strategies that are rationalizable. In the environment studied in this paper, this equivalence does not hold. Example 1, below, provides an example that illustrates the source of the problem: it is possible for a mixed strategy best response to place a positive weight on a dominated pure strategy.

Definition 1. A strategy σ_i is **dominated** by σ_i' if, for all $\Phi_i \subseteq \Delta(A_{-i})$, $U(\sigma_i', \Phi_i, \delta) > U(\sigma_i, \Phi_i, \delta)$.

Example 1. Consider the game in Fig. 2 where, for readability, we display only the payoffs of the row player. The strategy A is dominated by $\frac{B+C}{2}$ for $\delta > \frac{1}{2}$. Consider the case where the row player has complete uncertainty about the behavior of the column player, so that $\Phi_i = \Delta(\{X, Y, Z\})$, and fix $\delta = 0.6$. Basic calculations yield $U(\frac{B+C}{2}, \Phi_i, 0.6) = 12$, $U(A, \Phi_i, 0.6) = 10$ and $U(D, \Phi_i, 0.6) = 0$. Yet the mixed strategy $\frac{80A+9D}{89}$ produces a utility of $\frac{1520}{89} \approx 17$ at $\delta = 1$ and $\frac{800}{89} \approx 9$ at $\delta = 0$, and is a best response to Φ when $\delta = 0.6$.

There are three possible definitions of rationalizability that might be used in such an environment. First, rationalizability might be defined as the set of pure strategies that could reasonably be played. Second, rationalizability might be defined as the set of mixed strategies that could reasonably be played. Third, rationalizability might be defined as the support of the set of mixed strategies that could reasonably be played.

We reject the first possibility on the basis that an outside observer of the game could use this definition to erroneously infer that a player using a reasonable mixed strategy was behaving irrationally. We prefer the third definition given that it coincides with the set of pure strategy realizations that an outside observer might use to identify the rationality of a player. Readers who prefer the second definition can easily amend the definitions used below.

Definition 2. A strategy a_i is **justifiable with respect to beliefs** Φ_i if $a_i \in \text{supp}\big(\Sigma_i^*(\Phi_i, \delta)\big)$.

Definition 3. A strategy a_i is **justifiable** if $a_i \in \text{supp}\big(\Sigma_i^*(\delta)\big)$.

We define $\Sigma_i^0(\delta) = \Sigma_i^*(\delta)$, and then build the set of rationalizable strategies recursively. Let $\mathcal{P}_i^n$, representing the set of all feasible beliefs sets when each opponent plays a strategy in $\Sigma_j^n(\delta)$, be the set of all closed and convex subsets of $\Delta(A_{-i}|\Sigma_{-i}^n) = \{\sigma_{-i} \in \Delta(A_{-i}) : \forall j \neq i, \text{marg}_j(\sigma_{-i}) \in \text{co}(\Sigma_j^n(\delta))\}$ where co denotes the convex hull and $\text{marg}_j(\sigma_{-i})$ denotes the marginal distribution of σ_{-i} on A_j. Note that this definition implies that our definition of rationalizability is a form of correlated rationalizability as we do not impose independence of beliefs over opponent's strategies.[14] Finally, let $\Sigma_i^{n+1}(\delta) = \bigcup_{\Phi_i \in \mathcal{P}_i^n} \Sigma_i^*(\Phi_i, \delta)$.

Definition 4. The set of **rationalizable strategies** for player i is $Z_i^\delta = \text{supp}\big(\bigcap_{n=0}^{\infty} \Sigma_i^n(\delta)\big)$.

A further direct implication of Example 1 is that a key result of game theory with Expected Utility agents, often referred to as Pearce's lemma (Pearce, 1984), that a strategy is dominated if and only if it is not justifiable, does not hold in our environment. This suggests that a stricter notion of dominance will be useful.

We say that a (pure) strategy a_i is *totally dominated* if any mixed strategy that places a positive weight on a_i is dominated. We write $\Sigma_{a_i} = \{\sigma_i : \sigma_i(a_i) > 0\}$ to denote the set of mixed strategies for player i where the pure strategy a_i is played with a positive probability.

Definition 5. A strategy a_i is **totally dominated** if, for all $\sigma_{a_i} \in \Sigma_{a_i}$, σ_{a_i} is dominated.

Clearly, if a_i is totally dominated then a_i is dominated, and it follows from Example 1 that the converse does not hold universally.[15] Total dominance is, perhaps, not the most intuitive condition, particularly as each $\sigma_{a_i} \in \Sigma_{a_i}$ can be dominated by a different σ'. However, total dominance plays a critical role in our proof: total dominance respects Pearce's lemma.

Lemma 1. *A pure strategy a_i is totally dominated if and only if a_i is not justifiable.*

Proof. If a_i is totally dominated then, for all $\sigma_{a_i} \in \Sigma_{a_i}$, σ_{a_i} is dominated. Therefore σ_{a_i} is not a best response to any beliefs and a_i is not justifiable.

The proof of the 'only if' statement follows the proof of Lemma 3 in Pearce (1984) closely. Suppose that a_i is not justifiable. Then, for all $\sigma_{a_i} \in \Sigma_{a_i}$, there exists a function $b(\Phi_i) : \mathcal{P}_i \to \Delta(A_i)$ such that, for all Φ_i, $U(b(\Phi_i), \Phi_i, \delta) > U(\sigma_{a_i}, \Phi_i, \delta)$.

Consider the following pure strategy zero-sum game, played between players 1 and 2. The set of (pure) strategies for player 1 is $S_1 = \Delta(A_i)$ and the set of (pure) strategies for player 2 is $S_2 = \mathcal{P}_i$, with representative strategies denoted by s_1 and s_2 respectively. Note that the strategy sets for each player are complete metric spaces; for player 2, we can see this by defining a metric space with the Hausdorff metric on the collection of all compact sets of strategies in the original game (Henrikson, 1999; Barich, 2011). The payoffs for player 1 are given by $\overline{U}_1(s_1, s_2) = U(s_1, s_2, \delta) - U(\sigma_{a_i}, s_2, \delta)$ and the payoffs for player 2 are given by $\overline{U}_2(s_2, s_1) = -\overline{U}_1(s_1, s_2)$. $\overline{U}_1(s_1, s_2)$ inherits quasi-concavity in the first argument, and continuity, directly from the underlying preferences of the original game. Lemma 6 establishes that $\overline{U}_1(s_1, s_2)$ is linear in the second argument. Theorem 3.4 of Sion (1958) guarantees that the minimax theorem holds in pure strategies for this zero-sum game. Let (s_1^*, s_2^*) denote a pure strategy Nash equilibrium of this game. Then, for any $s_2 \in \mathcal{P}_i$:

$$\begin{aligned}\overline{U}_1(s_1^*, s_2) &\geq \overline{U}_1(s_1^*, s_2^*)\\ &\geq \overline{U}_1(b(s_2^*), s_2^*)\end{aligned}$$

[14] The distinction between correlated and independent rationalizability is discussed in both Brandenburger and Dekel (1987) and Tan and Werlang (1988).

[15] For the case of expected utility total dominance and dominance are, of course, equivalent. They are also equivalent at both $\delta = 1$ and $\delta = 0$, but may diverge for intermediate values of δ.

$$
\begin{aligned}
&> \overline{U}_1(\sigma_{a_i}, s_2^*) \\
&= U(\sigma_{a_i}, s_2^*, \delta) - U(\sigma_{a_i}, s_2^*, \delta) \\
&= 0
\end{aligned}
$$

Further, $\overline{U}_1(s_1^*, s_2) > 0$ for all $s_2 \in \mathcal{P}_i$ implies that $U(s_1^*, \Phi_i, \delta) > U(\sigma_{a_i}, \Phi_i, \delta)$ for all $\Phi_i \in \mathcal{P}_i$. Therefore, σ_{a_i} is dominated.

The above holds for all $\sigma_{a_i} \in \Sigma_{a_i}$. Therefore a_i is totally dominated. □

Next, we establish that if a strategy is dominated at some $\delta < 1$ then it is also dominated for all $\delta' > \delta$. To see this, notice that if σ_i is dominated for some δ then it is also dominated if we restrict the agent to hold only singleton beliefs.[16] Dominance with respect to singleton beliefs implies dominance at $\delta = 1$. Further, because Equation (1) is linear in δ, if a strategy is dominated for any two values of δ it must also be dominated for any intermediate value of δ.

Lemma 2. *For all $1 \geq \delta \geq \delta'$, if σ_i' dominates σ_i at δ' then σ_i' dominates σ_i at δ.*

Proof. First, we show that σ_i' dominates σ_i at $\delta = 1$. If σ_i' dominates σ_i at δ' then, in particular, for all *singleton* belief sets ϕ_i, $U(\sigma_i', \phi_i, \delta') > U(\sigma_i, \phi_i, \delta')$. For singleton belief sets, utility is constant in δ, so that $U(\sigma_i', \phi_i, 1) > U(\sigma_i, \phi_i, 1)$ for all singleton belief sets ϕ_i.

$U(\sigma_i', \phi_i, 1) > U(\sigma_i, \phi_i, 1)$ for all singleton belief sets implies σ_i' dominates σ_i for $\delta = 1$. To see this, note that for all Φ_i, writing $\underline{\phi}$ to denote the $\arg\min_{\phi \in \Phi_i} u(\sigma_i')$, we have the following:

$$
\begin{aligned}
U(\sigma_i', \Phi_i, 1) &= U(\sigma_i', \underline{\phi}, 1) \\
&> U(\sigma_i, \underline{\phi}, 1) \\
&\geq U(\sigma_i, \Phi_i, 1).
\end{aligned}
$$

Second, we show that σ_i' dominates σ_i at $\delta' \leq \delta \leq 1$. Utility is linear in δ. Therefore, for any σ^1, σ^2 if $U(\sigma^1, \Phi, 1) > U(\sigma^2, \Phi, 1)$ and $U(\sigma^1, \Phi, \delta') > U(\sigma^2, \Phi, \delta')$ then $U(\sigma^1, \Phi, \delta) > U(\sigma^2, \Phi, \delta)$ at all intermediate $\delta' \leq \delta \leq 1$. □

The proof of our key theorem follows directly from the previous lemmas.

Theorem 1. *Suppose that $\delta \geq \delta'$. If $a_i \in A_i$ is justifiable at δ then it is also justifiable at δ'.*

Proof. We prove the contrapositive. Suppose that a_i is not justifiable at δ'. Then, by Lemma 1, a_i is totally dominated at δ'. Lemma 2 implies that a_i is totally dominated at δ. Finally, Lemma 1 implies that a_i is not justifiable at δ. □

The cardinality of the set of justifiable strategies, and therefore cardinality of the set of rationalizable strategies, decreases in δ.

Corollary 1. *Suppose that $\delta \geq \delta'$. Then $Z^\delta \subseteq Z^{\delta'}$.*

Proof. Follows by application of the iterated definition of rationalizability. □

[16] That is, if a strategy is dominated for an agent with Saito preferences, it is dominated for an agent with Expected Utility preferences.

	X	Y	Z
A	10	10	19
B	40	0	20
C	0	40	20
D	80	80	0
E	11	11	11

Fig. 3. An extension of the game in Example 1.

To further understand the scope of this monotonocity result, consider the closely related case of an agent who always exhibits complete uncertainty about the behavior of an opponent: the agent only holds one feasible belief, $\Phi_i = \Delta(A_{-i})$. For this agent, the set of justifiable strategies is *not* monotonic in δ. That is, while the set of strategies that are *justifiable* is monotonic in δ, the set of strategies that are *justifiable with respect to* $\Phi_i = \Delta(A_{-i})$ is not monotonic in δ.

Example 2. Consider the game in Fig. 3 where, for readability, we display only the payoffs of the row player. We fix the row player's beliefs to be $\Phi_i = \Delta(\{X, Y, Z\})$. At $\delta = 1$ the unique best response to these beliefs is $\frac{B+C}{2}$, and at $\delta = 0$ the unique best response is E. As in Example 1, the best response at $\delta = 0.6$ is $\frac{80A+9D}{89}$. Because A is justifiable at $\delta = 0.6$, A must also be justifiable at $\delta = 0$ for some beliefs. Consider the set of beliefs $\Phi = \{\phi | \phi(Z) \geq 0.15\}$, for example.

3.1. The special cases of $\delta = 0$ and $\delta = 1$

There are several interesting points of intersection between the set of rationalizable set of strategies developed here, Z^δ, and previous results in the literature. The behavior of an agent with $\delta = 1$ is indistinguishable from the behavior of an Expected Utility agent. A version of this result can also be found in Chen and Luo (2012), and a decision theoretic analogue can be found in Kuzmics (2017).

Result 1. *$a_i \in Z^1$ if and only if a_i is rationalizable for a Subjective Expected Utility agent.*

Proof. First, we demonstrate that if strategy a_i is justifiable with respect to beliefs Φ_i then it is justifiable for a SEU agent.

When $\delta = 1$, preferences reduce to

$$U(\sigma_i, \Phi_i, 1) = \min_{\phi_i \in \Phi_i} \sum_{a_{-i} \in A_{-i}} \phi(a_{-i}) \sum_{a_i \in A_i} \sigma(a_i) v(a_i, a_{-i})$$
$$= \min_{\phi_i \in \Phi_i} \sum_{a_{-i} \in A_{-i}} \sum_{a_i \in A_i} \phi(a_{-i}) \sigma(a_i) v(a_i, a_{-i})$$

Let $\arg\max_{\sigma \in \Delta(A_i)} U(\sigma_i, \Phi_i, 1) = \sigma_i^*$, and let ϕ_i^* be the utility minimizing belief for $U(\sigma_i^*, \Phi_i, 1)$. Clearly, σ_i^* is also the best response for an agent with Expected Utility preferences and belief ϕ_i^*.

Next we show that if a_i is justifiable for an SEU agent then it is justifiable with respect to some beliefs Φ_i. Suppose that σ^{**} is the best response for an Expected Utility agent with beliefs ϕ_i^{**}. Let $\Phi_i^{**} = \{\phi_i^{**}\}$. Then $\sigma^{**} \in \arg\max_{\sigma_i \in \Delta(A_i)} U(\sigma_i, \Phi_i^{**}, 1)$.

Finally, equivalence of the sets of justifiable strategies implies equivalence of the set of correlated rationalizable strategies. $\square$

	X (0.66)	Y (0.34)
A (0.63)	25, 20	14, 12
B (0.09)	14, 20	25, 12
C (0.28)	18, 12	18, 22

Fig. 4. Example from Calford (2020). The proportion of subjects choosing each strategy is given in parentheses. Payoffs are in Canadian dollars.

The special case of $\delta = 0$ is also found in the prior literature (Epstein, 1997).[17] Thus, one of the contributions in the literature is providing a connection between the standard notion of rationalizability and the notion developed in Epstein (1997).

Result 2. *$a_i \in Z^0$ if and only if a_i is Epstein (1997) rationalizable.*

Proof. Immediate. □

4. Observability of model parameters

Theorem 1 implies that δ is partially observable. Specifically, suppose that there exists a strategy a_i that is justifiable for all $\delta \leq \delta_1$ but not justifiable for $\delta > \delta_1$. Then, if we observe a_i being played, and assume that the agent is rational (in the sense of maximizing utility as defined in Equation (1)), we must conclude that $\delta \leq \delta_1$. In addition, suppose that a_i' is rationalizable for all $\delta \leq \delta_1$ but not rationalizable for $\delta > \delta_1$. Then, if we observe a_i' being played, and assume rationality and common knowledge of rationality, we must conclude that $\delta \leq \delta_1$.

To illustrate, consider the following game in Fig. 4 from Calford (2020) where, for now, we assume that agents exhibit risk neutrality. We discuss the relaxation of risk neutrality below, after illustrating the procedure. The proportion of subjects choosing each pure strategy is given in parentheses.

Example 3. In this game, C is justifiable if and only if $\delta \leq \frac{8}{11}$ and B is rationalizable if and only if $\delta \leq \frac{8}{11}$. Therefore, if we assume rationality we can conclude that at least 28% of subjects have $\delta \leq \frac{8}{11}$ and if we assume common knowledge of rationality then we can conclude that at least 37% of subjects have $\delta \leq \frac{8}{11}$. We can conclude, given the stated assumptions, that *at most* 63% of subjects have $\delta = 1$. Of course, this proportion could be considerably lower, given that (A, X) is rationalizable even for subjects with $\delta = 0$.

Other experimental studies produce similar conclusions.

Example 4. Consider the game in Fig. 5, from Kelsey and le Roux (2015). Strategy R is never a best response whenever $\delta > \frac{4}{5}$. Given that 30% of subjects play R we can conclude, given an assumption of rationality, that at least 30% of subjects have $\delta \leq \frac{4}{5}$.

We now address the role of risk preferences. Generally, outside of a special case discussed below, we are only able to jointly identify risk preferences and δ. To illustrate, we shall assume that

[17] In the following, we use the shorthand "Epstein (1997) rationalizability" to refer to the specific application of Epstein's general formulation to the case of Maxmin Expected Utility agents.

	L (0.40)	M (0.30)	R (0.30)
U (0.50)	0, 0	6, 2	1, 1.2
D (0.50)	2, 6	0, 0	1.1, 1.2

Fig. 5. Example from Kelsey and le Roux (2015). The proportion of subjects choosing each strategy is given in parentheses. Payoffs are in British Pounds.

agents have a constant relative risk aversion utility function over monetary outcomes such that, with an abuse of notation, $v(a_i, a_{-i}) = v(m(a_i, a_{-i})) = \frac{(m(a_i,a_{-i})+w)^{1-\rho}-1}{1-\rho}$ where $m(a_i, a_{-i})$ is the monetary outcome associated with the outcome (a_i, a_{-i}) and w is background wealth.[18] To conserve on notation, we will typically write $v(m)$ instead of $v(m(a_i, a_{-i}))$. Given that Z^1 is indistinguishable from the set of rationalizable strategies for an ambiguity neutral agent, in many applications we are interested in identifying whether $\delta < 1$ or $\delta = 1$. That is, when can we conclude that an agent's behavior reveals ambiguity aversion – an event that can only occur when the agent has a less than complete preference for randomization?

Many experimental studies, including Calford (2020), include measures of risk aversion as controls. Fig. 6 depicts the preference for randomization, δ, on the vertical axis and risk aversion, ρ, on the horizontal axis and displays the set of rationalizable strategies at various parameter combinations. The figure depicts the function $\delta = \min(\frac{v(18)-v(14)}{v(l)-v(14)}, 1)$ where $v(l) = \frac{v(25)+v(14)}{2}$, which delineates the boundary at which the set of rationalizable strategies for the game in Fig. 4 changes: below this function all strategies are rationalizable, while above the function only $\{A, X\}$ is rationalizable. Using data from an auxillary risk measurement task, Calford (2020) can identify whether subjects have a CRRA risk aversion parameter that is above, or below, $\rho = 1.14$ (the vertical line in the figure).[19]

If we observe a subject who has risk aversion of $\rho < 1.14$, and the subject does not play A, then we can conclude that δ must be bounded away from 1 in the lower left hand region of the figure. Calford (2020) finds that, for subjects with $\rho < 1.14$, 70% of subjects play A. Thus, controlling for risk preferences, albeit in a crude manner, does not materially affect our conclusions about the proportion of subjects with $\delta < 1$.

Kelsey and le Roux (2015) does not include data on the risk preferences of subjects, yet we can still claim that $\delta < 1$ for subjects that play R by reflecting on the degree of risk aversion that we might expect subjects to hold. For example, the canonical Holt and Laury (2002) multiple price list method for measuring risk preferences is only designed to identify risk preference up to $\rho = 1.37$, describing higher degrees of risk aversion as "stay in bed" preferences. Fig. 7 depicts the preference for randomization, δ, on the vertical axis and risk aversion, ρ, on the horizontal axis and displays the set of rationalizable strategies at various parameter combinations. The boundary between the two regions in the figure is given by $\delta = \min(\frac{v(1.2)-v(0)}{v(l)-v(0)}, 1)$, where $v(l) = pv(2) + (1-p)v(0)$ and $p = \frac{v(6)-v(0)}{v(6)+v(2)-v(0)}$. It is clear from the figure that for any subject with $\rho < 1.37$, a choice of R implies that $\delta < 1$.

In a restricted class of games we can partially identify δ without any assumptions regarding risk preferences. Denote the set of strategies that can be rationalized by an Expected Utility agent, as identified by Börgers (1993), by B. That is, B contains all strategies that can be rationalized under any monotonic transformation of the payoffs of a game.

[18] It is well known that CRRA utility estimates are affected by assumptions regarding the integration of wealth into the utility function. In what follows we assume that subjects integrate all wealth earned during the experiment (e.g. show up fees, payment from other tasks) into the utility function, but do not integrate wealth from outside the experiment.

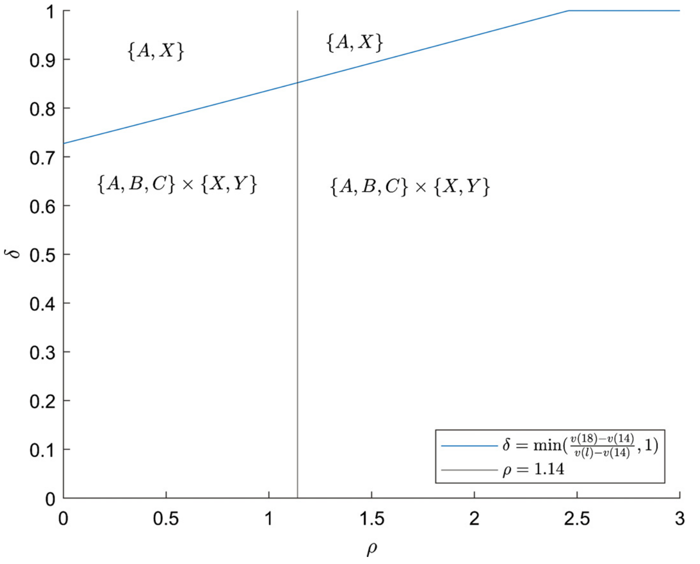

Fig. 6. Set of rationalizable strategies as a function of δ and ρ for the game in Fig. 4 from Calford (2020). The displayed functions are $\delta = \min(\frac{v(18)-v(14)}{v(l)-v(14)}, 1)$ and $\rho = 1.14$, where $v(l) = \frac{v(25)+v(14)}{2}$. Utility is calculated assuming that subjects integrate a C\$5 show-up fee into their utility function. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

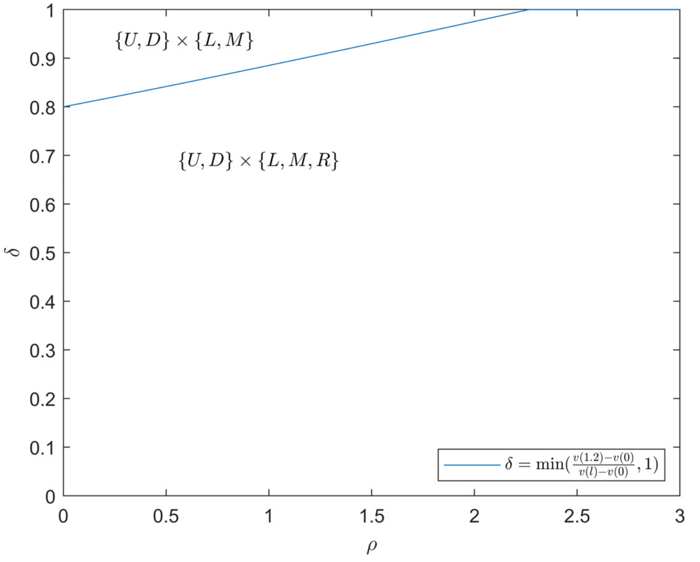

Fig. 7. Set of rationalizable strategies as a function of δ and ρ for the game in Fig. 5 from Kelsey and le Roux (2015). The displayed function is $\delta = \min(\frac{v(1.2)-v(0)}{v(l)-v(0)}, 1)$, where $v(l) = pv(2) + (1-p)v(0)$ and $p = \frac{v(6)-v(0)}{v(6)+v(2)-v(0)}$. Utility is calculated assuming that subjects integrate a £5 show-up fee, plus expected earnings from other tasks of £2, into their utility function.

The relationship between Z^δ and B is modulated by the relationship between two classes of dominance for Expected Utility agents.

Definition 6. A pure strategy a_i' is **strictly dominated** by σ_i for an Expected Utility agent if, for all $\phi \in \Delta(A_{-i})$,

$$\sum_{a_{-i}\in A_{-i}} \phi(a_{-i}) \sum_{a_i \in A_i} \sigma_i(a_i)v(a_i, a_{-i}) > \sum_{a_{-i}\in A_{-i}} \phi(a_{-i})v(a_i', a_{-i}).$$

Definition 7. A pure strategy a_i' is **weakly dominated** by σ_i for an Expected Utility agent if, for all $\phi \in \Delta(A_{-i})$,

$$\sum_{a_{-i}\in A_{-i}} \phi(a_{-i}) \sum_{a_i \in A_i} \sigma_i(a_i)v(a_i, a_{-i}) \geq \sum_{a_{-i}\in A_{-i}} \phi(a_{-i})v(a_i', a_{-i})$$

with a strict inequality for at least one $\phi \in \Delta(A_{-i})$.

We say that strict dominance and weak dominance coincide whenever the set of pure strategies that are strictly dominated and the set of pure strategies that are weakly dominated are identical.

Lemma 3. *In all games where strict dominance and weak dominance coincide, $Z^1 \subseteq Z^0 \subseteq B$.*

Proof. Suppose that a_i is dominated in the sense of Börgers (1993), so that $a_i \notin B$. Then, when strict dominance and weak dominance coincide, a_i is strictly dominated by a pure strategy (Börgers, 1993). Therefore a_i is dominated in the sense of Definition 1 for all δ. It follows that a_i is also totally dominated for all δ. To see this, take any mixed strategy that places positive weight on a_i and replace the weight on a_i with the dominating pure strategy (noting that if, instead, a_i was dominated by a mixed strategy this replacement argument would not be valid). Therefore, $a_i \notin Z^0$. □

If strict dominance and weak dominance do not coincide, however, it is possible to partially identify δ without any assumption on risk preferences. That is, there exist games where $B \subset Z^0$ and, in this case, the observation of a strategy a_i such that $a_i \in Z^0$ but $a_i \notin B$ implies that $\delta < 1$. To see this consider the game in Fig. 8, which is a simplified version of an example found in Epstein (1997).

Example 5. In the game in Fig. 8 the strategy M is weakly dominated by both T and D, but is not strictly dominated by either pure strategy. M is also dominated in the sense of Börgers (1993) and, because of this, it is dominated for an EU agent for any monotonic transformation of the payoffs.[20] Therefore, M is not justifiable for any EU agent. It is also the case that $M \in Z^0$ but $M \notin Z^\delta$ for $\delta > 0$. Therefore, if we observe an agent playing M we can conclude that $\delta = 0$ for that agent without making any assumptions about risk preferences.

Allowing for outside wealth to be integrated into the utility function would reduce the threshold δ values outlined below; as w increases the effects of curvature of the utility function are reduced.

19 This calculation assumes that $w = 5$ to reflect the \$5 show-up fee paid to subjects. In Calford (2020) subjects were paid for one randomly chosen task, so that there is no need to account for expected earnings from other tasks within the experiment.

20 Clearly, $\frac{T+D}{2}$ dominates M for any monotonic transformation of the row player's payoffs.

	L	R
T	1, 0.9	2, 1
M	1, 100	1, 1
D	2, 0.9	1, 1

Fig. 8. Simplified example from Epstein (1997).

5. Heterogeneous preferences for randomization

In previous sections we considered only a homogeneous preference for randomization across all players; here we consider the case of heterogeneous preferences for randomization.

The main results, including Theorem 1, all go through with heterogeneous preferences for randomization. Suppose that player i has preference parameter δ_i. Then the set of rationalizable strategies weakly increases as δ_i decreases, holding δ_j constant for $j \neq i$.

Changes in δ for different players may have drastically different effects on the set of rationalizable strategies, however. Consider again the game in Fig. 4. Suppose that the row player has preference for randomization δ_R, and the column player δ_C. The rationalizable set is $\{A, X\}$ if $\delta_R > \frac{8}{11}$ and the entire game if $\delta_R \leq \frac{8}{11}$. Importantly, note that the rationalizable set is independent of δ_C.

To see this notice that, for any value of δ_C, Y is rationalizable only if C is rationalizable. For the row player, C is justifiable with respect to $\Phi_R = \{\phi : 0 \leq \phi(X) \leq 1\}$ if $\delta_R \leq \frac{8}{11}$ and C is dominated by $\sigma(A) = \sigma(B) = \frac{1}{2}$ if $\delta_R > \frac{8}{11}$. Further, when C is justifiable then all strategies are justifiable and the entire game is rationalizable. When C is not justifiable then iterated elimination implies that Y is not rationalizable which implies that B is not rationalizable.

6. Extensions to other preference models

The proof of Theorem 1 depends on relatively few details of the underlying utility function, suggesting that we can replace the underlying MEU utility with an alternative model of ambiguity aversion. Here, we sketch the conditions on utility functions for which Theorem 1 will hold, and discuss the specific application of smooth ambiguity aversion (Klibanoff et al., 2005).

In general we can isolate two versions of the appropriate utility function, $u_{ep}(\sigma, \theta)$ and $u_{ea}(\sigma, \theta)$, representing utility over ex-post randomizations and ex-ante randomizations respectively, where σ denotes a mixed strategy and θ denotes an arbitrary belief (which may be a capacity, set of beliefs, second order distribution etc., depending on the model in question), and the set of all feasible beliefs is given by Θ. We can then write the general utility function as

$$U(\sigma, \theta, \delta) = \delta u_{ep}(\sigma, \theta) + (1 - \delta) u_{ea}(\sigma, \theta).$$

Given this, the following conditions are sufficient for Theorem 1 to hold:

1. There exists an appropriate restriction on beliefs, $\Theta' \subset \Theta$ such that
   (a) for all $\theta' \in \Theta'$ $u_{ep}(\sigma, \theta') = u_{ea}(\sigma, \theta')$; and
   (b) for all $\theta \in \Theta$, there exists a $\theta' \in \Theta'$ such that $u_{ep}(\sigma, \theta') = u_{ep}(\sigma, \theta)$.
2. The expression $\overline{U}(\sigma, \theta) = U(\sigma, \theta, \delta) - U(\sigma', \theta, \delta)$ is upper semi-continuous in the first argument and lower semi-continuous in the second argument and either
   (a) quasi-concave in the first argument and quasi-convex in the second argument, or

(b) concave-like in the first argument and convex-like in the second argument.[21]

For smooth ambiguity preferences Θ is the set of all second order probability distributions, and utility is linear in such beliefs. Condition 1(a) can be satisfied by restricting second order beliefs to be degenerate, in which case the preferences collapse back to Expected Utility preferences. If we restrict agents to be both weakly ambiguity averse and weakly risk averse then preferences are concave and, therefore, condition 1(b) also holds (Chen and Luo, 2012). The expression $\overline{U}(\sigma, \theta)$ is clearly continuous in both arguments, linear in the second argument, and, given the restriction to weak ambiguity aversion and weak risk aversion, concave in the first argument.

7. Discussion

There is a close relationship between the results presented in this and several recent papers that have established comparative statics results for ambiguity aversion in games. To understand the relationship, consider that preferences for randomization may be interpreted as an ability to build hedges against ambiguity. To an outside observer, as my preference for randomization increases I appear as if I am less ambiguity averse: my ability to hedge implies that I am less affected by the presence of ambiguity and, therefore, ambiguity has a lesser effect on my behavior. Given that we interpret rationalizable strategies as the set of strategies that an outside observer might observe, the close relationship between preference for ambiguity and preference for randomization seems natural.

The comparative statics results in the existing literature are neither special cases nor generalizations of the results presented here. The existing literature focuses on the special cases of $\delta = 0$ and $\delta = 1$, while this paper considers all $0 < \delta < 1$. On the other hand, the existing literature varies the strength of ambiguity aversion while this paper holds ambiguity aversion constant (and maximal).[22]

Several comparative static results for the smooth ambiguity model already exist in the literature. If $\delta = 1$, then ambiguity preference has no effect on the set of rationalizable strategies (Chen and Luo, 2012). If $\delta = 0$, then the set of rationalizable strategies increases as ambiguity aversion increases (Battigalli et al., 2016).[23] This paper provides a connection between these two results, by applying the extension discussed in Section 6: for a fixed degree of ambiguity aversion, the set of rationalizable strategies decreases in δ with the special case at $\delta = 1$ coinciding with Chen and Luo (2012) and the special case at $\delta = 0$ coinciding with Battigalli et al. (2016).

[21] $\overline{U}(\sigma, \theta)$ is concave-like in the first argument if for every $\sigma' \in \Delta(A_i)$ and $\sigma'' \in \Delta(A_i)$ and $0 \leq t \leq 1$ there exists $\sigma \in \Delta(A_i)$ such that $t\overline{U}(\sigma', \theta) + (1-t)\overline{U}(\sigma'', \theta) \leq \overline{U}(\sigma, \theta)$ for all $\theta \in \Theta$. The definition of convex-like reverses the sign of the inequality.

[22] For a fixed set of beliefs, the MEU model displays the greatest possible ambiguity aversion over those beliefs. However, the degree of ambiguity aversion exhibited by an MEU agent, as measured by the certainty equivalent of a prospect, varies as beliefs change. Nevertheless, for each set of beliefs, the MEU model can be formulated as the limit of the smooth ambiguity model as ambiguity aversion becomes infinite.

[23] The Battigalli et al. (2016) result is established for a more general set of games than the finite games considered in this paper. Battigalli et al. (2016) also take a conceptually different approach to mixed strategies than used in this paper. The authors argue that any mixed strategy which might provide a hedge against ambiguity should instead be defined explicitly in the set of available strategies. Any mixed strategy that is not defined in this manner is assumed to not provide any hedge against ambiguity. Thus, in the language of the current paper, they assume $\delta = 1$ for a fixed set of mixed strategies and $\delta = 0$ for any other mixed strategy.

Hanany et al. (2020) consider equilibrium in extensive form games with incomplete information for agents with smooth ambiguity preferences. They establish several distinct comparative statics. First, if priors are held fixed, then the set of equilibrium strategies can shift arbitrarily as ambiguity aversion changes. Example 2 demonstrates a similar effect in the current environment: if beliefs are held fixed, then best response correspondences can vary arbitrarily in δ. Second, if a common prior restriction is imposed, but beliefs are otherwise allowed to vary, increasing ambiguity aversion increases the set of equilibrium outcomes. There is no comparable result in the current framework. Third, if the common prior restriction is removed, changes in ambiguity aversion have no effect on the equilibrium set. Finally, if agents are restricted to pure strategies, then increases in ambiguity aversion increase the set of equilibrium strategies. These final two cases have parallels to the rationalizability results in the prior literature, discussed above, for the $\delta = 1$ and $\delta = 0$ cases respectively.[24]

Also related is Weinstein (2016) who studies the comparative statics of risk aversion in games. In particular, Weinstein (2016) establishes that the set of rationalizable strategies increases as risk aversion increases and that in the limit as risk aversion becomes infinite the set of rationalizable strategies is the set B as identified by Börgers (1993). The relationship between B and Z^δ is formalized in Lemma 3.

Liu et al. (2020) study the role of randomization in exchange economies with MEU agents. There is a close mapping between the results demonstrated here and those in Liu et al. (2020), despite the differences in the structure of the environment. Liu et al. (2020) show that the set of feasible allocations when mixed reports are allowed and $\delta = 0$ is identical to the set of feasible allocations when only pure reports are allowed. Further, they show that the set of feasible allocations at $\delta = 1$ is smaller than at any $\delta < 1$ including, in particular, $\delta = 0$. If beliefs are restricted to represent complete uncertainty, which they refer to as Wald's maxmin preferences, then every efficient allocation is feasible for any $\delta \in [0, 1]$.[25]

Recent experimental results, in the domain of individual decision making and using Ellsberg urns as the source of ambiguity, have investigated whether human subjects hold a preference for randomization or not. Kuzmics et al. (2020) conclude that a proportion of subjects ($\sim 35\%$) appear to have a preference for randomization devices, but an inability to randomize without a randomization device. They also find that a smaller proportion ($\sim 23\%$) refuse to use a randomization device even after the hedging properties of the randomization are explained.[26] The remaining subjects are either Expected Utility maximizers or exhibit a preference for randomization. Dominiak and Schnedler (2011) provide evidence that only a minority of ambiguity averse subjects have a preference for randomization. Oechssler et al. (2016) find that timing of the resolution of uncertainty has no effect on the prevalence of hedging. In direct contrast, Spears (2009) finds that subjects are more likely to use a randomization device when subjective uncertainty is resolved prior to objective uncertainty. In addition, Kuzmics et al. (2020) find that many subjects are sensitive to advice on how randomization should be used to make decisions suggesting that

[24] Note that the restriction to pure strategies is effectively an assumption that $\delta = 0$: it is not possible to use mixed strategies to construct a hedge if mixed strategies are unavailable.

[25] It is well known, in the Expected Utility case, that not all efficient allocations are feasible. Thus, the restriction to complete uncertainty plays an important role by ruling out singleton belief sets (which return the economy to the EU case).

[26] The objective randomization was resolved after the subjective uncertainty, and the authors interpret this as a violation of the motivation for the models given in Saito (2015) and Ke and Zhang (2020). A purely subjective interpretation of these models could be viewed as consistent with the data.

“preferences” in this domain may not be strongly held. Taken in aggregate, this branch of the literature suggests substantial variation in δ across subjects.

8. Conclusion

In this paper, we apply a model with ambiguity averse preferences and partial preference for randomization (Saito, 2015) to game theory. We prove that the set of rationalizable strategies increases (in the sense of set inclusion) as preference for randomization decreases. Using carefully designed games, such as those studied in Calford (2020), it is possible to partially identify the preference for randomization of an ambiguity averse subject in a game theoretic experiment.

Appendix A. Auxiliary lemmas

Consider the function $\overline{U}_1 : \Delta(A_i) \times \mathcal{P}_i \to \mathbb{R}$ given by $\overline{U}_1(\sigma, \Phi) = U(\sigma, \Phi, \delta) - U(\sigma', \Phi, \delta)$, for some fixed σ' and δ. We prove that $\overline{U}_1(\sigma, \Phi)$ is linear in the second argument.

Take two arbitrary elements $\Phi_A, \Phi_B \in \mathcal{P}_i$ and consider $\Phi_\alpha = \alpha\Phi_A + (1-\alpha)\Phi_B \in \mathcal{P}_i$.[27]

Lemma 4. $U(\sigma, \Phi_\alpha, \delta) \le \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta)$

Proof. We begin with the $\delta = 1$ case, such that utility is given by Equation (2). Write ϕ_A^* and ϕ_B^* for the belief that minimizes the utility functions $U(\sigma, \Phi_A, \delta)$ and $U(\sigma, \Phi_B, \delta)$, respectively. Let $\phi_\alpha = \alpha\phi_A^* + (1-\alpha)\phi_B^*$, and note that $\phi_\alpha \in \Phi_\alpha$. Then we have

$$\begin{aligned} U(\sigma, \Phi_\alpha, 1) &\le U(\sigma, \phi_\alpha, 1) \\ &= U(\sigma, \alpha\phi_A^* + (1-\alpha)\phi_B^*, 1) \\ &= \alpha U(\sigma, \phi_A^*, 1) + (1-\alpha)U(\sigma, \phi_B^*, 1) \\ &= \alpha U(\sigma, \Phi_A, 1) + (1-\alpha)U(\sigma, \Phi_B, 1) \end{aligned}$$

where the third line follows from linearity of the utility function for singleton beliefs.

We continue with the $\delta = 0$ case. Write $u(a_i, \Phi)$ to denote the utility of a pure strategy a_i in Equation (3), where the additional argument emphasizes the beliefs with which the utility is being calculated. By the same argument as above, $u(a_i, \Phi_\alpha) \le \alpha u(a_i, \Phi_A) + (1-\alpha)u(a_i, \Phi_B)$ for all $a_i \in A_i$. Therefore, $U(\sigma, \Phi_\alpha, 0) \le \alpha U(\sigma, \Phi_A, 0) + (1-\alpha)U(\sigma, \Phi_B, 0)$.

Piecing the two cases together, we have that $U(\sigma, \Phi_\alpha, \delta) \le \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta)$ for all δ. □

Lemma 5. $U(\sigma, \Phi_\alpha, \delta) \ge \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta)$.

Proof. We begin with the $\delta = 1$ case, such that utility is given by Equation (2). Write ϕ_α^* for the belief that minimizes the utility function $U(\sigma, \Phi_\alpha, \delta)$. There exists a $\phi_A \in \Phi_A$ and $\phi_B \in \Phi_B$ such that $\phi_\alpha^* = \alpha\phi_A + (1-\alpha)\phi_B$. We have

[27] The linear combination is to be taken element by element. That is, $\Phi_\alpha = \{\phi \in \Delta(A_{-i}) : \exists\phi_A \in \Phi_A, \exists\phi_B \in \Phi_B, \phi = \alpha\phi_A + (1-\alpha)\phi_B\}$.

$$
\begin{aligned}
U(\sigma, \Phi_\alpha, 1) &= U(\sigma, \phi^*_\alpha, 1) \\
&= U(\sigma, \alpha\phi_A + (1-\alpha)\phi_B, 1) \\
&= \alpha U(\sigma, \phi_A, 1) + (1-\alpha)U(\sigma, \phi_B, 1) \\
&\geq \alpha U(\sigma, \Phi_A, 1) + (1-\alpha)U(\sigma, \Phi_B, 1)
\end{aligned}
$$

where the third line follows from linearity of the utility function for singleton beliefs.

We continue with the $\delta = 0$ case. Write $u(a_i, \Phi)$ to denote the utility of a pure strategy a_i in Equation (3), where the additional argument emphasizes the beliefs with which the utility is being calculated. By the same argument as above, $u(a_i, \Phi_\alpha) \geq \alpha u(a_i, \Phi_A) + (1-\alpha)u(a_i, \Phi_B)$ for all $a_i \in A_i$. Therefore, $U(\sigma, \Phi_\alpha, 0) \geq \alpha U(\sigma, \Phi_A, 0) + (1-\alpha)U(\sigma, \Phi_B, 0)$.

Piecing the two cases together, we have that $U(\sigma, \Phi_\alpha, \delta) \geq \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta)$ for all δ. □

Lemma 6. $\overline{U}_1(\sigma, \Phi)$ *is linear in the second argument.*

Proof. Lemma 4 and Lemma 5 imply that $U(\sigma, \Phi_\alpha, \delta) = \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta)$ for all σ. Therefore we have

$$
\begin{aligned}
\overline{U}_1(\sigma, \Phi_\alpha) &= U(\sigma, \Phi_\alpha, \delta) - U(\sigma', \Phi_\alpha, \delta) \\
&= \alpha U(\sigma, \Phi_A, \delta) + (1-\alpha)U(\sigma, \Phi_B, \delta) \\
&\quad - \alpha U(\sigma', \Phi_A, \delta) - (1-\alpha)U(\sigma', \Phi_B, \delta) \\
&= \alpha\big[U(\sigma, \Phi_A, \delta) - U(\sigma', \Phi_A, \delta)\big] + (1-\alpha)\big[U(\sigma, \Phi_B, \delta) - U(\sigma', \Phi_B, \delta)\big] \\
&= \alpha\overline{U}_1(\sigma, \Phi_A) + (1-\alpha)\overline{U}_1(\sigma, \Phi_B) \quad \square
\end{aligned}
$$

References

Agranov, Marina, Ortoleva, Pietro, 2017. Stochastic choice and preferences for randomization. J. Polit. Econ. 125 (1), 40–68.

Agranov, Marina, Healy, Paul J., Nielsen, Kirby, 2020. Non-random randomization. Mimeo. February.

Bade, Sophie, 2015. Randomization devices and the elicitation of ambiguity-averse preferences. J. Econ. Theory 159, 221–235. Part A.

Baillon, Aurelien, Halevy, Yoram, Li, Chen, 2019. Experimental elicitation of ambiguity attitude using the random incentive system. Mimeo.

Barich, Katie, 2011. Proving completeness of the Hausdorff induced metric space. Mimeo.

Battigalli, P., Cerreia Vioglio, S., Maccheroni, F., Marinacci, M., 2016. A note on comparative ambiguity aversion and justifiability. Econometrica 84 (5), 1903–1916.

Börgers, Tilman, 1993. Pure strategy dominance. Econometrica 61 (2), 423–430.

Brandenburger, Adam, Dekel, Eddie, 1987. Rationalizability and correlated equilibria. Econometrica 55 (6), 1391–1402.

Calford, Evan M., 2020. Uncertainty aversion in game theory: experimental evidence. J. Econ. Behav. Organ. 176, 720–734.

Chen, Yi-Chun, Luo, Xiao, 2012. An indistinguishability result on rationalizability under general preferences. Econ. Theory 51, 1–12.

Dominiak, Adam, Schnedler, Wendelin, 2011. Attitudes toward uncertainty and randomization: an experimental study. Econ. Theory 48, 289–312.

Eichberger, Jurgen, Kelsey, David, 1996. Uncertainty aversion and preference for randomization. J. Econ. Theory 71, 31–43.

Epstein, Larry G., 1997. Preference, rationalizability and equilibrium. J. Econ. Theory 73, 1–29.

Epstein, Larry G., Marinacci, Massimo, Seo, Kyoungwon, 2007. Coarse contingencies and ambiguity. Theor. Econ. 2, 355–394.

Gilboa, Itzhak, Schmeidler, David, 1989. Maxmin expected utility with non-unique prior. J. Math. Econ. 18, 141–153.
Hanany, Eran, Klibanoff, Peter, Mukerji, Sujoy, 2020. Incomplete information games with ambiguity averse players. Am. Econ. J. Microecon. 12 (2), 135–187.
Henrikson, Jeff, 1999. Completeness and total boundedness of the Hausdorff metric. MIT Undergrad. J. Math. 1, 69–80.
Holt, Charles A., Laury, Susan K., 2002. Risk aversion and incentive effects. Am. Econ. Rev. 92 (5), 1644–1655.
Ivanov, Asen, 2011. Attitudes to ambiguity in one-shot normal-form games: an experimental study. Games Econ. Behav. 71, 366–394.
Ke, Shaowei, Zhang, Qi, 2020. Randomization and ambiguity aversion. Econometrica 88 (3), 1159–1195.
Kelsey, David, le Roux, Sara, 2015. An experimental study on the effect of ambiguity in a coordination game. Theory Decis. 79 (4), 667–688.
Klibanoff, Peter, Marinacci, Massimo, Mukerji, Sujoy, 2005. A smooth model of decision making under ambiguity. Econometrica 73 (6), 1849–1892.
Kuzmics, Christoph, 2017. Abraham Wald's complete class theorem and Knightian uncertainty. Games Econ. Behav. 104, 666–673.
Kuzmics, Christoph, Rogers, Brian W., Zhang, Xiannong, 2020. Is Ellsberg behavior evidence of ambiguity aversion? Mimeo.
Li, Chen, Turmunkh, Uyanga, Wakker, Peter P., 2019. Trust as a decision under ambiguity. Exp. Econ. 22, 51–75.
Liu, Zhiwei, Song, Xinxi, Yannelis, Nicholas C., 2020. Randomization under ambiguity: efficiency and incentive compatibility. J. Math. Econ. 90, 1–11.
Lo, Kin Chung, 2009. Correlated Nash equilibrium. J. Econ. Theory 144, 722–743.
Machina, Mark J., Siniscalchi, Marciano, 2013. Ambiguity and ambiguity aversion. In: Machina, Mark J., Kip Viscusi, W. (Eds.), Handbook of the Economics of Risk and Uncertainty. Newnes, pp. 730–807.
Oechssler, Jörg, Rau, Hannes, Roomets, Alex, 2016. Hedging and ambiguity. Mimeo.
Pearce, David G., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica 52 (4), 1029–1050.
Saito, Kota, 2015. Preferences for flexibility and randomization under uncertainty. Am. Econ. Rev. 105 (3), 1246–1271.
Sion, Maurice, 1958. On general minimax theorems. Pac. J. Math. 8 (1), 171–176.
Spears, Dean, 2009. Preference for randomization?: Anscombe–Aumann inconsistency in the lab. Mimeo.
Tan, Tommy Chin-Chiu, Werlang, Sergio Ribeiro da Costa, 1988. The Bayesian foundations of solution concepts of games. J. Econ. Theory 45, 370–391.
Weinstein, Jonathan, 2016. The effect of changes in risk attitude on strategic behavior. Econometrica 84 (5), 1881–1902.