Australian National University

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

by

Evan M. Calford

ANU Working Papers in Economics and Econometrics
675

JULY 2020

JEL: D81, C70, C72

ISBN: 0 86831 675 X

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

Evan M. Calford †

July 17, 2020

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 preferences for randomization weaken. 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 preferences 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.

Keywords: Ambiguity Aversion, Mixed Strategies, Game Theory
JEL Codes: D81, C70, C72

*I would like to thank Pierpaolo Battigalli, Adam Dominiak, Yoram Halevy and Terri Kneeland for helpful comments and suggestions.

†John Mitchell Fellow, Research School of Economics, Australian National University, Canberra, Australia. Evan.Calford@anu.edu.au

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 agents, with ambiguity averse subjects choosing 'safe' strategies more frequently, and (2) that subjects display a persistent preference for randomization.[1] The previous literature on rationalizabiilty 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).[2] In contrast, this paper presents a model that is consistent with both (1) and (2).

We assume throughout that ambiguity preferences can be represented by Maxmin Expected Utility (Gilboa and Schmeidler, 1989). 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

[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.

[2]See below for references and extensive discussion. Note that Battigalli et al. (2016) partially straddle this divide by using a non-standard implementation of mixed strategies. For now, we note that Battigalli et al. (2016) do not allow for agents to hold a strict preference for randomizing over strategies using an external randomization device, but defer the discussion until Section 6.

the minimum utility against strategies in the belief set. Maxmin Expected Utility 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

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

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

Consider the modernized Battle of the Sexes game in Figure 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 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 (1) but not (2). 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

[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.

uncertainty. All strategies are justifiable, and the entire game is therefore rationalizable.

Next, consider the literature that is consistent with (2) but not (1). 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.

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. 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

[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. This interpretation of mixed strategies is discussed further in Section 6.

[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.

Colin's strategy choice. 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.[7] 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 $\delta \leq \frac{x}{2}$.

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 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 remarkably simple. 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. Lemma 3 from Pearce (1984) continues to hold in our environment, and rules out the possibility that a strategy is justifiable only when used as part of a mixed strategy. Therefore, a strategy is justifiable if and only if it is a pure best response to some belief. Hence, as δ increases, less strategies are justifiable.

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

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.

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 discusses the interpretation of mixed strategies for agents with ambiguity averse preferences and Section 7 concludes.

2 Ambiguity aversion and preferences over randomization in games

We define a game by $G = \langle I, (A_i)_{i\in I}, (v_i)_{i\in I}, \delta\rangle$ where $I = \{1, \ldots, n\}$ 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-Nuemann Morgenstern utility function over outcomes, and δ is a preference parameter defined below.[8] 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.

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.

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 $\mathcal{P}$ and note that $\mathcal{P}$ 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 $\Phi = \times_{i\in N}\Phi_i$.[9]

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

[9] Saito (2015) considers the primitives of behavior to be a set of feasible actions, rather than an explicit random-

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

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

and, with an abuse of notation,

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

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.[10]

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.

3 Rationalizability

We define justifiable and rationalizable strategies in the standard way. A strategy is justifiable if there exists some beliefs for which the strategy lies in the support of the best response, and a profile

ization, 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.

[10] 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.

of strategies is rationalizable if each component is justifiable when beliefs are restricted to only the opponents' rationalizable strategies.

Definition 1. 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 2. A strategy a_i is justifiable if there exists a $\Phi_i \subseteq \Delta(A_{-i})$ such that a_i is justifiable with respect to Φ_i.

Definition 3. The set of rationalizable strategies is the largest set $Z = \times_{i\in N} Z_i \subseteq \times_{i \in N} A_i$ such that for every $k \in N$ each $a_k \in Z_k$ is justifiable with respect to beliefs Φ_k where $\Phi_k \subseteq \Delta(Z_{-k})$.

To emphasize the dependence of the rationalizable set on the underlying parameter δ, we will sometimes write Z_δ to denote the set of rationalizable strategies. When $\delta = 1$ this definition of rationalizability collapses to the standard notion of correlated rationalizability, given that beliefs are not restricted to be stochastically independent. When $\delta = 0$ this definition collapses to Epstein (1997) rationalizability for Maxmin Expected Utility agents. Further, note that, following standard arguments, the set of rationalizable strategies can be found by iterated elimination of non-justifiable strategies.

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

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

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.

□

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.

Result 2. *$a_i \in Z_0$ if and only if a_i is Epstein (1997) rationalizable.*

Proof. Immediate. □

Next we show that a result from Pearce (1984), that any undominated strategy must be a pure best response to some beliefs, extends to the current environment. That Pearce's result extends to our environment is not without surprise: given a preference for mixing, it appears plausible that a strategy might only be justifiable when used as part of a mixed strategy to build a hedge against uncertainty. It turns out, however, that a strategy is only useful as a hedge against uncertainty if that strategy is also able to 'stand on its own' for at least one set of beliefs.

Lemma 1. A pure strategy a_i is undominated if and only if there exists Φ_i such that, for all $\sigma_i \in \Delta(A_i)$, $U(a_i, \Phi_i, \delta) \geq U(\sigma_i, \Phi_i, \delta)$.

Proof. The 'only if' statement is obvious. The proof of the 'if' statement follows. The proof follows the proof of Lemma 3 in Pearce (1984) closely.

We prove the contrapositive. Suppose that α is never a best response: i.e. $\forall\ \Phi_i$, $\exists\ \sigma_i \in \Delta(A_i)$, $U(\alpha, \Phi_i, \delta) < U(\sigma_i, \Phi_i, \delta)$. Then there exists a function $b(\Phi_i) : \mathcal{P} \to \Delta(A_i)$ such that, for all Φ_i, $U(b(\Phi_i), \Phi_i, \delta) > U(\alpha, \Phi_i, \delta)$.

Consider the following zero-sum game, played between players 1 and 2. The set of (pure) strategies for player 1 is $S_1 = A_i$ and the set of (pure) strategies for player 2 is $S_2 = \mathcal{P}$, and representative mixed strategies are denoted by σ_1 and σ_2 respectively. The payoffs for player 1 are given by $\overline{U}_1(\sigma_1, \sigma_2) = U(\sigma_1, \sigma_2, \delta) - U(\alpha, \sigma_2, \delta)$, and the payoffs for player 2 are given by $\overline{U}_2(\sigma_1, \sigma_2) = -\overline{U}_1(\sigma_1, \sigma)$. Existence of a Nash equilibrium in this zero-sum game is guaranteed by the existence theorem of Zhou et al. (2011), given that S_1 and S_2 are compact metric spaces and the utility functions are continuous. Let (σ_1^*, σ_2^*) denote a Nash equilibrium of this game. Then, for any $\sigma_2 \in \mathcal{P}$:

$$
\begin{aligned}
\overline{U}_1(\sigma_1^*, \sigma_2) &\geq \overline{U}_1(\sigma_1^*, \sigma_2^*) \\
&\geq \overline{U}_1(b(\sigma_2^*), \sigma_2^*) \\
&> \overline{U}_1(\alpha, \sigma_2^*) \\
&= 0
\end{aligned}
$$

where the final line is implied by the construction of the zero-sum game. Further, $\overline{U}_1(\sigma_1^*, \sigma_2) > 0$ for all $\sigma_2 \in \mathcal{P}$ implies that $U(\sigma_1^*, \Phi_i, \delta) > U(\alpha, \Phi_i, \delta)$ for all $\Phi_i \in \mathcal{P}$. Therefore, α is dominated.

□

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. First, Lemma 1 implies that a_i must be a pure strategy best response for some beliefs at δ. That is, there exists a Φ_i such that, for all $\sigma \in \Delta(A_i)$, $U(a_i, \Phi_i, \delta) \geq U(\sigma, \Phi_i, \delta)$. Next, note that $U(\sigma, \Phi_i, \delta)$ is weakly decreasing in the third argument (because utility for mixed strategies falls as preference for mixing decreases) and $U(a_i, \Phi_i, \delta)$ is independent of the third argument (because utility for a pure strategy is independent of any preference for mixing). Therefore, for all $\sigma \in \Delta(A_i)$, $U(a_i, \Phi_i, \delta') \geq U(\sigma, \Phi_i, \delta')$. □

The proof is straightforward. Given Lemma 1, a strategy is justifiable if and only if it is a pure best response to some beliefs. The utility from playing any pure strategy is independent of δ, while the utility of a mixed strategy is weakly decreasing in δ. Therefore, if a_i is justifiable at δ then it must also be justifiable at $\delta' \leq \delta$. The converse is false.

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

Proof. Follows by application of iterated elimination of non-justifiable strategies. □

The cardinality of the set of justifiable strategies, and therefore cardinality of the set of rationalizable strategies decreases in δ. There is a clear relationship between this result and the main result of Battigalli et al. (2016), who demonstrate that the cardinality of the set of justifiable strategies increases with ambiguity aversion in the smooth ambiguity aversion model of Klibanoff et al. (2005). Here, however, ambiguity preferences are held constant. Instead the *expression* of ambiguity preferences changes with δ: an agent with large δ need not play a "safe" strategy becuase they may hedge their ambiguity via mixing.

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 (Figure 2) from Calford (2020) where payoffs have been normalized to lie between 0 and 1.

	X (0.66)	Y (0.34)
A (0.63)	$1, \frac{4}{5}$	$0, 0$
B (0.09)	$0, \frac{4}{5}$	$1, 0$
C (0.28)	$\frac{4}{11}, 0$	$\frac{4}{11}, 1$

Figure 2: Example from Calford (2020). Payoffs are normalized to lie between 0 and 1 and the proportion of subjects choosing each strategy is given in parentheses.

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 $\delta \leq \frac{8}{11}$.[11] We can conclude that, 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. Consider the game in Figure 3, from Kelsey and le Roux (2015), where, again, payoffs have been normalized to lie between 0 and 1. Strategy R is dominated by a mix between L and M 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 $\delta \leq \frac{4}{5}$.

[11]Note that the assumption of rationality with respect to the preferences outlined in Equation 1 implicitly assumes risk neutrality, given that payoffs in the original experiment represented monetary payoffs to subjects. However, risk preferences cannot explain the substantial fraction of row players who fail to choose the unique correlated rationalizable strategy in this game. To see this, we can restrict the sample to only those subjects who exhibit both ambiguity aversion and low levels of risk aversion – i.e. the subjects who would be the most likely to use mixing as a hedge against strategic uncertainty. Among this subsample, the proportion of subjects choosing A is 58%.

	L (0.40)	M (0.30)	R (0.30)
T (0.50)	$0, 0$	$1, \frac{1}{3}$	$\frac{1}{6}, \frac{1}{5}$
B (0.50)	$\frac{1}{3}, 1$	$0, 0$	$\frac{11}{60}, \frac{1}{5}$

Figure 3: Example two from Kelsey and le Roux (2015). Payoffs are normalized to lie between 0 and 1 and the proportion of subjects choosing each strategy is given in parentheses.

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 Figure 2. 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 $\delta > \frac{8}{11}$. 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 Internal and external mixed strategies

Battigalli et al. (2016) (henceforth BCVMM) allow for two types of mixed strategies.[12] Primarily, they consider only mixed strategies that are defined in the action space of the game; we will refer to these as internal mixed strategies. To complete some of their proofs, they require a more standard form of mixed strategy; we will refer to these as external mixed strategies given they are mixed strategies that the agent can implement externally, by randomizing over the pre-defined set of

[12]The working paper version of Battigalli et al. (2016) contains a more substantial discussion of mixed strategies than does the final published version.

available strategies. BCVMM assume that agents have a preference for randomization over internal mixed strategies, but do *not* have a preference for randomization over external mixed strategies.

The main result of BCVMM is that the set of justifiable strategies grows as ambiguity aversion increases in the smooth ambiguity aversion model (Klibanoff et al. (2005)). However, if the action space includes all possible internal mixed strategies (e.g. the strategy space is a simplex representing the mixed extension of a finite game) then the main result of BCVMM collapses and the set of justifiable strategies is independent of ambiguity preference. Thus, the BCVMM result relies on a restriction of the preference for mixing – via the assumption that there is no preference for randomization over external mixed strategies, and by restricting the availability of internal mixed strategies.

Chen and Luo (2012) apply a similar approach to mixing. In the case where Chen and Luo (2012) apply their model to finite games (i.e. the strategy space is a simplex in a finite-dimensional Euclidian space) the most natural interpretation is that mixed strategies are internal mixed strategies. Chen and Luo (2012) assume that preferences satisfy a "concave-like" condition that implies a maximal preference for mixing over these internal mixed strategies. In this case, the game is then isomorphic to a standard finite game with external mixed strategies and a maximal preference for randomization. The assumption that agents have a maximal preference for mixing leads to their main result: that the set of rationalizable strategies for ambiguity averse and ambiguity neutral agents are identical.

The position taken in this paper is that it is never possible to prevent an agent from using an external mixed strategy. That is, in every game, whether a laboratory experiment or real-world interaction, agents can always randomize before playing the game by either consulting a randomization device or randomizing inside their own heads. Further, preferences for randomization over external mixed strategies are subjective and cannot be restricted by the game designer. Therefore, we must at least allow for the possibility that agents have a preference for randomization with respect to external mixed strategies.

It would be, however, possible to reinterpret the framework used here to accommodate internal mixed strategies. To do so requires that the payoff associated with any internal mixed strategy is *defined* to be consistent with Equation 1. This implies that the game theorist must specify the strength of preference for randomization as part of the game. In most settings, including lab experiments, it would be unusual for the game theorist to have knowledge of δ prior to the game being played. In such an environment, it seems more natural for the game theorist to define only the payoffs of outcomes and allow the agents to choose, or not choose, to implement external mixed strategies as desired.

7 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.

References

Marina Agranov and Pietro Ortoleva. Stochastic choice and preferences for randomization. *Journal of Political Economy*, 125(1):40–68, 2017. → pages 2

Marina Agranov, Paul J. Healy, and Kirby Nielsen. Non-random randomization. Mimeo, February 2020. → pages 2

Sophie Bade. Randomization devices and the elicitation of ambiguity-averse preferences. *Journal of Economic Theory*, 159:221–235, 2015. Part A. → pages 4

Aurelien Baillon, Yoram Halevy, and Chen Li. Experimental elicitation of ambiguity attitude using the random incentive system. Mimeo, 2019. → pages 4

P. Battigalli, S. Cerreia Vioglio, F. Maccheroni, and M. Marinacci. A note on comparative ambiguity aversion and justifiability. *Econometrica*, 84(5):1903–1916, 2016. → pages 2, 6, 10, 12

Evan M. Calford. Uncertainty aversion in game theory: Experimental evidence. *Journal of Economic Behavior and Organization*, 2020. Forthcoming. → pages 2, 3, 11, 14

Yi-Chun Chen and Xiao Luo. An indistinguishability result on rationalizability under general preferences. *Economic Theory*, 51:1–12, 2012. → pages 4, 5, 6, 13

Jurgen Eichberger and David Kelsey. Uncertainty aversion and preference for randomization. *Journal of Economic Theory*, 71:31–43, 1996. → pages 4

Larry G. Epstein. Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73: 1–29, 1997. → pages 3, 5, 8, 9

Larry G. Epstein, Massimo Marinacci, and Kyoungwon Seo. Coarse contingencies and ambiguity. *Theoretical Economics*, 2:355–394, 2007. → pages 4

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

Shaowei Ke and Qi Zhang. Randomization and ambiguity aversion. *Econometrica*, 88(3): 1159–1195, 2020. → pages 4

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 11, 12

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005. → pages 6, 10, 13

Chen Li, Uyanga Turmunkh, and Peter P. Wakker. Trust as a decision under ambiguity. *Experimental Economics*, 22:51–75, 2019. → pages 2

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

Mark J. Machina and Marciano Siniscalchi. Ambiguity and ambiguity aversion. In Mark J. Machina and W. Kip Viscusi, editors, *Handbook of the Economics of Risk and Uncertainty*, pages 730–807. Newnes, 2013. → pages 2

David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984. → pages 5, 9

Kota Saito. Preferences for flexibility and randomization under uncertainty. *American Economic Review*, 105(3):1246–1271, 2015. → pages 1, 4, 5, 6, 14

Y.H. Zhou, J. Yu, and L. Wang. A new proof of existence of equilibria in infinite normal form games. *Applied Mathematics Letters*, 24:253–256, 2011. → pages 9