

Random Belief Equilibrium in Normal Form Games¹

James W. Friedman

Department of Economics, CB #3305

University of North Carolina

Chapel Hill, NC 27599-3305

email: Jim_Friedman@unc.edu

Claudio Mezzetti

Department of Economics, CB #3305

University of North Carolina

Chapel Hill, NC 27599-3305

email: mezzetti@email.unc.edu

May 1, 2002

¹We would like to thank two anonymous referees for useful suggestions. This research was completed while the second author was visiting the University of Venice Ca' Foscari, Department of Applied Mathematics, whose financial support is gratefully acknowledged.

Abstract

We define *Random Belief Equilibrium* (RBE) in finite, normal form games. We assume that a player's beliefs about the strategy choices of others are randomly drawn from a belief distribution. This distribution is dispersed around a central strategy profile, the *focus*. At an RBE: (1) Each player chooses a best response relative to her drawn beliefs. (2) The expected choice of each player coincides with the focus of the other players' belief distributions. RBE provides a statistical framework for estimation which we apply to data from three experimental games.

We provide a characterization of the limit-RBE as players' beliefs converge to certainty. Let $D_i(s_i)$ be the set of mixed strategies of player i 's opponents to which s_i is a best reply. When atoms in the belief distributions vanish in the limit, all pure strategies s_i used with positive probability by player i in a limit-RBE (called a *Robust Equilibrium*) must have the following property. The intersection of $D_i(s_i)$ with any ball around the equilibrium mixed strategies of player i 's opponents must have positive Lebesgue measure. We show that this implies that not all trembling hand perfect Nash equilibria are robust and not all robust equilibria are perfect.

1 Introduction

Equilibrium concepts typically include, either implicitly or explicitly, beliefs that players have about how rival players will play. For Nash equilibrium these beliefs must be exactly correct; that is, they must put all probability mass on the precise equilibrium strategies followed by others. On the other hand rationalizability, due to Bernheim [1] and Pearce [11], does not require coincidence of a player's beliefs with her rivals' true strategies; the only restriction on beliefs is that they attach positive probability only to those pure strategies that rivals could play if the rationality of all players were common knowledge. Thus, a player's rationalizable behavior is typically associated with incorrect beliefs, but, at the same time, these beliefs are not contradicted by the strategies chosen by the other players.

In this paper we provide a generalization of Nash equilibrium for finite normal form games, called *random belief equilibrium* (RBE), based upon players having uncertainty concerning the strategies chosen by rival players at equilibrium. RBE retains the assumption that each player chooses a best reply to her beliefs. Unlike rationalizability, RBE requires that players' beliefs be statistically consistent with the equilibrium strategies. We postulate that players' beliefs about the mixed strategy choices of each of their opponents are randomly drawn from some belief distribution that is defined over the mixed strategy set of the opponent and view this distribution as being dispersed around a central strategy profile that we call the *focus*. The focus could be, for example, the mean or the mode of the distribution. At an RBE: (i) Each player chooses a best response relative to the beliefs she draws (thus, the belief distribution induces a distribution over the set of possible choices). (ii) Beliefs are statistically consistent in the sense that the *expected choice* of each player i coincides with the focus of the belief distributions of the other players about player i 's choices. Because the beliefs of a player are drawn from a non-degenerate probability measure over the other players' strategies, a player's drawn beliefs are typically incorrect and an RBE is not a Nash equilibrium.

We parameterize the dispersion of beliefs with a positive scalar ε . Less dispersed beliefs are more concentrated around the focus and, as $\varepsilon \rightarrow 0$, the belief distributions converge to the degenerate distribution having unit probability mass on the focus. As uncertainty vanishes, RBE converges to a Nash equilibrium. We permit the belief distributions to have atoms under either of two conditions: (i) The probability mass on atoms goes to one as $\varepsilon \rightarrow 0$ or (ii) the probability mass on atoms goes to zero as $\varepsilon \rightarrow 0$.¹ As $\varepsilon \rightarrow 0$, a sequence of RBE is generated; the cluster points of this sequence are called limit-RBE. Under case (i) the set of limit-RBE coincides with the set of Nash equilibria. Under case (ii) we call the limit-RBE *robust* equilibria. Theorem 4 provides a characterization of robust equilibrium showing that a Nash equilibrium is robust if and only if every pure strategy that is played with positive probability by a player is a best reply to all the mixed strategies of the other players that belong to a set having positive Lebesgue measure within any small ball around the equilibrium profile of the other players.

Robust equilibria have an appealing robustness: They only utilize pure strategies that are best replies to all mixed strategies of opponents in a set with positive measure in a neighborhood of the equilibrium. We show that not all perfect equilibria are robust. For two player games, the set of robust equilibria is a subset of the set of trembling hand perfect equilibria, but if there are at least three players not all robust equilibria are perfect.

An important property of RBE is that it yields a statistical framework for estimation using either field or experimental data. RBE can be fruitfully compared with the quantal response equilibrium (QRE) of McKelvey and Palfrey [6]. QRE differs from Nash equilibrium by supposing that players have uncertainty about payoffs or, equivalently, players' choices are subject to decision errors. The error dispersion parameter λ varies from purely random choice, $\lambda = 0$, to no error at all as $\lambda \rightarrow \infty$.² Thus both

¹Note that in the limit, as $\varepsilon \rightarrow 0$, there can be at most one non-vanishing atom. Moreover, if there is a non-vanishing atom, it must coincide with the focus of the distribution.

²An alternative interpretation of the mathematical structure used by McKelvey and Palfrey is

RBE and QRE differ from Nash equilibrium and from rationalizability by including a parameter in the equilibrium concept that can vary, yielding Nash equilibrium at one extreme, and offering the ability to be estimated from data. The key difference is that RBE is based upon a model in which players know their own payoff functions precisely, but have stochastic beliefs about other players' strategy choices, while QRE is based upon a model in which players have exact, correct beliefs about other players' strategy choices, but have stochastic knowledge of their own payoff functions.

QRE retains an empirical advantage over RBE in games having weakly dominated strategies, because RBE always assigns zero probability to such strategies while, in practice, experimental subjects sometimes choose them. Consequently the likelihood function associated with RBE is always zero in these games. QRE permits dominated strategies to be played with positive probability. One can estimate a model in which RBE and QRE both appear (both payoff uncertainty and uncertainty about opponents' strategies are present). Then one can test whether removing the RBE component leads to a significant difference. This, however, precludes any direct comparison between RBE and QRE, although it permits testing for the significance of (i) RBE in the presence of QRE and of (ii) QRE alone. It remains possible to estimate RBE directly in games lacking weakly dominated strategies.

Using a Dirichlet distribution to represent beliefs, we fit RBE to three sets of experimental data by maximizing a likelihood function with respect to the degree of belief uncertainty ε . The data are borrowed from the experiments in Ochs [8], McKelvey, Palfrey and Weber [7], and O'Neill [9] in which subjects played a game with a unique, completely mixed, Nash equilibrium.³ We find that random beliefs help to explain the subjects behavior; the belief dispersion ε is significantly different from zero in all three sets of games. In two cases the Dirichlet RBE fits the data better than the logistic QRE and in the O'Neill experiment the QRE fits better. We

that the players are boundedly rational (see Chen et al. [2]) in the manner of Luce [5]. Then the estimated parameter is interpreted as measuring the degree of rationality of players.

³We thank Jack Ochs for making his experimental data available to us.

had no à priori notion of whether to expect belief uncertainty to be more empirically fruitful than payoff uncertainty; these are two differing sources of randomness and either, neither, or both could be present in a given situation. In related work Haruvy, Stahl and Wilson [3] introduced and tested an econometric model in which players have heterogeneous priors about their opponents choices, as in RBE, and random payoffs, as in QRE. They did not develop an equilibrium concept, however, but instead considered boundedly rational players in the spirit of Stahl and Wilson [13], [14], finding that both heterogeneous priors and payoff errors are present in the data sets they study.

The paper proceeds as follows. In Section 2 the model is described. We define (i) *random belief games*, consisting of a game and a belief structure, (ii) random belief equilibrium (RBE), consisting of a strategy profile that is both consistent with the belief structure and optimal relative to it, (iii) *limit-RBE* which is the limit of a sequence of RBE strategy profiles as the dispersion of beliefs goes to zero, and (iv) *robust equilibrium* which is a limit-RBE when the probability mass concentrated on atoms converges to zero. In Section 3 we prove that RBE and limit-RBE exist and provide characterizations of limit-RBE with non-vanishing atoms in the belief distributions, and of robust equilibrium. In Appendix C we show that robust equilibrium coarsens essential equilibrium and is unrelated to weakly proper equilibrium. In Section 4 we use maximum likelihood estimation to fit RBE to the experimental data and compare the performance of RBE with QRE. Section 5 concludes.

2 The Model

A *random belief game*, defined in Section 2.1, is a pair consisting of a finite game and a belief structure. In a random belief game, players' beliefs about the strategies of their opponents are random variables. More precisely, before choosing their strategies in an n -player game Γ , players draw belief realizations about their rivals' strategy

choices from a belief measure, then each player chooses a best reply to her belief realization. The belief measure $\mu_{j\varepsilon}^i(\sigma_j^f)$ of player i concerning the choice of rival player j is a probability measure over the set of j 's mixed strategy profiles Δ_j that is parameterized by $\varepsilon \in \mathbb{R}_+$ and by $\sigma_j^f \in \Delta_j$. The parameter σ_j^f , called the *focus* about j , is a measure of central tendency (e.g., it could be the expected value or the mode); ε parameterizes the dispersion of $\mu_{j\varepsilon}^i(\sigma_j^f)$ around σ_j^f .⁴ We permit $\mu_{j\varepsilon}^i(\sigma_j^f)$ to have atoms and denote the total probability of the atoms by $\alpha_{j\varepsilon}^i(\sigma_j^f)$.

In Section 2.2 *random belief equilibrium* is defined as a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ that is both an expected best reply to the belief structure and is statistically consistent with it. Statistical consistency means that, for each player i , the strategy σ_i equals the focus of all the other player's beliefs about i ; that is, $\sigma_i = \sigma_i^f$. Thus, at a random belief equilibrium the expected best reply of each player equals the focus about her strategy of all other players' beliefs (e.g., if the focus is the expected value, then the other players' beliefs about the strategy chosen by player i are correct on average). This consistency condition is a weakening of the usual Nash equilibrium consistency condition that the strategy chosen by player i equals the strategy every other player expects i to choose.

As $\varepsilon \rightarrow 0$ the belief measures collapse to certainty so that, as described in Section 2.3, the random belief game converges to a conventional game and the random belief equilibrium converges to a Nash equilibrium. Existence and properties of this limiting equilibrium are deferred to Section 3.

2.1 Random Belief Games

A *random belief game* is a pair $\langle \Gamma, \mu_\varepsilon \rangle$ where Γ is a finite game and $\mu_\varepsilon = (\mu_\varepsilon^1, \dots, \mu_\varepsilon^n)$ is a *belief map profile* where μ_ε^i is a *belief map* of player i which specifies a family

⁴We assume that the belief measures of all players depend on the same parameter ε . This is just for notational convenience. Nothing of substance would change by having a separate parameter ε_i for each player i .

of probability measures over the set of mixed strategy profiles of the other players. The class of games Γ is specified in Definition 1 and the class of belief map profiles is specified in Definition 3 with the aid of the standard definition of a probability space in Definition 2.

Definition 1 *The finite game $\widehat{\Gamma} = \langle N, S, U \rangle$ is given by a finite set of players $N = \{1, \dots, n\}$, a set of pure strategy profiles $S = \times_{i \in N} S_i$, where $S_i = \{1, \dots, m_i\}$ is the finite pure strategy set of player i , and the utility profile $U = (U_1, \dots, U_n)$, where $U_i : S \rightarrow \mathbb{R}$ is the payoff function of player i . Denote by $\Gamma = \langle N, \Delta, u \rangle$ the mixed extension of $\widehat{\Gamma}$, where Δ_i is the set of mixed strategies of player i , $\Delta = \times_{i \in N} \Delta_i$, $u = (u_1, \dots, u_n)$, and $u_i : \Delta \rightarrow \mathbb{R}$ is derived from U_i via von Neumann-Morgenstern utility.*

Henceforth *finite game* means a mixed extension $\Gamma = \langle N, \Delta, u \rangle$ as described above in Definition 1. The following conventional notations are used throughout the paper: $\Delta_{-i} = \times_{j \neq i} \Delta_j$, σ_{-i} is a typical element of Δ_{-i} , $S_{-i} = \times_{j \neq i} S_j$, and s_{-i} is a typical element of S_{-i} . The probability assigned by σ_i to $s_i \in S_i$ is denoted $\sigma_i(s_i)$.

Definition 2 *Let X be a complete, separable metric space, \mathcal{B} the Borel sets of X , and $\mathcal{M}(X)$ the set of probability measures on the measurable space $\langle X, \mathcal{B} \rangle$. Given $\eta \in \mathcal{M}(X)$, the triple $\langle X, \mathcal{B}, \eta \rangle$ is a probability space.*

Each player i has a *belief realization* $\sigma_j^R \in \Delta_j$ about the strategies the players $j \neq i$ will choose; each σ_j^R is chosen randomly according to the *belief measure* $\mu_{j\varepsilon}^i(\sigma_j^f) \in \mathcal{M}(\Delta_j)$, where $\sigma_j^f \in \Delta_j$ is the *focus* about player j , and $\varepsilon \in \mathbb{R}_+$ is a parameter. For any $\Delta_j^0 \subset \Delta_j$, the probability that $\sigma_j^R \in \Delta_j^0$ under $\mu_{j\varepsilon}^i(\sigma_j^f)$ is denoted $\mu_{j\varepsilon}^i(\sigma_j^f)(\Delta_j^0)$. The belief measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ is the product measure of the $\mu_{j\varepsilon}^i(\sigma_j^f)$ for $j \neq i$. That is, let $\mathcal{B}^i = \otimes_{j \neq i} \mathcal{B}_j$, and $\mu_\varepsilon^i(\sigma_{-i}^f)(A^i) = \times_{j \in N \setminus \{i\}} \mu_{j\varepsilon}^i(\sigma_j^f)(A_j)$ for any $A^i = \times_{j \neq i} A_j$ where $A_j \subset \Delta_j$ and $\langle \Delta_{-i}, \mathcal{B}^i, \mu_\varepsilon^i(\sigma_{-i}^f) \rangle$ denotes the product space $\times_{j \neq i} \langle \Delta_j, \mathcal{B}_j, \mu_{j\varepsilon}^i(\sigma_j^f) \rangle$. Each player thus has a *belief map* $\mu_\varepsilon^i : \Delta_{-i} \rightarrow \mathcal{M}(\Delta_{-i})$, where $\mathcal{M}(\Delta_{-i})$ is the subset of

probability measures on Δ_{-i} for which the respective σ_j^R , for $j \neq i$, are independently distributed. The belief map μ_ε^i associates a belief measure to each focus σ_{-i}^f ; $\mu_\varepsilon = (\mu_\varepsilon^1, \dots, \mu_\varepsilon^n)$ denotes a *belief map profile*.

Thus if a focus σ^f is specified, each player i has well defined probabilistic beliefs about the choice that will be made by any other player j . These are given by the probability measure $\mu_{j\varepsilon}^i(\sigma_j^f)$ on Δ_j . The probabilistic beliefs of player i about the behavior of all remaining players $N \setminus \{i\}$ are assumed to be a product measure; that is, we assume that all players believe that their rivals are acting independently of one another. There is no coordination among them and no correlating device is available. We make this assumption because we want to stay close to the spirit of Nash equilibrium. Looking at random beliefs when correlating devices are available is left to future research.

Definition 3 specifies that the belief measures $\mu_{j\varepsilon}^i(\sigma_j^f)$ center in a weak sense around the focus σ_j^f by requiring that, as $\varepsilon \in \mathbb{R}_+$ goes to zero, $\mu_{j\varepsilon}^i(\sigma_j^f)$ weakly converges to the measure that puts all mass on the focus σ_j^f .⁵ For all $\varepsilon > 0$, we permit the belief measure $\mu_{j\varepsilon}^i(\sigma_j^f)$ to have a finite number of atoms. Clearly, weak convergence to the focus implies that, as $\varepsilon \rightarrow 0$, all but possibly one atom must vanish. The definition also requires continuity of $\mu_{j\varepsilon}^i(\sigma_j^f)$ with respect to σ_j^f , and that any subset of Δ_j has positive measure under $\mu_{j\varepsilon}^i(\sigma_j^f)$ if it has positive Lebesgue measure. The latter condition ensures that the expected strategy of player j under the beliefs of player i must be completely mixed.

Definition 3 Let \mathcal{B}_j be the Borel sets of Δ_j . For each $\sigma_j^f \in \Delta_j$, each $i \in N$, and

⁵On weak convergence, see, for example, Parthasarathy [10]. Let $C(X)$ be the set of all bounded, real valued, continuous functions on X . Then $\{\mu_\varepsilon\}$ weakly converges to μ if and only if $\int f d\mu_\varepsilon \rightarrow \int f d\mu$ for all $f \in C(X)$. The topology generated by weak convergence is called the *weak topology*. The *Prohorov metric* p on $\mathcal{M}(X)$ also generates the weak topology. Let $B^\xi = \{x \in X \mid \delta(x, B) < \xi\}$ and $\delta(x, B)$ be the distance between x and B ; for $\mu, \gamma \in \mathcal{M}(X)$, the Prohorov metric is

$$p(\mu, \gamma) = \inf\{\xi > 0 \mid \mu(B) \leq \gamma(B^\xi) + \xi \text{ and } \gamma(B) \leq \mu(B^\xi) + \xi \text{ for all } B \in \mathcal{B}\}.$$

each $j \neq i$, denote by $\langle \Delta_j, \mathcal{B}_j, \mu_{j\varepsilon}^i(\sigma_j^f) \rangle$ a collection of probability spaces parameterized by $\varepsilon \in \mathbb{R}_+$ having the properties that (i) for any $\varepsilon \in \mathbb{R}_{++}$ and any $\sigma_j^f \in \Sigma_j$, the non-atomic part of $\mu_{j\varepsilon}^i(\sigma_j^f)$ is absolutely continuous with respect to Lebesgue measure on Δ_j and $\mu_{j\varepsilon}^i(\sigma_j^f)(A) > 0$ whenever $A \in \mathcal{B}_j$ has positive Lebesgue measure; (ii) for any $\varepsilon \in \mathbb{R}_{++}$ and any $A \in \mathcal{B}_j$, $\mu_{j\varepsilon}^i(\sigma_j^f)(A)$ is continuous with respect to $\sigma_j^f \in \Delta_j$; (iii) as $\varepsilon \rightarrow 0$, $\mu_{j\varepsilon}^i(\sigma_j^f)$ weakly converges to $\mathbf{1}_{\sigma_j^f}$, where $\mathbf{1}_{\sigma_j^f}$ is the singular probability measure that puts all mass on σ_j^f ; (iv) the measure $\mu_{j\varepsilon}^i(\sigma_j^f)$ may have at most a finite number of atoms.

For $\sigma^f \in \Delta$, denote by $\mathbf{1}_{\sigma_{-i}^f}$ the limit of $\mu_{\varepsilon}^i(\sigma_{-i}^f)$ as $\varepsilon \rightarrow 0$ and let $\mathbf{1}_{\sigma^f} = (\mathbf{1}_{\sigma_{-1}^f}, \dots, \mathbf{1}_{\sigma_{-n}^f})$.

The measure $\mu_{\varepsilon}^i(\sigma_{-i}^f)$ is equal to the sum of an absolutely continuous measure $\mu_{\varepsilon}^{iA}(\sigma_{-i}^f)$ and a discrete measure $\mu_{\varepsilon}^{iD}(\sigma_{-i}^f)$. Let $\alpha_{j\varepsilon}^i(\sigma_j^f)$ be the total probability mass of the discrete measure, (i.e., the total mass on atoms), $\alpha_j^i(\sigma_j^f) = \lim_{\varepsilon \rightarrow 0} \alpha_{j\varepsilon}^i(\sigma_j^f)$, and $\alpha_0 = \sup_{i,j,\sigma_j^f} \alpha_j^i(\sigma_j^f)$. We are interested in two different cases. In case (i) there is no restriction on α_0 . In particular, we allow $\alpha_0 = 1$ so that the non-atomic part of beliefs vanishes as $\varepsilon \rightarrow 0$. In case (ii) we require $\alpha_0 = 0$, so that atoms vanish in the limit measures.

Definition 4 A random belief game is given by $\langle \Gamma, \mu_{\varepsilon} \rangle$ where $\Gamma = \langle N, \Delta, u \rangle$ satisfies Definition 1 and μ_{ε} is a belief map profile satisfying (i), (ii), and (iv) of Definition 3.

2.2 Random Belief Equilibrium

The definition of random belief equilibrium relies upon the conventional best reply mapping, $b = (b_1, \dots, b_n)$, which, for player i , is

$$b_i(\sigma_{-i}) = \{\sigma_i \in \Delta_i \mid u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) \forall \sigma'_i \in \Delta_i\} \quad \forall \sigma_{-i} \in \Delta_{-i}.$$

In the game $\langle \Gamma, \mu_{\varepsilon} \rangle$, given the focus σ^f , player i chooses a best reply to the belief realization σ_{-i}^R drawn from the belief measure $\mu_{\varepsilon}^i(\sigma_{-i}^f)$. If the conventional best reply

mapping, b_i , is single-valued everywhere except on a set of measure zero, then integrating the best reply with respect to the belief measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ yields a single-valued expected best reply function $\psi_{\mu_\varepsilon^i}(\sigma_{-i}^f)$ for the player. Namely,

$$\psi_{\mu_\varepsilon^i}(\sigma_{-i}^f) = \int_{\Delta_{-i}} b_i(\sigma_{-i}^R) d\mu_\varepsilon^i(\sigma_{-i}^f).$$

Denoting the expected best reply function of the game by $\psi_{\mu_\varepsilon} = (\psi_{\mu_\varepsilon^1}, \dots, \psi_{\mu_\varepsilon^n})$, a random belief equilibrium of the game $\langle \Gamma, \mu_\varepsilon \rangle$ is a fixed point σ' of ψ_{μ_ε} . That is,

$$\sigma' = (\sigma'_1, \dots, \sigma'_n) = \psi_{\mu_\varepsilon}(\sigma') = [\psi_{\mu_\varepsilon^1}(\sigma'_{-1}), \dots, \psi_{\mu_\varepsilon^n}(\sigma'_{-n})].$$

Thus for each player $i \in N$, σ'_i is an expected best reply to the belief measure with focus σ'_{-i} .

If b_i is multi-valued on a set of positive measure, or at a point corresponding to an atom in the belief measure, then the expected best reply mapping is a correspondence denoted Ψ_{μ_ε} . This mapping is derived in Appendix A. In Appendix B it is shown that the expected best reply mapping is single-valued whenever $\mu_\varepsilon^i(\sigma_{-i}^f)$ has no atoms and there is no player in the game having a pair of *indifferent strategies*. A pair of strategies, s_i, s'_i , of player i are said to be indifferent if $u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.

We are now ready to formally define random belief equilibrium.

Definition 5 *Let $\langle \Gamma, \mu_\varepsilon \rangle$ be a random belief game satisfying Definition 1 and parts (i), (ii) and (iv) of Definition 3, $\sigma(\varepsilon)$ be a strategy profile, and Ψ_{μ_ε} be the expected best reply correspondence given by equations (4) to (6) in Appendix A. If $\sigma(\varepsilon) \in \Psi_{\mu_\varepsilon}(\sigma(\varepsilon))$, then $\sigma(\varepsilon)$ is a random belief equilibrium or RBE.*

At an RBE, $\sigma(\varepsilon)$, for each $i \in N$ the strategy $\sigma_i(\varepsilon)$ is an expected best reply with the expectation taken over Δ_{-i} using the belief measure $\mu_\varepsilon^i[\sigma_{-i}(\varepsilon)]$.

2.3 Robust Equilibrium

A limit-RBE is a limit point of strategy profiles associated with random belief equilibria as $\varepsilon \rightarrow 0$. If $\alpha_0 = 0$, so that the atoms in the belief measures vanish in the limit, we call the limit-RBE a *robust equilibrium*. Let $M = \{\mu_\varepsilon\}_{\varepsilon \rightarrow 0}$ be a sequence of belief map profiles associated with a sequence $\{\varepsilon_h\}_{h=1}^\infty$ having the property that $\lim_{h \rightarrow \infty} \varepsilon_h = 0$. Denote by $\langle \Gamma, M \rangle = \{\langle \Gamma, \mu_\varepsilon \rangle\}_{\varepsilon \rightarrow 0}$ the associated sequence of random belief games.

Definition 6 *Let $\langle \Gamma, M \rangle$ be a sequence of random belief games satisfying Definitions 1 and 3 and, for each ε , let $\sigma(\varepsilon)$ be a random belief equilibrium of $\langle \Gamma, \mu_\varepsilon \rangle$. A cluster point σ^0 of the sequence of random belief equilibria $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ is an M -limit-RBE. A strategy profile σ^0 is a limit-RBE if σ^0 is an M -limit-RBE for some M . The profile σ^0 is called a robust equilibrium if it is a limit-RBE for some sequence of belief map profiles such that $\alpha_0 = 0$.*

Note that σ^0 is a limit-RBE as long as one can find a sequence of belief map profiles M with respect to which σ^0 is an M -limit-RBE. This is reminiscent of the specification of Selten's [12] perfect equilibrium; recall that σ^0 is a *perfect Nash equilibrium* of Γ as long as one can find a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ of completely mixed strategy profiles that converges to σ^0 and such that σ_i^0 is a best reply to $\sigma_{-i}(\varepsilon)$ for all $i \in N$ and all ε . In the profile $\sigma(\varepsilon)$ players use completely mixed strategies; this captures the idea that a player may tremble when choosing her action, so that each choice must have positive probability. In our setup no player ever trembles in her choice of action. However, the beliefs $\mu_{j\varepsilon}^i(\sigma_j)$ associated with an RBE are random; that is, trembles occur in each player's beliefs about her opponents' strategies.

3 Existence and Characterizations

Section 3.1 uses standard arguments to establish the existence of both RBE and limit-RBE. In Section 3.2 we first prove that if there are no restrictions on the atoms in the

belief measures, then the set of limit-RBE coincides with the set of Nash equilibria. Then we characterize robust equilibria; that is, limit-RBE when atoms in the belief distributions vanish in the limit.

Before proceeding we need to introduce some definitions. The inverse of the conventional best reply correspondence for player i in the game Γ is given by

$$D_i(s_i) = \left\{ \sigma_{-i} \in \Delta_{-i} \mid u_i(s_i, \sigma_{-i}) = \max_{s'_i \in S_i} u_i(s'_i, \sigma_{-i}) \right\} \quad (1)$$

$D_i(s_i) \subset \Delta_{-i}$ is the set of mixed strategy profiles of player i 's opponents to which s_i is a best reply in the game Γ . Let $d_i(\emptyset) = \emptyset$ and for all non-empty $T_i \subset S_i$, let

$$d_i(T_i) = \{ \sigma_{-i} \in \Delta_{-i} \mid \sigma_{-i} \in \cap_{s_i \in T_i} D_i(s_i), \sigma_{-i} \notin D_i(s_i) \text{ for } s_i \notin T_i \} \quad (2)$$

The set $d_i(T_i) \subset \Delta_{-i}$ consists of mixed strategy profiles of player i 's opponents to which precisely the pure strategies in T_i , and only those strategies, are best replies for i in the game Γ . The collection of sets $\{d_i(T_i)\}_{T_i \subset S_i}$ partitions Δ_{-i} ; the sets $\{d_i(T_i)\}$ and $\{D_i(s_i)\}$ are related by

$$D_i(s_i) = \cup_{\substack{T_i \subset S_i \\ s_i \in T_i}} d_i(T_i)$$

3.1 Existence of RBE and Limit-RBE

Theorem 1 proves the existence of an RBE; Theorem 2 establishes the existence of a limit-RBE. Both proofs are conventional and straightforward. The existence proof for the RBE relies on the best reply mapping satisfying the Kakutani fixed point theorem. The proof for limit-RBE uses a standard convergence argument along with continuity properties of the payoff functions.

Theorem 1 *Let $\langle \Gamma, \mu_\varepsilon \rangle$ be a random belief game satisfying Definition 1 and parts (i), (ii), and (iv) of Definition 3. Then there exists a random belief equilibrium $\sigma(\varepsilon)$.*

Proof. A strategy profile $\sigma(\varepsilon)$ is an equilibrium if and only if it is a fixed point of $\Psi_{\mu_\varepsilon} : \Delta \rightarrow \Delta$. It is easy to verify that Ψ_{μ_ε} satisfies the conditions of the Kakutani fixed point theorem and hence that there exists $\sigma(\varepsilon) \in \Psi_{\mu_\varepsilon}(\sigma(\varepsilon))$. First, Δ is a compact and convex set. Second, by Equations (4)-(6) in Appendix A, $\Psi_{\mu_\varepsilon}(\sigma)$ is a non-empty and convex set for all $\sigma \in \Delta$. Finally, from Definition 3, for any $d_i(T_i) \subset \Delta_{-i}$, $\mu_\varepsilon^i(\sigma_{-i})(d_i(T_i))$ is continuous with respect to $\sigma_{-i} \in \Delta_{-i}$ and thus Ψ_{μ_ε} is upper hemicontinuous. ■

To prove existence of a limit-RBE we show that a sequence of random belief equilibria $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ has a cluster point σ^0 . This cluster point is a limit-RBE.

Theorem 2 *Every game Γ satisfying Definitions 1 and 3 has a limit-RBE.*

Proof. Given any game Γ , let $\sigma(\varepsilon)$ be an RBE of $\langle \Gamma, \mu_\varepsilon \rangle$ for $\varepsilon \rightarrow 0$. The sequence $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ is contained in the compact set Δ ; hence it has a convergent subsequence $\{\sigma(\varepsilon_\ell)\}_{\ell=1}^\infty$ with $\varepsilon_\ell \rightarrow 0$. Let $\sigma^0 = \lim_{\ell \rightarrow \infty} \sigma(\varepsilon_\ell)$ be the limit of such subsequence. Then, setting $M = \{\mu_{\varepsilon_\ell}\}_{\ell=1}^\infty$, σ^0 is a limit-RBE of Γ . ■

3.2 Characterizations of Limit-RBE and Robust Equilibrium

By letting the mass on atoms corresponding to Nash equilibria go to one in the limit, Theorem 3 shows that it is straightforward to obtain any Nash equilibrium as a limit-RBE when no restrictions are placed on atoms.

Define $\delta(\sigma, \sigma') = \max_{\substack{i \in N \\ s_i \in S_i}} |\sigma_i(s_i) - \sigma'_i(s_i)|$ as the distance between two strategy profiles σ and $\sigma' \in \Delta$ and $\delta(\sigma_{-i}, \sigma'_{-i}) = \max_{\substack{j \in N \setminus \{i\} \\ s_j \in S_j}} |\sigma_j(s_j) - \sigma'_j(s_j)|$ as the distance between σ_{-i} and $\sigma'_{-i} \in \Delta_{-i}$.⁶ Let $B_r(\sigma_{-i}^0)$ denote an open ball of radius r around σ_{-i} : $B_r(\sigma_{-i}) = \{\sigma'_{-i} \in \Delta_{-i} : \delta(\sigma_{-i}, \sigma'_{-i}) < r\}$.

⁶This metric is chosen for convenience. Since there are only finitely many pure strategies, any other metric could be used without affecting any of our conclusions.

Theorem 3 *If there are no restrictions on α_0 , then the set of limit-RBE of a game Γ coincides with the set of Nash equilibria.*

Proof. (Limit-RBE implies Nash.) Suppose, to the contrary, that σ^0 is the limit of a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ of RBE, but it is not a Nash equilibrium. Let $\{\mu_\varepsilon\}_{\varepsilon \rightarrow 0}$ be the belief map sequence associated with $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$. Then there exists i and $s_i \in S_i$ such that $\sigma^0(s_i) > 0$ and s_i is not a best reply to σ_{-i}^0 . Since the sets $D_i(s_i)$ are closed, there must be an open neighborhood of σ_{-i}^0 within which s_i is never a best reply; that is, we can find $r^* > 0$ such that for $r \in (0, r^*)$ the ball $B_r(\sigma_{-i}^0) \subset \Delta_{-i}$ of radius r around σ_{-i}^0 satisfies $B_r(\sigma_{-i}^0) \cap D_i(s_i) = \emptyset$. Moreover, for all $\theta > 0$, there exists ε^* such that, for $\varepsilon < \varepsilon^*$, $\mu_\varepsilon^i(\sigma_{-i}(\varepsilon))(B_r(\sigma_{-i}^0)) > 1 - \theta$; that is, the measure $\mu_\varepsilon^i(\sigma_{-i}(\varepsilon))$ with focus $\sigma_{-i}(\varepsilon)$ puts probability of at least $1 - \theta$ on the ball $B_r(\sigma_{-i}^0)$. This implies that $\sigma_i(\varepsilon)(s_i) < \theta$ and that $\sigma_i(\varepsilon)(s_i) \rightarrow 0$ as $\varepsilon \rightarrow 0$. This contradicts the assumption that $\sigma_i(\varepsilon) \rightarrow \sigma_i^0$ and $\sigma_i^0(s_i) > 0$.

(Nash implies limit-RBE) Let σ^* be a Nash equilibrium. Define

$$\sigma_i^\omega = (\sigma_i^\omega(1), \dots, \sigma_i^\omega(s_i), \dots, \sigma_i^\omega(m_i)) \in \mathbb{R}^{m_i}$$

as the vector with $\sigma_i^\omega(s_i) = \omega_i(D(s_i))$, where ω_i is Lebesgue measure over Δ_{-i} . Let $\sigma^\omega = (\sigma_1^\omega, \dots, \sigma_n^\omega)$. For each player i , define the following probability measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ with focus σ_{-i}^f . The non-atomic part of $\mu_\varepsilon^i(\sigma_{-i}^f)$ is Lebesgue measure with a total mass equal to ε . The remaining probability mass of $(1 - \varepsilon)$ goes on an atom at $\left[\left(\sigma_{-i}^f - \varepsilon \sigma_{-i}^\omega \right) / (1 - \varepsilon) \right]^+$ where $\left[\left(\sigma_{-i}^f - \varepsilon \sigma_{-i}^\omega \right) / (1 - \varepsilon) \right]^+ = \lambda^M \sigma_{-i}^f + (1 - \lambda^M) \sigma_{-i}^\omega$ and $\lambda^M = \sup \left\{ \lambda \in [1, 1/(1 - \varepsilon)] : \lambda \sigma_{-i}^f + (1 - \lambda) \sigma_{-i}^\omega \in \Delta_{-i} \right\}$. The probability measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ satisfies the conditions of Definition 3 (in particular, note that the non-atomic part is a product measure) and for all ε , $(1 - \varepsilon)\sigma^* + \varepsilon\sigma^\omega$ is an RBE. Finally, $(1 - \varepsilon)\sigma^* + \varepsilon\sigma^\omega$ converges to σ^* as $\varepsilon \rightarrow 0$. ■

Let $S_i^0 \subset S_i$ be the set of pure strategies of player i for which $D_i(s_i)$ has positive

Lebesgue measure and let $S^0 = \times_{i \in N} S_i^0$. If the belief measures have no atoms, then the support of any RBE must be contained in S^0 . This implies that the support of a robust equilibrium must also be contained in S^0 . In other words, a robust equilibrium will put zero probability on any pure strategy s_i that is a best reply only on a measure zero subset of Δ_{-i} . This explains why some perfect equilibria are not robust: a strategy s_i could be in the support of a perfect equilibrium strategy profile and, at the same time, be a best reply only on a set of measure zero.

Weakly dominated strategies never belong to S_i^0 . If a strategy s_i is weakly dominated, then there is some pure strategy profile $s_{-i} \in S_{-i}$ against which s_i is not a best reply and s_i is never a best reply to any mixed strategy profile $\sigma_{-i} \in \Delta_{-i}$ that places positive probability on s_{-i} . Consequently, the Lebesgue measure of $D_i(s_i)$ is zero. Thus, no pure strategy played with positive probability in a robust equilibrium can be weakly dominated. Since in two-person games an equilibrium is perfect if and only if no weakly dominated strategy is played with positive probability, the next lemma easily follows.

Lemma 1 *A robust equilibrium of a two-person game Γ is a perfect Nash equilibrium of Γ .*

In Theorem 5, we will use the game in Table 2 to show that the converse of Lemma 1 is false; not all perfect equilibria of two-player games are robust.

Suppose now that the strategy s_i is equivalent to a convex combination of strategies in a set $S_i^K \subset S_i$ containing at least two distinct strategies, then $D_i(s_i) \subset \bigcap_{s_i^k \in S_i^K} D_i(s_i^k)$ and by Lemma 2 in Appendix B the latter must have Lebesgue measure zero. This implies that two games with the same reduced normal form have the same set of robust equilibria.⁷

⁷The *reduced normal form* of a game Γ is the game that results from the semi-reduced normal form (see fn. 13) after deleting all pure strategies that are mixtures of other strategies. Formally, strategy s_i of player i is deleted if there exists a mixed strategy σ_i of i with $\sigma_i(s_i) = 0$ and such that $u_j(s_i, s_{-i}) = u_j(\sigma_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ and all $j \in N$.

| | | | | | |
|-----|---------|---------|-----|---------|---------|
| | L | R | | L | R |
| T | 1, 1, 1 | 1, 0, 1 | T | 1, 1, 0 | 0, 0, 0 |
| B | 1, 1, 1 | 0, 0, 1 | B | 0, 1, 0 | 1, 0, 0 |
| | ℓ | | | r | |

Table 1: Positive Measure Does Not Imply Robust

The next theorem characterizes robust equilibrium. It shows that a Nash equilibrium σ^* is robust if and only if, for any pure strategy s_i in the support of σ_i^* , the set of strategies in any small neighborhood around σ_{-i}^* to which s_i is a best reply has positive Lebesgue measure. Note that it is not sufficient for each strategy in the support of a Nash equilibrium σ to be a best reply to a set of positive Lebesgue measure for σ to be robust. This illustrated by the game in Table 1 where (B, L, ℓ) is a Nash equilibrium using only strategies that are best replies to a set of positive Lebesgue measure, but the equilibrium is not robust.⁸ To see this, denote by p_1 the probability that player 1 chooses T , by p_2 the probability that player 2 chooses L , and by p_3 the probability that player 3 chooses ℓ so that (B, L, ℓ) is represented as $(0, 1, 1)$. What drives the example is that the set

$$D_1(B) = \{(p_2, p_3) \in [0, 0.5]^2 \mid p_3 \leq (1 - 2p_2)/(2 - 3p_2)\} \cup \{(1, 1)\}$$

has positive measure; however, it consists of the union of two disconnected sets, one of which is the isolated point $(1, 1)$ corresponding to the equilibrium strategies of players 2 and 3. In preparation for Theorem 4 denote by $C_i(\sigma_i)$ the carrier of σ_i .⁹

Theorem 4 *Let σ^0 be a Nash equilibrium of Γ and, for $i \in N$, let $B_r(\sigma_{-i}^0)$ be an open ball of radius r around σ_{-i}^0 . Then σ^0 is a robust equilibrium if and only if there exists $r' > 0$ such that for all $0 < r < r'$ and for all i , $D_i(s_i) \cap B_r(\sigma_{-i}^0)$ has positive Lebesgue measure for all $s_i \in C_i(\sigma_i^0)$.*

⁸ (B, L, ℓ) is also not perfect.

⁹The carrier of a mixed strategy σ_i is the subset of S_i consisting of the pure strategies to which σ_i assigns positive probability.

Proof. (Only if) Let σ^0 be a Nash equilibrium of Γ and suppose that for all $r' > 0$ there exists $r'' < r'$ such that, for some $s_i^* \in C_i(\sigma_i^0)$, $D_i(s_i^*) \cap B_{r''}(\sigma_{-i}^0)$ has zero Lebesgue measure. Then, it is also true that $D_i(s_i^*) \cap B_r(\sigma_{-i}^0)$ has zero Lebesgue measure for all $r < r''$. We will show that σ^0 cannot be a robust equilibrium. Suppose, to the contrary, that σ^0 is robust. Then there must exist a sequence of RBE's $\{\sigma(\varepsilon)\}$ such that $\sigma(\varepsilon)$ converges to σ^0 as ε goes to zero. This implies that there must exist ε' such that, for all $\varepsilon < \varepsilon'$, (i) $\sigma_i(\varepsilon)(s_i^*) > 0$ and (ii) $\sigma_{-i}(\varepsilon) \in B_r(\sigma_{-i}^0)$ for all $r < r''$. However, as ε converges to zero, $\mu_\varepsilon^i(\sigma_{-i}(\varepsilon))(B_r(\sigma_{-i}^0))$ converges to one and hence $\mu_\varepsilon^i(\sigma_{-i}(\varepsilon))(D_i(s_i^*))$ converges to zero. Thus, the probability attached by player i to strategy s_i^* must converge to zero; that is, $\sigma_i(\varepsilon)(s_i^*)$ cannot converge to $\sigma_i^0(s_i^*) > 0$ and σ^0 cannot be a robust equilibrium.

(If) Let σ^0 be a Nash equilibrium of Γ and suppose that there exists r' such that for all $0 < r < r'$ and for all i , $D_i(s_i) \cap B_r(\sigma_{-i}^0)$ has positive Lebesgue measure for all s_i such that $\sigma_i^0(s_i) > 0$. We will show that σ^0 is a robust equilibrium by constructing a sequence of belief map profiles $\{\mu_\varepsilon\}_{\varepsilon \rightarrow 0}$ and an associated sequence $\{\sigma(\varepsilon)\}$ of RBE's converging to σ^0 . To simplify the exposition, we will only define the map μ_ε^i for the case in which its focus σ_{-i}^f is contained in a ball of radius $\varepsilon < r'/2$ around σ_{-i}^0 : $\sigma_{-i}^f \in B_\varepsilon(\sigma_{-i}^0)$. It is always possible to extend this map continuously to all possible foci. Note that $\sigma_{-i}^f \in B_\varepsilon(\sigma_{-i}^0)$ implies that $D_i(s_i) \cap B_{2\varepsilon}(\sigma_{-i}(\varepsilon))$ has positive Lebesgue measure for all s_i such that $\sigma_i^0(s_i) > 0$. Thus, for all ε , for all s_i such that $\sigma_i^0(s_i) > 0$ and for all $j \neq i$ we can find a closed, convex set $A_{j\varepsilon}(s_i, \sigma_{-i}^f) \in \Delta_j$ such that $A_\varepsilon(s_i, \sigma_{-i}^f) = \times_{j \neq i} A_{j\varepsilon}(s_i, \sigma_{-i}^f) \in \Delta_{-i}$ is a subset of $D_i(s_i) \cap B_{2\varepsilon}(\sigma_{-i}^f)$. Let k_i be the cardinality of $C_i(\sigma_i^0)$, the carrier of σ_i^0 . Let the measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ be composed of two parts. The first part is ω_i , the Lebesgue measure over Δ_{-i} ; this part has a total mass of ε . The second part puts mass $\sigma_i^0(s_i) - \varepsilon/k_i$ on the product set $A_\varepsilon(s_i, \sigma_{-i}^f)$ for all s_i such that $\sigma_i^0(s_i) > 0$. Note that μ_ε^i satisfies Definition 3 (in particular, it is a

| | L | R |
|-----|------|------|
| T | 2, 0 | 0, 1 |
| M | 1, 1 | 1, 1 |
| B | 0, 1 | 2, 0 |

Table 2: Perfect Does Not Imply Robust

product measure). For all i , define $\sigma_i(\varepsilon)(s_i)$ as follows

$$\sigma_i(\varepsilon)(s_i) = \begin{cases} \sigma_i^0(s_i) - \varepsilon/k_i + \varepsilon\omega_i(D_i(s_i)) & \text{if } s_i \in C_i(\sigma_i^0) \\ \varepsilon\omega_i(D_i(s_i)) & \text{if } s_i \notin C_i(\sigma_i^0) \end{cases}$$

It is clear that $\sigma(\varepsilon)$ is an RBE, because $\sigma_i(\varepsilon)$ is the expected best reply when the focus is $\sigma_{-i}(\varepsilon)$. Since $\sigma(\varepsilon)$ converges to σ^0 as ε goes to zero, σ^0 is a robust equilibrium. ■

The next theorem shows that perfect equilibria need not be robust and that if there are at least three players robust equilibria need not be perfect.

Theorem 5 *A perfect Nash equilibrium of a game Γ need not be robust; if the game has at least three players, then a robust equilibrium of Γ need not be perfect.*

Proof. (Perfect does not imply robust.) Consider the game in Table 2: In this game $((p, 1 - 2p, p); (0.5, 0.5))$ is a perfect equilibrium for all $p \in (0, 0.5)$, because no player has a weakly dominated strategy. However, only $((0.5, 0, 0.5), (0.5, 0.5))$ is an accumulation point of a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ and thus a robust equilibrium. To see this, note that M is only a best response to $(0.5, 0.5)$ which is a zero Lebesgue measure point in the set of strategies of player 2; T is the unique best reply when the probability of L exceeds 0.5 and R is the unique best reply when it is less than 0.5. Thus, in any RBE the probability on strategy M is bounded above by the mass on an atom on $(0.5, 0.5)$ in the belief distribution of player 1 about 2' s strategy. Any such atom must vanish in the limit and the probability attached to M must go to zero.

(Robust does not imply perfect.) Consider the game in Table 3: In this game $\sigma^{p,q}$

| | | | |
|-----|---------|---------|--|
| | L | R | |
| T | 1, 0, 1 | 1, 1, 1 | |
| B | 0, 0, 1 | 0, 1, 1 | |
| | ℓ | | |

| | | | |
|-----|---------|---------|--|
| | L | R | |
| T | 0, 1, 1 | 0, 0, 1 | |
| B | 1, 1, 1 | 1, 0, 1 | |
| | c | | |

| | | | |
|-----|---------|---------|--|
| | L | R | |
| T | 1, 1, 0 | 1, 0, 0 | |
| B | 0, 1, 0 | 0, 0, 0 | |
| | r | | |

Table 3: Robust Does Not Imply Perfect

defined by

$$\begin{aligned}
& [(\sigma_1(T), \sigma_1(B)), (\sigma_2(L), \sigma_2(R)), (\sigma_3(\ell), \sigma_3(c), \sigma_r(r))] \\
& = [(p, 1 - p), (q, 1 - q), (0.5, 0.5, 0)]
\end{aligned}$$

with $p, q \in (0, 1)$ is a robust equilibrium, but it is not a perfect equilibrium. To see that it is robust note that

$$\begin{aligned}
D_1(T) &= \{\sigma_{-1} \in \Delta_{-1} : \sigma_3(c) \leq 1/2\} \\
D_1(B) &= \{\sigma_{-1} \in \Delta_{-1} : \sigma_3(c) \geq 1/2\} \\
D_2(L) &= \{\sigma_{-2} \in \Delta_{-2} : \sigma_3(\ell) \leq 1/2\} \\
D_2(R) &= \{\sigma_{-2} \in \Delta_{-2} : \sigma_3(\ell) \geq 1/2\} \\
D_3(\ell) &= D_3(c) = \Delta_{-3}
\end{aligned}$$

and hence for all i , for all $p, q \in (0, 1)$ and for all s_i such $\sigma_i^{p,q}(s_i) > 0$, there exists $r' > 0$ such that $0 < r < r'$ implies that $D_i(s_i) \cap B_r(\sigma_i^{p,q})$ has positive Lebesgue measure. To see that $\sigma^{p,q}$ is not a perfect equilibrium consider the following completely mixed strategy profile

$$\sigma(\varepsilon) = ((p + \varepsilon_T, 1 - p - \varepsilon_T), (q + \varepsilon_L, 1 - q - \varepsilon_L), (0.5 - \varepsilon_\ell, 0.5 - \varepsilon_c, \varepsilon_\ell + \varepsilon_c))$$

where $\varepsilon_\ell + \varepsilon_c > 0$. Note that as $\varepsilon_h \rightarrow 0$, for $h = T, L, \ell, c$, $\sigma(\varepsilon) \rightarrow \sigma^{p,q}$. For $\sigma^{p,q}$ to be perfect, there must exist a sequence $\sigma(\varepsilon)$, such that (i) T and B are best replies for player 1 against $\sigma_{-1}(\varepsilon)$, (ii) L and R are best replies for player 2 against $\sigma_{-2}(\varepsilon)$,

and (iii) ℓ and c are best replies for player 3 against $\sigma_{-3}(\varepsilon)$. Condition (i) requires that $\sigma_3(c) = 1/2$, or $\varepsilon_c = 0$, while condition (ii) requires that $\sigma_3(\ell) = 1/2$, or $\varepsilon_\ell = 0$; thus, it must be $\varepsilon_\ell + \varepsilon_c = 0$, a contradiction. ■

Completely mixed equilibria satisfy many of the known refinements, e.g., they are strictly perfect and strictly proper (see van Damme [15]). It is interesting to note that although the game in Table 2 has several completely mixed equilibria, none of them are robust. Robust equilibria have an appealing robustness property. Any pure strategy that is chosen with positive probability in a robust equilibrium must be a best reply to the mixed strategies in a set of positive measure contained in a small neighborhood of the equilibrium.

In the proof of Theorem 5 $\sigma^* = [(p, 1 - p), (q, 1 - q), (0.5, 0.5, 0)]$ was shown to be a robust but not perfect equilibrium of the game in Table 3. The absence of perfection was demonstrated by showing that having a positive tremble on $\sigma_3(r)$ implies either $\sigma_3(\ell) < 0.5$ or $\sigma_3(c) < 0.5$. The latter implies that $(p, 1 - p)$ (for $p \in (0, 1)$) cannot be a best reply to the tremble and the former implies that $(q, 1 - q)$ (for $q \in (0, 1)$) cannot be a best reply to the tremble. There are two special features of this game. Player 3 has one dominated strategies and also two duplicate strategies. Might all robust equilibria be perfect in any game in which there are no duplicate strategies and no dominated strategies? The answer to this is no. It is possible to augment the game in Table 2 by adding a third strategy for players 1 and 2 and constructing the payoffs so that (i) player 3 has no duplicate or dominated strategies, (ii) $\sigma^* = [(p, 1 - p, 0), (q, 1 - q, 0), (0.5, 0.5, 0)]$ is a robust equilibrium, and (iii) σ^* is not perfect.

A weak restriction on the belief map of each player that eliminates pathological beliefs is that the belief map can be represented as a mixture of uniform distributions. It is easy to check that even with such a restriction not all robust equilibria are perfect.

Consider the following robust, but not perfect, equilibrium of the game in Table 3

$$\sigma^* = [(0.75, 0.25), (0.75, 0.25), (0.5, 0.5, 0)].$$

For $\varepsilon \in (0, 1)$, define the following belief map for player i

$$\mu_\varepsilon^i(\sigma_{-i}^f) = \varepsilon\omega(\Delta_{-i}) + (1 - \varepsilon)\omega(B_\varepsilon(\sigma_{-i}^f))$$

where $\omega(\Delta_{-i})$ is the Lebesgue probability measure over the set Δ_{-i} of mixed strategies of player i 's opponents and $\omega(B_\varepsilon(\sigma_{-i}^f))$ is the Lebesgue probability measure over $B_\varepsilon(\sigma_{-i}^f)$, the ball of radius ε around σ_{-i}^f in Δ_{-i} . Let $\mu_\varepsilon = (\mu_\varepsilon^1, \mu_\varepsilon^2, \mu_\varepsilon^3)$ and note that the robust equilibrium σ^* is also an RBE of $\langle \Gamma, \mu_\varepsilon \rangle$ for all ε . In other words, σ^* can be obtained as the limit-RBE of a sequence of random belief games in which, for all ε , the belief map of each player is a simple mixture of uniform distributions.

3.3 Interior Atoms, Limit-RBE and Perfection

Suppose we allow the belief distributions to have non-vanishing atoms in the interior of the strategy spaces, but we require atoms on the boundary to vanish in the limit. In this section we show that with such an assumption the set perfect equilibria is a subset of the set of limit-RBE.

Let $\alpha_{j\varepsilon}^i(\sigma_j^f) = \alpha_{j\varepsilon}^{iI}(\sigma_j^f) + \alpha_{j\varepsilon}^{iB}(\sigma_j^f)$ be the total probability mass on atoms, where $\alpha_{j\varepsilon}^{iI}(\sigma_j^f)$ is the total probability mass on atoms in the interior of Δ_j and $\alpha_{j\varepsilon}^{iB}(\sigma_j^f)$ is the total probability mass on atoms on the boundary of Δ_j . Let $\alpha_j^{iK}(\sigma_j^f) = \lim_{\varepsilon \rightarrow 0} \alpha_{j\varepsilon}^{iK}(\sigma_j^f)$ and $\alpha_0^K = \sup_{i,j,\sigma_j^f} \alpha_j^{iK}(\sigma_j^f)$, with $K \in \{I, B\}$, and let $\alpha_0 = \alpha_0^I + \alpha_0^B$. We are interested in the case in which $\alpha_0^B = 0$, so that atoms on the boundary vanish in the limit measures, but no restrictions are imposed on α_0^I .

Theorem 6 *If $\alpha_0^B = 0$ and there are no restrictions on α_0^I , then any perfect equilibrium of a game Γ is a limit-RBE. If Γ is a two-person game, then the set of limit-RBE*

coincides with the set of perfect equilibria.

Proof. First we show that any perfect equilibrium σ^0 of Γ is a limit-RBE when $\alpha_0^B = 0$ and no restrictions are imposed on α_0^I . Perfection of σ^0 implies that there exists a sequence $\{\sigma^P(k)\}_{k=1}^\infty$ of completely mixed strategies, each with an associated error profile $\varepsilon^k \gg 0$, that converges to σ^0 and such that, for all i and all k , σ_i^0 is a best reply to $\sigma_{-i}^P(k)$. In the error profile

$$\varepsilon^k = (\varepsilon_{1,1}^k, \dots, \varepsilon_{1,m_1}^k, \dots, \varepsilon_{n,1}^k, \dots, \varepsilon_{n,m_n}^k)$$

$\varepsilon_{i,j}^k > 0$ is the tremble associated with the j^{th} pure strategy of player i . As $k \rightarrow \infty$ the error profile converges to the profile with all zero errors.

Let $W_i \subset S_i$ be the pure strategies s_i of player i such that $D_i(s_i)$ has zero Lebesgue measure and $\sigma_i^0(s_i) = 0$. Construct $\sigma^U(k)$ from $\sigma^P(k)$ by removing all probability from pure strategies in W_i and redistributing it proportionately to the player's other pure strategies. Denoting $w_i(k) = \sum_{s_i \in W_i} \sigma_i^P(k)(s_i)$, then

$$\sigma_i^U(k)(s_i) = \begin{cases} 0 & \text{if } s_i \in W_i \\ \sigma_i^P(k)(s_i)/(1 - w_i(k)) & \text{if } s_i \notin W_i \end{cases}$$

By construction $\sigma^U(k) \rightarrow \sigma^0$ as $k \rightarrow \infty$.

For each k we now construct a belief map profile $\mu_k = (\mu_k^1, \dots, \mu_k^n)$ such that $\sigma^U(k)$ is an RBE. Let

$$\theta_k = \min_{\substack{i \in N \\ s_i \in S_i}} \{\varepsilon_{i,s_i}^k\}$$

be the smallest tremble in the error profile ε^k . Choose the non-atomic part of the belief measure of player i , with focus $\sigma_{-i}^U(k)$, as follows for all sets $D_i(s_i)$ having

positive Lebesgue measure:

$$\mu_k^i(\sigma_{-i}^U(k))(D_i(s_i)) = \begin{cases} \sigma_i^U(k)(s_i) & \text{if } \sigma_i^0(s_i) = 0 \\ \theta_k & \text{if } \sigma_i^0(s_i) > 0 \end{cases}$$

It is immaterial how the measure $\mu_k^i(\sigma_{-i}^U(k))$ is defined on each of the $D_i(s_i)$ as long as there are no atoms, the conditions of Definition 3 are satisfied, and the measures assigned to the sets $D_i(s_i)$ are as specified above. The remaining probability, $1 - \mu_k^i(\sigma_{-i}^U(k))(\Delta_i)$, is assigned to the interior atom $\sigma_{-i}^P(k)$.

The best reply to the measure $\mu_k^i(\sigma_{-i}^U(k))$ assigns probability $\sigma_i^U(k)(s_i)$ to any strategy that is not in the carrier of σ_i^0 . Meanwhile, the non-atomic part assigns probabilities to strategies in the carrier that are no greater than the probabilities under σ_i^0 . The remaining probability is on the atom $\sigma_{-i}^P(k)$. As all strategies in the carrier are best replies to $\sigma_{-i}^P(k)$, the remaining probability can be parcelled out to yield $\sigma_i^U(k)$ as a best reply to the belief measure $\mu_k^i(\sigma_{-i}^U(k))$, making $\sigma^U(k)$ an RBE.

It remains to show that for two-person games all limit-RBE are perfect equilibria. This follows from the fact that a weakly dominated strategy cannot be a best reply to a completely mixed strategy. Since non-vanishing atoms in the limit belief distributions can only be on completely mixed strategies, a weakly dominated strategy cannot be played in a limit-RBE. Since an equilibrium of a two-player game is perfect if and only if no weakly dominated strategy is used with positive probability, a limit-RBE must be perfect. ■

4 Analysis of Experimental Data

As we pointed out in the introduction, an important property of RBE (and of the QRE of McKelvey and Palfrey [6]) is that it yields a statistical framework for estimation using either field or experimental data. To show the potential of RBE for

empirical applications, in this section we estimate RBE using data from three sets of experiments in which subjects played a game with a unique, completely mixed, Nash equilibrium. In Section 4.1 we introduce the Dirichlet belief structure used in our empirical implementation. In Section 4.2 we present the logistic version of QRE introduced by McKelvey and Palfrey [6]. In Sections 4.3, 4.4, and 4.5 we estimate the Dirichlet RBE using data from the experiments in Ochs [8], McKelvey, Palfrey and Weber [7], and O’Neill [9]; we also compare the fit of our Dirichlet RBE estimates with the fit of the logistic QRE.

All three experiments utilize games with completely mixed Nash equilibria; hence, these are games in which no player has a weakly dominated strategy. In experimental games with weakly dominated strategies one always find a small proportion of players using them, while RBE predicts that this should not happen. Applying RBE to such games would force us to deal with a likelihood that is always equal to zero. Taking care of this and other empirical issues goes beyond the scope of the present paper and is better left to future research.¹⁰

4.1 Dirichlet Belief Structure

To implement RBE empirically, we specify a belief structure under which the beliefs of player i about player j ’s strategy $\sigma_j = (\sigma_j(1), \dots, \sigma_j(m_j)) \in \Delta_j$ have a Dirichlet distribution with parameters $\alpha_{jk}^i > 0$, $k = 1, \dots, m_j$. The density function of σ_j is

¹⁰Weakly dominated strategies can be dealt with by incorporating payoff errors as in QRE by building a model with both random beliefs and random payoffs. With Dirichlet beliefs and logistic payoff errors, this model contains two parameters (ε and λ); its equilibrium, called RBQRE, includes RBE and QRE as special cases. Integrating the quantal response function $b_i^Q(\sigma_{-i})$ with respect to the belief measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ yields the expected quantal response function $\psi_{\mu_\varepsilon^i}^Q(\sigma_{-i}^f)$ of player i . RBQRE is defined analogously to RBE. The strategy profile σ is a *random belief, quantal response equilibrium* or *RBQRE* if $\sigma = \left(\psi_{\mu_\varepsilon^1}^Q(\sigma_{-1}), \dots, \psi_{\mu_\varepsilon^n}^Q(\sigma_{-n})\right)$.

given by

$$f_{j\varepsilon_i}^i(\sigma_j) = \frac{\Gamma\left(\sum_{k=1}^{m_j} \alpha_{jk}^i\right)}{\prod_{k=1}^{m_j} \Gamma\left(\alpha_{jk}^i\right)} \prod_{k=1}^{m_j} \sigma_j(k)^{\alpha_{jk}^i - 1}$$

where $\Gamma(\cdot)$ is the gamma function. When σ_j is two-dimensional, then the random variable $\sigma_j(1)$ has a beta distribution. It is natural to assume that the focus of player i 's beliefs about player j , σ_j^f , coincides with the expected value of σ_j . That is

$$\sigma_j^f(k) = \frac{\alpha_{jk}^i}{\sum_{h=1}^{m_j} \alpha_{jh}^i}$$

Let ε_{ij}^2 be the average variance of the components of the vector σ_j (i.e., $m_j \varepsilon_{ij}^2$ is the trace of the covariance matrix), then

$$m_j \varepsilon_{ij}^2 = \frac{\left(\sum_{k=1}^{m_j} \alpha_{jk}^i\right)^2 - \sum_{k=1}^{m_j} (\alpha_{jk}^i)^2}{\left(\sum_{k=1}^{m_j} \alpha_{jk}^i\right)^2 \left(\sum_{k=1}^{m_j} \alpha_{jk}^i + 1\right)}$$

Simple algebra shows that if the beliefs of player i about player j 's strategy choice have a Dirichlet distribution with trace $m_j \varepsilon_{ij}^2$ and focus equal to the expected strategy, then they can be represented by the following probability density function

$$f_{j\varepsilon_i}^i(\sigma_j^f)(\sigma_j) = G \prod_{h=1}^{m_j} \sigma_j(h)^{-\sigma_j^f(h) - 1 + (1 - \sum_{k=1}^{m_j} \sigma_j^f(k)^2) \sigma_j^f(h) / m_j \varepsilon_{ij}^2}$$

where

$$G = \frac{\Gamma\left(-1 + \left(1 - \sum_{k=1}^{m_j} \sigma_j^f(k)^2\right) / m_j \varepsilon_{ij}^2\right)}{\prod_{h=1}^{m_j} \Gamma\left(-\sigma_j^f(h) + \left(1 - \sum_{k=1}^{m_j} \sigma_j^f(k)^2\right) \sigma_j^f(h) / m_j \varepsilon_{ij}^2\right)}$$

Using the Dirichlet distribution and defining

$$\mu_{j\varepsilon_i}^i(\sigma_j^f)(A) = \int_A f_{j\varepsilon}^i(\sigma_j^f)(\sigma_j) d\sigma_j$$

for all $A \subset \Delta_j$ we obtain a belief measure satisfying all the required conditions.

As $\varepsilon_{ij} \rightarrow 0$, the Dirichlet belief measure converges to the degenerate measure with all mass on the focus and RBE converges to Nash equilibrium. To simplify, we assume $\varepsilon_{ij} = \varepsilon$ for all i and j . Then the variance ε^2 is the only parameter to estimate.

4.2 Logistic QRE

McKelvey and Palfrey [6] suppose that player i 's payoff from pure strategy s_i in game Γ when the other players are playing σ_{-i} is given by $u_i(s_i, \sigma_{-i}) + \phi_i(s_i)$ where $\phi_i(s_i)$ is a random variable. The vector of random variables $\phi_i = (\phi_i(1), \dots, \phi_i(m_i))$ is assumed to have a joint density $f_\lambda^i(\phi_i)$ with zero mean; the parameter λ measures the concentration of the density around zero. Let $f_\lambda = (f_\lambda^1, \dots, f_\lambda^n)$ be the payoff error density profile. A *random payoff game* $\langle \Gamma, f_\lambda \rangle$ is a pair consisting of a finite game Γ and a payoff error profile. In a random payoff game, each player observes the realization of his own error vector ϕ_i and then chooses a best reply. (Players need not observe the payoff realizations of their opponents.) The error density f_λ^i induces a probability distribution $b_i^Q(\sigma_{-i}) = (b_i^Q(\sigma_{-i})(1), \dots, b_i^Q(\sigma_{-i})(m_i))$ over the choices of player i , where $b_i^Q(\sigma_{-i})(s_i)$ is the probability that player i chooses action s_i when the other players choose σ_{-i} . The probability distribution $b_i^Q(\sigma_{-i})$ is called the *quantal response* of player i to σ_{-i} .

In the logistic version of QRE, which is what has been used in applications, the error vector has an extreme value distribution and the quantal response function of player i is given by

$$b_i^Q(\sigma_{-i})(s_i) = \frac{\exp\{\lambda_i u_i(s_i, \sigma_{-i})\}}{\sum_{k=1}^{m_i} \exp\{\lambda_i u_i(k, \sigma_{-i})\}} \quad (3)$$

Note that when $\lambda = 0$, the player puts probability $1/m_i$ on each pure strategy. For $\lambda_i > 0$, strategies with higher expected payoffs receive higher probability and as $\lambda_i \rightarrow \infty$ the logistic quantal response function converges to a conventional best reply.

| | | |
|----------|----------|----------|
| | <i>L</i> | <i>R</i> |
| <i>T</i> | 45/26, 0 | 0, 25/18 |
| <i>B</i> | 0, 25/18 | 5/26, 0 |

| | | |
|----------|----------|----------|
| | <i>L</i> | <i>R</i> |
| <i>T</i> | 25/17, 0 | 0, 5/4 |
| <i>B</i> | 0, 5/4 | 25/68, 0 |

Table 4: Ochs' Experiments: Games 2 and 3

The mixed strategy profile σ is a *quantal response equilibrium* of the random payoff game $\langle \Gamma, f_\lambda \rangle$ if it satisfies $\sigma = \left(b_1^Q(\sigma_{-1}), \dots, b_n^Q(\sigma_{-n}) \right)$. Thus, QRE gives us the parameter λ to estimate from the data.

To appreciate the difference between RBE and QRE, it is useful to consider the game in Table 2. In this game the profile $((1/3, 1/3, 1/3), (1/2, 1/2))$ is a QRE for all possible values of the parameter λ . On the contrary any RBE will attach probability zero to strategy M , no matter what the value of the parameter ε . At a QRE, the strategy a player chooses with the highest probability is the one that does best against the opponents' average behavior; the second highest probability goes to the second best strategy, and so on. On the contrary, at an RBE the probability attached to strategy s_i of player i depends on the measure of the set of mixed strategies to which s_i is a best response, in a neighborhood of the equilibrium.

4.3 Ochs

Ochs [8] studied three, 2×2 games in three separate experimental sessions. For our purposes, the data pertaining to the first game is uninteresting and we do not analyze it, because the Nash equilibrium also corresponds to an RBE for all values of the variance parameter ε and to a QRE for all values of the error parameter λ . Games 2 and 3 are displayed in Table 4.¹¹ In Game 2 the unique Nash equilibrium is $[(0.5, 0.5), (0.1, 0.9)]$, and in Game 3 it is $[(0.5, 0.5), (0.2, 0.8)]$. Subjects were randomly assigned to be either the row or the column player; in each round of play a

¹¹In Session 2 (3) players played Game 2 (3) after some initial rounds of Game 1 and were paid at the end of the session according to a lottery procedure with a high (\$20) and a low (\$10) prize. The payoffs in Games 2 and 3 have been normalized in order to isolate the effect on expected payoff of a single round of play of the game.




Figure 1: The Focus as a Function of epsilon: Ochs' Game 2

row and a column player were randomly matched. In Session Two, 16 subjects played 56 rounds of Game 2. In Session Three, 16 subjects played 64 rounds of Game 3. Each round consisted of ten plays of a game.

Figure 1 shows the RBE as a function of ϵ in Game 2. The curve marked by T is the probability of choosing T for player 1 and the curve marked L is the probability of choosing L for player 2. The likelihood maximizing values appear on the vertical broken line at $\epsilon = 0.2331$. A comparable illustration of QRE for this game can be found in McKelvey and Palfrey [6, Figure 6]

Table 5 contains data and maximum likelihood estimates for Games 2 and 3. $-\mathcal{L}$ is the value of the negative log likelihood. In both games, RBE performs better than QRE. RBE and QRE are not nested equilibrium notions. However, we can combine them by formulating a two-parameter (ϵ and λ) equilibrium model (RBQRE) which includes RBE and QRE as special cases (see fn. 10). Interestingly, estimating this two-parameter model we find that QRE is rejected in favor of RBE at any significance

| | | Game 2 | Game 3 |
|------|----------------|---------|---------|
| Data | T | 0.5949 | 0.5416 |
| | L | 0.2583 | 0.3363 |
| Nash | T | 0.5000 | 0.500 |
| | L | 0.1000 | 0.200 |
| | $-\mathcal{L}$ | 6119.50 | 7078.61 |
| RBE | T | 0.6593 | 0.6090 |
| | L | 0.2565 | 0.3504 |
| | ε | 0.2331 | 0.2795 |
| | $-\mathcal{L}$ | 5621.95 | 6849.02 |
| QRE | T | 0.6638 | 0.6280 |
| | L | 0.2387 | 0.3239 |
| | λ | 2.3977 | 2.2228 |
| | $-\mathcal{L}$ | 5631.33 | 6879.75 |

Table 5: Estimates For Ochs, Games 2 and 3

| | | | | | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|------|-----|-----|-----|-----|
| | L | R | | L | R | | L | R | | L | R |
| T | 9,0 | 0,1 | T | 9,0 | 0,4 | T | 36,0 | 0,4 | T | 4,0 | 0,1 |
| B | 0,1 | 1,0 | B | 0,4 | 1,0 | B | 0,4 | 4,0 | B | 0,1 | 1,0 |

Table 6: McKelvey and Palfrey's Experiments: Games A, B, C, and D

level in both games. Nash equilibrium is always rejected at any significance level.

4.4 McKelvey, Palfrey and Weber

McKelvey, Palfrey and Weber [7] studied four 2×2 games, displayed in Table 8, that are similar to those studied by Ochs [8].¹² In Games A, B, and C, the Nash equilibrium is $[(0.5, 0.5), (0.1, 0.9)]$, while in Game D it is $[(0.5, 0.5), (0.2, 0.8)]$. The experiments were conducted in eight separate sessions. In each session 12 subjects played 50 iterations of each of two of the games in Table 8 in the order shown in Table 7. At the beginning of each session subjects were randomly assigned row or the

¹²McKelvey, Palfrey and Weber [7] are interested in studying the effects of payoff magnitude and heterogeneity on behavior. If λ is constant across different games, then QRE predicts that a change in the magnitude of payoffs has a big effect on behavior. They find little support for this hypothesis; the maximum likelihood values of λ vary substantially across games.

| | | | | | | | | |
|-------------|---|---|---|---|---|---|---|---|
| Session | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| First Game | A | B | B | C | A | C | A | D |
| Second Game | B | A | C | B | C | A | D | A |

Table 7: Order of Play in the Sessions of McKelvey, Palfrey and Weber’s Experiment

| Game | | A | B | C | D |
|------|----------------|---------|---------|---------|--------|
| Data | T | 0.643 | 0.630 | 0.594 | 0.550 |
| | L | 0.241 | 0.244 | 0.257 | 0.328 |
| Nash | T | 0.5000 | 0.5000 | 0.5000 | 0.5000 |
| | L | 0.1000 | 0.1000 | 0.1000 | 0.2000 |
| | $-\mathcal{L}$ | 2390.47 | 1601.56 | 1635.83 | 822.60 |
| RBE | T | 0.6580 | 0.6582 | 0.6593 | 0.6117 |
| | L | 0.2392 | 0.2411 | 0.2557 | 0.3360 |
| | ε | 0.2129 | 0.2151 | 0.2321 | 0.2589 |
| | $-\mathcal{L}$ | 2167.93 | 1459.64 | 1505.39 | 797.35 |
| QRE | T | 0.6895 | 0.7078 | 0.6297 | 0.5896 |
| | L | 0.1148 | 0.2120 | 0.1065 | 0.2098 |
| | λ | 5.3700 | 0.7551 | 1.9401 | 7.3425 |
| | $-\mathcal{L}$ | 2287.33 | 1477.74 | 1604.66 | 817.05 |

Table 8: Estimates For McKelvey, Palfrey and Weber, Games A,B,C and D

column players and each time a game was played, subjects were randomly matched. The data and maximum likelihood estimates for each of the four games are in Table 8

As in Ochs [8], RBE performs better than QRE. Estimating a nested, two-parameter, RBQRE model, we find that in all games QRE is rejected in favor of RBE at any significance level. Nash equilibrium is always rejected at any significance level by the two alternative equilibria.

4.5 O’Neill

In O’Neill’s [9] experiment, 25 pairs of subjects played the zero-sum game in Table 9 for 105 times in succession. The game has a unique Nash equilibrium with both players using the strategy $(.4, .2, .2, .2)$. Table 10 reports our maximum likelihood

| | | | | |
|-------|-------|-------|-------|-------|
| | B_1 | B_2 | B_3 | B_4 |
| A_1 | 5 | -5 | -5 | -5 |
| A_2 | -5 | -5 | 5 | 5 |
| A_3 | -5 | 5 | -5 | 5 |
| A_4 | -5 | 5 | 5 | -5 |

Table 9: O’Neill’s Game

| | Data | Nash | RBE | QRE |
|----------------|--------|----------|---------|---------|
| A_1 | 0.3615 | 0.4000 | 0.3626 | 0.3770 |
| A_2 | 0.2206 | 0.2000 | 0.2124 | 0.2077 |
| A_3 | 0.2152 | 0.2000 | 0.2125 | 0.2077 |
| A_4 | 0.2027 | 0.2000 | 0.2125 | 0.2077 |
| B_1 | 0.4263 | 0.4000 | 0.4093 | 0.4179 |
| B_2 | 0.2255 | 0.2000 | 0.1968 | 0.1940 |
| B_3 | 0.1790 | 0.2000 | 0.1969 | 0.1940 |
| B_4 | 0.1691 | 0.2000 | 0.1969 | 0.1940 |
| λ | | ∞ | | 1.2354 |
| ε | | 0 | 0.2536 | |
| $-\mathcal{L}$ | | 7016.12 | 7005.74 | 7004.18 |

Table 10: Estimates For O’Neill

estimates. In this game QRE fits better than RBE. Estimating RBQRE reveals that RBE is rejected by a maximum likelihood ratio test in favor of QRE, but only at the 10% significance level.

5 Concluding Comments

We have introduced a variant of Nash equilibrium, called random belief equilibrium, under which the beliefs of a player about the opponents’ strategies are randomly drawn from non-degenerate probability measures over the other players’ mixed strategies. These measures are centered around the other players’ equilibrium choices. Thus, the equilibrium beliefs of each player are sensitive to the equilibrium choices of the others, but are not so precise as under Nash equilibrium. We believe this to

be a very intuitive and appealing relaxation of the belief condition associated with Nash equilibrium. That condition, exact and correct beliefs, seems to us much too stringent.

We examined limiting equilibrium behavior as beliefs converge to the measure that puts all mass on the opponents' equilibrium choices, $\varepsilon \rightarrow 0$, and found that when atoms in the belief measures vanish, the limit equilibrium, called a robust equilibrium, refines perfect equilibrium for two-player games, but neither refines nor coarsens perfect equilibrium in general games. Robust equilibria utilize only strategies that are best replies to mixed strategies of the opponents belonging to sets of positive measure, implying an appealing robustness of the equilibrium. If, on the other hand, we do not impose restrictions on atoms in the limit belief measures, then the set of limit-RBE coincides with the set of Nash equilibria.

An important feature of RBE is that it yields a statistical framework which can be used to estimate experimental or field data. We fit RBE to three sets of experimental data that were also used by McKelvey and Palfrey. We find that random beliefs help to explain the subjects behavior; however, the same can be said for the random payoffs model (QRE). It is possible that some game situations are more naturally characterized by random beliefs and others by random payoffs; this is a matter for future research.

Appendix

We undertake three tasks here. First, in Appendix A we derive the expected best reply mapping Ψ_{μ_ε} when this mapping is not everywhere single-valued. Next, in Appendix B we show that the best reply mapping is single-valued whenever $\mu_\varepsilon^i(\sigma_{-i}^f)$ has no atoms and there are no indifferent strategies. Finally, in Appendix C we show that robustness is neither necessary nor sufficient for properness and that all essential equilibria are robust.

A The Best Reply Correspondence

Recall the definitions of $D_i(s_i)$ and $d_i(T_i)$ from equations (1) and (2), respectively. Any strategy that mixes over T_i and puts zero probability on strategies in $S_i \setminus T_i$ is a best reply to any profile $\sigma_{-i} \in D_i(T_i)$. This must be accounted for in constructing the expected best reply correspondence $\Psi_{\mu_\varepsilon^i}(\sigma_{-i})$. For $T_i \subset S_i$, $T_i \neq \emptyset$, the set $\Lambda_{T_i} \subset \Delta_i$, given by

$$\Lambda_{T_i} = \{\lambda_{T_i} \in \Delta_i \mid \lambda_{T_i}(s_i) = 0 \text{ if } s_i \notin T_i\}$$

is the set of all mixed strategies of player i that place zero probability on any strategy in $S_i \setminus T_i$. Thus σ_i is a best reply to $\sigma_{-i} \in d(T_i)$ if and only if $\sigma_i \in \Lambda_{T_i}$. Let $\Lambda_i = \times_{T_i \subset S_i, T_i \neq \emptyset} \Lambda_{T_i}$, and let λ_i be a generic element of Λ_i ; thus $\lambda_i = \{\lambda_{T_i}\}_{T_i \subset S_i}$ is a collection of strategy profiles of player i , one element for each non-empty $T_i \subset S_i$.

To find all the expected best replies of player i to the belief map μ_ε^i , first choose an arbitrary $\lambda_i \in \Lambda_i$ and an arbitrary focus σ_{-i} . This λ_i is used to find one specific expected best reply to $\mu_\varepsilon^i(\sigma_{-i})$:

$$\psi_{\mu_\varepsilon^i(\sigma_{-i})}^{\lambda_i}(s_i) = \sum_{\substack{T_i \subset S_i \\ T_i \neq \emptyset}} \lambda_{T_i}(s_i) \mu_\varepsilon^i(\sigma_{-i})(d_i(T_i)) \text{ for all } s_i \in S_i \quad (4)$$

and

$$\psi_{\mu_\varepsilon^i(\sigma_{-i})}^{\lambda_i} = (\psi_{\mu_\varepsilon^i(\sigma_{-i})}^{\lambda_i}(1), \dots, \psi_{\mu_\varepsilon^i(\sigma_{-i})}^{\lambda_i}(m_i)) \quad (5)$$

The set of expected best replies to $\mu_\varepsilon^i(\sigma_{-i})$ is

$$\Psi_{\mu_\varepsilon^i(\sigma_{-i})} = \{\psi_{\mu_\varepsilon^i(\sigma_{-i})}^{\lambda_i} \mid \lambda_i \in \Lambda_i\} \quad (6)$$

Thus, for each $i \in N$ equations (4) to (6) define an expected best reply correspondence $\Psi_{\mu_\varepsilon^i} : \Delta_{-i} \rightarrow \Delta_i$; this correspondence maps each focus $\sigma_{-i}^f \in \Delta_{-i}$ into the set of expected best replies to $\mu_\varepsilon^i(\sigma_{-i})$. Let $\Psi_{\mu_\varepsilon} : \Delta \rightarrow \Delta$ be the correspondence defined by $\Psi_{\mu_\varepsilon}(\sigma) = \times_{i \in N} \Psi_{\mu_\varepsilon^i}(\sigma_{-i})$.

B Indifferent Strategies

If the set $d_i(T_i)$ has zero Lebesgue measure, then any belief measure $\mu_\varepsilon^i(\sigma_{-i}^f)$ that satisfies Definition 3 and has no atoms must attach zero probability to the set of strategies that are completely mixed over precisely T_i . Thus, in this case if $d_i(T_i)$ has zero measure whenever T_i contains more than one strategy then the expected best reply $\Psi_{\mu_\varepsilon^i}(\sigma_{-i})$ is necessarily a singleton and will be written as a single-valued function $\psi_{\mu_\varepsilon^i}(\sigma_{-i})$. Lemma 2 below establishes that if player i has no indifferent strategies, then $d_i(T_i)$ has Lebesgue measure zero for all $T_i \subset S_i$ that have more than one element. Recall that two strategies, s_i and s'_i are *indifferent strategies* for player i if $u_i(s_i, s_{-i}) = u_i(s'_i, s_{-i})$ for all $s_{-i} \in S_{-i}$.¹³ If s_i and s'_i are not indifferent, then we say that they are *distinct strategies* for player i . Denote by $e_i^k \in \Delta_i$ the vector whose k -th element equals one and whose other elements are zero and let $e_L^k = (e_i^k)_{i \in L}$.

¹³Note that two indifferent strategies s_i and s'_i of player i may give different payoffs to the other players; that is, for some s_{-i} and some j it may be $u_j(s_i, s_{-i}) \neq u_j(s'_i, s_{-i})$. The strategies s_i and s'_i are called *duplicate strategies* if they give the same payoffs to all the players; that is, if $u_j(s_i, s_{-i}) = u_j(s'_i, s_{-i})$ for all $s_{-i} \in S_{-i}$ and for all $j \in N$. The *semi-reduced normal form* of Γ is the game that results after removing all duplications from Γ .

Lemma 2 *Let $\langle \Gamma, \mu_\varepsilon \rangle$ be a random belief game satisfying Definition 1 and parts (i), (ii) and (iv) of Definition 3. If player $i \in N$ has no indifferent strategies, then for all non-empty $T_i \subset S_i$, $d_i(T_i)$ has zero Lebesgue measure whenever $T_i \subset S_i$ is not a singleton. Hence, if for all players i the measures $\mu_\varepsilon^i(\sigma_{-i}^f)$ have no atoms and i has no indifferent strategies, then the expected best reply mapping is single-valued, and Ψ_{μ_ε} reduces to a single-valued function, $\psi_{\mu_\varepsilon} : \Delta \rightarrow \Delta$.*

Proof. Suppose, without loss of generality, that strategies 1 and 2 of player 1 are not indifferent, more precisely, suppose that $u_1(e_1^1, e_{-1}^1) \neq u_1(e_1^2, e_{-1}^1)$. Also suppose that, contrary to the lemma, there exists some $T_1 \subset N$ such that $1, 2 \in T_1$ and $d_1(T_1)$ has positive Lebesgue measure. Then there exists a completely mixed strategy profile $\sigma_{-1} \in D_1(1) \cap D_1(2)$ such that, for some positive ρ_0 and all $\rho \in (-\rho_0, \rho_0)$ the strategy profile $\sigma'_{-1} = \rho e_{-1}^1 + (1 - \rho)\sigma_{-1} \in D_1(1) \cap D_1(2)$. For $L \subset N \setminus \{1\}$ and $L^C = N \setminus (\{1\} \cup L)$ define

$$u_1(e_1^k, e_{L^C}^1, \sigma_L) \equiv \sum_{s_L \in S_L} u_1(e_1^k, e_{L^C}^1, s_L) \sigma_L(s_L) \quad (7)$$

where $S_L = \times_{i \in L} S_i$ and $\sigma_L(s_L) = \prod_{i \in L} \sigma_i(s_i)$ is the probability assigned to $s_L \in S_L$. Equation (7) is the payoff to player 1 when she uses pure strategy k , the players in L^C all use pure strategy 1, and the players in L use the mixed strategy profile σ_L . Denoting with $|L|$ the number of players in L and using equation (7), the payoff to player 1 when she uses pure strategy k , and the remaining players use the mixed profile σ'_{-1} is

$$u_1(e_1^k, \sigma'_{-1}) = \sum_{L \subset N \setminus \{1\}} \rho^{n-|L|-1} (1 - \rho)^{|L|} u_1(e_1^k, e_{L^C}^1, \sigma_L) \quad (8)$$

| | 1 | 2 |
|---|------|------|
| 1 | 2, 0 | 3, 1 |
| 2 | 1, 1 | 2, 1 |
| 3 | 1, 1 | 2, 0 |

Table 11: Indifferent Strategies

That $\sigma'_{-1} \in D_1(1) \cap D_1(2)$ implies that $u_1(e_1^1, \sigma'_{-1}) = u_1(e_1^2, \sigma'_{-1})$ which, with the aid of equation (8), implies

$$\begin{aligned}
& u_1(e_1^1, e_{-1}^1) - u_1(e_1^2, e_{-1}^1) \\
= & \sum_{\substack{L \subset N \setminus \{1\} \\ L \neq N \setminus \{1\}}} \left(\frac{1-\rho}{\rho} \right)^{|L|} [u_1(e_1^2, e_{LC}^1, \sigma_L) - u_1(e_1^1, e_{LC}^1, \sigma_L)] \quad (9)
\end{aligned}$$

The left hand side of equation (9) is non-zero by assumption; however, the right hand side cannot equal the left hand side for all $\rho \in (-\rho_0, \rho_0)$. To see this, note first that equality does not hold if (i) all the terms in square brackets on the right hand side are zero, or if (ii) for each coefficient $[(1-\rho)/\rho]^{|L|}$, the sum of terms sharing this coefficient equals zero. Second, if neither (i) nor (ii) holds then the right hand side changes in value as ρ varies and it must go to infinity in absolute value as $\rho \rightarrow 0$. Consequently, if $u_1(e_1^1, e_{-1}^1) \neq u_1(e_1^2, e_{-1}^1)$, then $d_i(T_i)$ has zero Lebesgue measure, which completes the proof. ■

The converse of the lemma is not true. In the game in Table 11 the sets $d_1(\{1, 2\})$, $d_1(\{1, 3\})$, $d_1(\{2, 3\})$, $d_1(\{1, 2, 3\})$ and $d_2(\{1, 2\})$ have all zero Lebesgue measure; however, strategies 2 and 3 are indifferent for player 1.

C Robustness, Properness and Essentiality

Note that in the game in Table 2 $[(p, 1-2p, p), (0.5, 0.5)]$ is a proper equilibrium for all $p \in (0, 0.5)$; however, we saw in the proof of Theorem 5 that only $[(0.5, 0, 0.5), (0.5, 0.5)]$

| | L | C | R |
|-----|--------|--------|--------|
| T | 1, 1 | 0, 0 | -8, -8 |
| M | 0, 0 | 0, 0 | -7, -7 |
| B | -8, -8 | -7, -7 | -7, -7 |

Table 12: Robust Does Not Imply Weakly Proper

| | L | C | R |
|-----|------|------|------|
| T | 2, 0 | 0, 1 | 0, 0 |
| M | 1, 1 | 1, 1 | 1, 0 |
| B | 0, 1 | 2, 0 | 0, 0 |

Table 13: Robust Equilibrium is not Invariant to Adding Dominated Strategies

is an accumulation point of a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \rightarrow 0}$ and thus a robust equilibrium. Thus, not all proper equilibria are robust. The converse result follows by noting that the pure strategy profile (M, C) of the game in Table 12 is perfect and robust, but not weakly proper.¹⁴ That robust equilibria need not be proper is unsurprising, because the random beliefs that a player i has about a player j say nothing about the relative probability that j uses strategy s_j versus strategy s'_j when neither is a best reply for player j .

The game in Table 13 is obtained by adding the dominated strategy R to the game in Table 2. By Theorem 4, in this game $[(p, 1 - 2p, p), (0.5, 0.5, 0)]$ is a robust equilibrium for all $p \in [0, 0.5]$. Any ball $B_r(0.5, 0.5, 0)$ around $(0.5, 0.5, 0)$ contains a subset of player 2's strategies with positive Lebesgue measure to which M is a best reply. This implies that after adding a dominated strategy the set of robust equilibria may expand. Note that the set of perfect, or proper, equilibria is not invariant either to adding or deleting dominated strategies.

Define the *payoff distance* between two games $\Gamma = \langle N, \Delta, u \rangle$ and $\Gamma' = \langle N, \Delta, u' \rangle$ as the maximal payoff difference between them: $\delta(\Gamma, \Gamma') = \max_{\substack{\sigma \in \Delta \\ i \in N}} |u_i(\sigma) - u'_i(\sigma)|$. The strategy profile σ^0 is an *essential equilibrium* of Γ if it is a Nash equilibrium of Γ and for all $\varepsilon > 0$ there exists a $\rho > 0$ such that every game with payoff distance

¹⁴The pure strategy equilibrium (T, L) is the only *stable set* of the game in Table 12. Thus, robust equilibria need not satisfy Kohlberg and Mertens' [4] definition of strategic stability.

less than ρ from Γ has a Nash equilibrium with distance less than ε from σ^0 . Van Damme [15] has shown that every essential equilibrium is strictly perfect (and hence perfect); we will now show that essential equilibria are also robust.

Lemma 3 *Every essential equilibrium is robust.*

Proof. Let σ^0 be an essential Nash equilibrium of $\Gamma = \langle N, \Delta, u \rangle$. We claim that there exists $r' > 0$ such that for all $0 < r < r'$ and for all i , $D_i(s_i) \cap B_r(\sigma_{-i}^0)$ has positive Lebesgue measure for all s_i such that $\sigma_i^0(s_i) > 0$. Then, by Theorem 4, σ^0 is robust. Suppose, to the contrary, that for all $r > 0$ there exist $i^* \in N$, $s_{i^*}^* \in S_{i^*}$ and $r'' < r'$ such that $\sigma_{i^*}^0(s_{i^*}^*) > 0$ and $D_{i^*}(s_{i^*}^*) \cap B_{r''}(\sigma_{-i^*}^0)$ has zero Lebesgue measure. Then $D_{i^*}(s_{i^*}^*) \cap B_\varepsilon(\sigma_{-i^*}^0)$ also has zero Lebesgue measure for all $\varepsilon < r''$. Define u_i^ρ as follows:

$$u_i^\rho(s_i, \sigma_{-i}) = \begin{cases} u_i(s_i, \sigma_{-i}) - \rho & \text{for } i = i^* \text{ and } s_{i^*} = s_{i^*}^* \\ u_i(s_i, \sigma_{-i}) & \text{otherwise} \end{cases}$$

Let $\Gamma(\rho) = \langle N, \Delta, u^\rho \rangle$, and note that $\delta(\Gamma, \Gamma(\rho)) = \rho$. In the game Γ , for all $\varepsilon < r''$, $s_{i^*}^*$ is only a best reply to a zero measure subset of $B_\varepsilon(\sigma_{-i^*}^0)$. Hence, by continuity of u_{i^*} , for all $\rho > 0$, $s_{i^*}^*$ is not a best reply to any element of $B_\varepsilon(\sigma_{-i^*}^0)$ in the game $\Gamma(\rho)$. This implies that any equilibrium σ' of $\Gamma(\rho)$ with $\sigma'_{-i^*} \in B_\varepsilon(\sigma_{-i^*}^0)$ must have $\sigma'_{i^*}(s_{i^*}^*) = 0$. Thus, there is no $\rho > 0$ such that the game $\Gamma(\rho)$ has a Nash equilibrium with distance from σ^0 less than $\max(\varepsilon, \sigma_{i^*}^0(s_{i^*}^*))$. Hence σ^0 cannot be an essential equilibrium of Γ . ■

Essential equilibria need not exist, which, together with Theorem 2, implies that not all robust equilibria are essential. Wu and Jiang [16] showed that all games with finitely many equilibria have at least one essential equilibrium. They also showed that the games in which all equilibria are essential form an open and dense subset of the set of games with a given strategy space S .

References

- [1] Bernheim, B. Douglas. (1984). “Rationalizable Strategic Behavior,” *Econometrica* **52**, 1007–1028.
- [2] Chen, Hsiao-Chi, Friedman, James W., and Thisse, Jacques-François. (1997). “Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach,” *Games and Economic Behavior* **18**, 32–54.
- [3] Haruvy, Ernan, Dale O. Stahl and Wilson, Paul W. (2000). “Modeling and Testing for Heterogeneity in Observed Strategic Behavior,” *Review of Economics and Statistics*. forthcoming.
- [4] Kohlberg, Elon and Mertens, Jean-François. (1986). “On the Strategic Stability of Equilibria,” *Econometrica* **54**, 1003–1037.
- [5] Luce, R. Duncan. (1959). *Individual Choice Behavior: A Theoretical Analysis*. New York: Wiley.
- [6] McKelvey, Richard D. and Palfrey, Thomas R. (1995). “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior* **10**, 6–38.
- [7] McKelvey, Richard D., Thomas R. Palfrey and Weber, Roberto A. (2000). “The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2×2 Games with Unique Mixed Strategy Equilibria,” *Journal of Economic Behavior and Organization* **42**, 523–548.
- [8] Ochs, Jack. (1995). “Games with Unique Mixed Strategy Equilibria: An Experimental Study,” *Games and Economic Behavior* **10**, 174–201.
- [9] O’Neill, Barry. (1987). “Nonmetric Test of the Minimax Theory of Two-Person Zerosum Games,” *Proceedings of the National Academy of Sciences, USA* **84**, 2106–2109.

- [10] Parthasarathy, K. R. (1967). *Probability Measures on Metric Spaces*. New York: Academic Press.
- [11] Pearce, David G. (1984). “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica* **52**, 1029–1050.
- [12] Selten, Reinhard. (1975). “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory* **4**, 25–55.
- [13] Stahl, Dale O. and Wilson, Paul W. (1994). “Experimental Evidence of Players’ Models of Other Players,” *Journal of Economic Behavior and Organization* **25**, 309–327.
- [14] Stahl, Dale O. and Wilson, Paul W. (1995). “On Players’ Models of Other Players Theory and Experimental Evidence,” *Games and Economic Behavior* **10**, 218–254.
- [15] van Damme, Eric. (1991). *Stability and Perfection of Nash Equilibria*, 2nd ed. Berlin: Springer-Verlag.
- [16] Wu, Wen-Tsün and Jiang, Jia-He. (1962). “Essential Equilibrium Points n-Person Non-Cooperative Games,” *Scientia Sinica* **11**, 1307–1322.