Random Belief and Ultraperfect Equilibria

Department of Economics Working Paper 99-03

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

November 16, 1999

Abstract

We examine noncooperative games in which players have random beliefs about the strategy choices that other players will make. Each player chooses a best response relative to her beliefs and, at equilibrium, the probability measure that represents her beliefs has some dependence on the other players' actual choices. These beliefs and the associated equilibrium are a generalization of the beliefs and equilibrium conditions underlying Nash equilibrium in games of complete information.

We study the limiting behavior of equilibrium as beliefs converge to certainty, finding that the limiting equilibrium is trembling hand perfect. We call such a limiting equilibrium ultraperfect and find the following: 1) not all perfect equilibria are ultraperfect, 2) being ultraperfect is neither necessary nor sufficient for being weakly proper.

1 Introduction

In a game of complete information Nash equilibrium satis…es two requirements. First, each player believes the other players will choose their equilibrium strategies. Second, each player's equilibrium strategy is a best reply to her beliefs. The combination of beliefs that are correct and behavior that is optimal relative to those beliefs de…nes Nash equilibrium. In this paper we retain intact the best reply requirement, but we reduce the precision of players' beliefs. We endow each player with probabilistic beliefs about what other players will do and require that in equilibrium beliefs be approximately correct on average. We call this a random belief equilibrium (RBE). As in Nash equilibrium, in an RBE beliefs are consistent with the actual choices made by the other players.

How players arrive at exactly correct beliefs and coordinate on a Nash equilibrium is a puzzle of game theory. A standard justi…cation is that the players will choose a Nash equilibrium, because all know each player's strategy space and payo¤ function, all are rational, and all of this is common knowledge. This justi…cation has some credibility in a game having a unique Nash equilibrium that is also strict. Take away uniqueness and, in general, one cannot be sure that all players have the same Nash equilibrium in mind when they choose strategies. If one player believes the others may have various Nash equilibria in mind, then she would presumably choose a best reply to those beliefs, which would not typically be a strategy associated with any Nash equilibrium. If Nash equilibrium were not strict, then any best reply to the equilibrium strategy pro…le would be a rational choice for a player. Thus, even when all players have the same Nash equilibrium in mind, they need not select equilibrium strategies even though each selects a best reply to the equilibrium pro…le.

Harsanyi [2] assumed that a player had some doubts about her own payo¤ function, permitting him to analyze the game as one of incomplete information in which each agent-normal player would have a unique best reply at equilibrium. This serves to justify a mixed strategy Nash equilibrium of the original complete information game while taking the view that any specific player will definitely choose a pure strategy. While this is quite ingenious, it does rest on the supposition that a player does not know her own payoff function. For games in which this is not a viable assumption, it provides no help, and it rests on the same perfect, precise expectations as the Nash equilibrium.

Even in the ideal situation of a unique and strict equilibrium, rationality need not imply that every player expects Nash equilibrium behavior. Bernheim [1] and Pearce [8] have taken a more relaxed view of what constitutes rationality; they say any pure strategy of a player is a rationalizable (hence rational) choice as long as it is a best reply to some profile that the other players in the game might use (i.e., that is rationalizable for the other player). Rationalizability is as agnostic as you can get about players' beliefs. The RBE is based on beliefs that are, in spirit, more precise than rationalizability beliefs and more sensitive to the equilibrium choices of the players, but they are less precise than Nash equilibrium beliefs and less sensitive to equilibrium choices.[1]

We look at games in which each player has no doubt about her own payoff function; however, each player i has beliefs, represented by a probability measure over the mixed strategy set of any other player j. We propose an equilibrium in which each player is choosing a strategy that is a best reply to her beliefs. Because the beliefs of a player are, in general, characterized by a non-degenerate probability measure over the other players' strategies, the resulting RBE is not a Nash equilibrium.

The limit behavior of the RBE as the dispersion of beliefs vanishes proves to be very interesting. To capture the responsiveness of the beliefs of player i over the mixed strategy profile, σ_{-i}, of the other players, we introduce a special mixed strategy profile σ^f, called the focus. A player's beliefs depend on the focus and, at equilibrium, the chosen strategy profile, σ^*, must equal the focus. Therefore, as the

[1] In contrast to rationalizability, a player may put positive probability on strategies that another player would never use; however, the RBE beliefs could easily be reformulated to put zero probability on such strategies.

dispersion of beliefs vanishes, the requirement that $\sigma^{*} = \sigma^{f}$ becomes precisely the Nash equilibrium condition that the players have exact beliefs that are correct and to which their respective strategies are best replies.

Thus, not surprisingly, the limit equilibrium is Nash; however, it is more than that. It is a subset of perfect equilibrium that rules out the use of any pure strategies that are best replies to only a very few strategy profiles. For a Nash equilibrium to be the limit of a sequence of RBE as the dispersion goes to zero, it must satisfy a robustness property: A strategy played with positive probability must be a best reply to a set of the other players' strategy profiles that has positive Lebesgue measure within any small ball around the equilibrium profile of the other players. This is not required by perfect equilibrium.

In Section 2 the model is described. We define random belief games, consisting of a game and a belief structure, random belief equilibrium (RBE), consisting of beliefs and a strategy profile that is both consistent with those beliefs and optimal relative to them, and ultraperfect equilibrium which is the limit of a sequence of RBE strategy profiles as the dispersion of beliefs goes to zero. In Section 3 we prove that ultraperfect equilibria exist, are perfect, that some perfect equilibria are not ultraperfect, and we provide a characterization of ultraperfect equilibrium. The relationship to both proper and essential equilibrium is also explored. Section 4 concludes.

2 The Model

A random belief game, defined in Section 2.2, is a pair consisting of (the mixed extension 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 $\mu^{i}_{j\varepsilon}(\sigma^{f}_{j})$ of player i concerning

	L	R
T	$a_1; a_2$	$b_1; c_2$
B	$c_1; b_2$	$d_1; d_2$

Table 1: A Generic 2 by 2 Game

the choice of rival player j is a probability measure over the set of j's mixed strategy profiles Σ_j that is parameterized by $\varepsilon \in \mathbb{R}_+$ and by $\sigma_j^f \in \Sigma_j$. Roughly speaking, the parameter σ_j^f, called the focus about j, represents the center of the probability measure and ε parameterizes its dispersion around σ_j^f.

In Section 2.3 random belief equilibrium is defined as an assessment which consists of a belief structure for a given ε and a strategy profile $\sigma = (\sigma_1; \ldots; \sigma_n)$ that is both an expected best reply to the belief structure and is consistent with it. Consistency means that, for each player i, the strategy σ_i equals the focus of all the other player's beliefs about i, $\sigma_i = \sigma_i^f$. Thus, at a random belief equilibrium the expected best reply of each player equals the focus about her strategy of all other players' beliefs. This consistency condition is a generalization 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 $\varepsilon \to 0$ the belief measures collapse to certainty so that, as shown in Section 2.4, the random belief game converges to a conventional game and the random belief equilibrium converges to a Nash equilibrium which is called an ultraperfect equilibrium. Existence and properties of ultraperfect equilibrium are deferred to Section 3.

We begin in Section 2.1 by presenting an example that illustrates the concepts of random belief games, random belief equilibrium, and ultraperfect equilibrium.

2.1 An Example

Consider the two by two game shown in Table 1. Each player chooses a mixed strategy, $p_i^\varepsilon \in [0; 1] = \Sigma_i$, where p_1^ε is the probability that player 1 chooses T and p_2^ε is the probability that player 2 chooses L. Let $M([0; 1])$ be the set of non-atomic probability measures on $[0;1]$ and $\mu_\varepsilon^i(x_j) \in M([0;1])$ be player i's belief measure concerning player j's mixed strategy p_j^C; $x_j \in [0;1] = \Sigma_j$ is the focus about player j's strategy. The focus is a measure of central tendency of the probability measure $\mu_\varepsilon^i(x_j)$. Thus $\mu_\varepsilon(x) = (\mu_\varepsilon^1(x_2); \mu_\varepsilon^2(x_1))$ is a specific pair of belief measures, $\mu_\varepsilon^i : \Sigma_j \to M([0;1])$ is a map, called a belief map, that associates a belief measure for player i to each focus $x_j \in [0;1] = \Sigma_j$ about player j, and $\mu_\varepsilon = (\mu_\varepsilon^1; \mu_\varepsilon^2)$ is a belief map profile.

Player i draws a belief realization p_j^R from $[0;1]$ according to $\mu_\varepsilon^i(x_j)$ and then chooses the mixed strategy that maximizes her expected payoff given beliefs p_j^R. The payoff of player i for an arbitrary $p_i^C \in \Sigma_i$ is

$$V_i(p_i^C; p_j^R) = p_i^C p_j^R a_i + p_i^C(1 - p_j^R) b_i + (1 - p_i^C) p_j^R c_i + (1 - p_i^C)(1 - p_j^R) d_i$$

Denote by $b_i(p_j^R)$ the conventional best reply mapping of player i.[2] That is

$$V_i(b_i(p_j^R); p_j^R) = \max_{p_i^C \in \Sigma_i} V_i(p_i^C; p_j^R)$$

The expected best reply of player i, $\theta_i(x_j)$, is the integral of $b_i(p_j^R)$ with respect to the belief measure $\mu_\varepsilon^i(x_j)$. That is,

$$\theta_i(x_j) = \int_{[0;1]} b_i(p_j^R)\, d\mu_\varepsilon^i(x_j)$$

The strategy profile $\theta(x) = (\theta_1(x_2); \theta_2(x_1))$ is the expected best reply to the profile of belief measures $\mu_\varepsilon(x) = (\mu_\varepsilon^1(x_2); \mu_\varepsilon^2(x_1))$. Note that $\theta : \Sigma \to \Sigma$. We define random belief equilibrium (RBE) as an assessment $\langle \mu_\varepsilon; p^0 \rangle$ where μ_ε is a belief map profile and p^0 is a fixed point of θ. Thus at an RBE

$$p_i^0 = \theta_i(p_j^0) = \int_{[0;1]} b_i(p_j^R)\, d\mu_\varepsilon^i(p_j^0)$$

[2] In a generic two by two game the best reply of player i is non-unique on, at most, one point in $[0;1]$ and hence $p_i^C(p_j^R)$ is unique on a set of measure one. We assume genericity to simplify the exposition of this example.

for $i, j = 1, 2$ and $i \neq j$.

At a random belief equilibrium the focus x_j equals the expected strategy of player j. Dependence of ν^i_ε upon x_j is merely a weakened version of the consistency condition embodied in Nash equilibrium. Two conditions characterize Nash equilibrium: First, each player chooses a strategy that is a best reply to what the player believes her rivals are doing and, second, these beliefs are correct. We weaken these requirements by having belief realizations randomly chosen and thus not necessarily correct. As in Nash equilibrium, players choose best replies to these belief realizations. The belief realization of each player i is drawn from a belief measure whose focus x_j at equilibrium equals the expected strategy $\mu_j(x_i)$ of the other player. We require ν^i_ε to satisfy several conditions: i) it is non-atomic, ii) its support is Σ_j, iii) it varies continuously with x_j, and iv) as $\varepsilon \to 0$ it converges to the degenerate measure with all mass on x_j.

Suppose for each $\varepsilon \in R_+$ that $\langle \nu_\varepsilon, p_\varepsilon \rangle$ is a random belief equilibrium and let $\varepsilon \to 0$. The sequence $\{p_\varepsilon\}$ necessarily has an accumulation point, p_0, that is a Nash equilibrium of the game in Table 1. This accumulation point is an ultraperfect equilibrium. As an illustration of the belief structure, for any $\varepsilon \in (0, 1]$ define the truncated normal probability density function $f_{j\varepsilon}(x_j)$ by

$$f_{j\varepsilon}(x_j)(y) = \frac{\exp\left\{-\frac{1}{2}\left(\frac{y-x_j}{\varepsilon}\right)^2\right\}}{\int_0^1 \exp\left\{-\frac{1}{2}\left(\frac{z-x_j}{\varepsilon}\right)^2\right\} dz}$$

Then let

$$\nu^i_{j\varepsilon}(x_j)(A) = \int_A f_{j\varepsilon}(x_j)(y)\,dy$$

These beliefs satisfy the required conditions.

It is useful to draw a comparison between RBE and the quantal response equilibrium (QRE) defined by McKelvey and Palfrey [5]. In a standard game, the expected payoff associated with a given pure strategy s_i of player i is a deterministic function

$u_i(s_i; \sigma_{-i})$ of the mixed strategies σ_{-i} chosen by i's rivals. McKelvey and Palfrey assume that the expected payoff is subject to a random error; that is, it is given by $u_i(s_i; \sigma_{-i}) + \varepsilon_i(s_i)$, where $\varepsilon_i(s_i)$ is an independent random variable. Players draw these random variables for each pure strategy and choose the strategy with the highest expected payoff realization. The distributions of random errors, together with the given profile of mixed strategies σ, generate a distribution of choices by the players; let $\sigma^C(\sigma)$ be the expectation of such a distribution. A QRE is a mixed strategy profile that is consistent in the sense of being equal to the expectation of its induced distribution of players choices, $\sigma^C(\sigma) = \sigma$. Thus, QRE and RBE are both statistical notions requiring average consistency of behavior when players choose best replies after observing the realization of random variables (expected payoffs for QRE, beliefs about rivals' choices for RBE).[3] The key difference is that the 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 the 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.

2.2 Random Belief Games

A random belief game is a pair $\langle \Gamma; \mu^* \rangle$ where Γ is a finite game and μ^* is a belief map profile. 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 $\hat{\Gamma} = \langle N; S; (U_i) \rangle$ is given by a finite set of players $N = \{1; \ldots; n\}$, a set of pure strategy profiles $S = \times_{i \in N} S_i$, where $S_i = \{1; \ldots; m_i\}$ is the finite pure strategy set of player i, and the utility profile $U = (U_1; \ldots; U_n)$, where $U_i : S \to \mathbb{R}$ is the payoff function of player i. Denote by $\Gamma = \langle N; \Sigma; (u_i) \rangle$ the

[3] See McKelvey and Palfrey [5] for a discussion of the connection between QRE and the literature on purification of Nash equilibria that originates from Harsanyi [2].

mixed extension of $\hat{\Gamma}$, where Σ_i is the set of mixed strategies of player i, $\Sigma = \times_{i \in N} \Sigma_i$, $u = (u_1, \dots, u_n)$, and $u_i : \Sigma \to \mathbb{R}$ is derived from U_i via von Neumann-Morgenstern utility.

Henceforth finite game means a mixed extension $\Gamma = \langle N, \Sigma, (u_i) \rangle$ as described above in Definition 1. Various conventional notations are used throughout the paper: $\Sigma_{-i} = \times_{j \neq i} \Sigma_j$, σ_{-i} is a typical element of Σ_{-i}, $S_{-i} = \times_{j \neq i} S_j$, and s_{-i} is a typical element of S_{-i}. For any subset of players $L \subset N$, $L^C = N \setminus L$ is the complement of L and $S_L = \times_{i \in L} S_i$ is the strategy space of the members of L. $|L|$ denotes the number of players in L. The probability assigned by σ_i to $s_i \in S_i$ is denoted $\sigma_i(s_i)$ and the probability assigned to $s_L \in S_L$ is $\sigma_L(s_L) = \prod_{i \in L} \sigma_i(s_i)$.

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

Each player i has a belief realization $\sigma_j^R \in \Sigma_j$ about the strategies the players $j \neq i$ will choose; each σ_j^R is chosen randomly according to the atomless belief measure $\mu^i_{j\varepsilon}(\sigma^f_j) \in M(\Sigma_j)$, where $\sigma^f_j \in \Sigma_j$ is the focus about player j, and $\varepsilon \in \mathbb{R}_+$ is a parameter. For any $\Sigma^0_j \subset \Sigma_j$, the probability that $\sigma^R_j \in \Sigma^0_j$ under $\mu^i_{j\varepsilon}(\sigma^f_j)$ is denoted $\mu^i_{j\varepsilon}(\sigma^f_j)(\Sigma^0_j)$. The belief measure $\mu^i_\varepsilon(\sigma^f_{-i})$ is the product measure of the $\mu^i_{j\varepsilon}(\sigma^f_j)$ for $j \neq i$. That is, let $B^i = \otimes_{j \neq i} B_j$, and $\mu^i_\varepsilon(\sigma^f_{-i})(A^i) = \times_{j \in N \setminus \{i\}} \mu^i_{j\varepsilon}(\sigma^f_j)(A_j)$ for any $A^i = \times_{j \neq i} A_j$ where $A_j \subset \Sigma_j$ and $\langle \Sigma_{-i}, B^i, \mu^i_\varepsilon(\sigma^f_{-i}) \rangle$ denotes the product space $\times_{j \neq i} \langle \Sigma_j, B_j, \mu^i_{j\varepsilon}(\sigma^f_j) \rangle$. Each player thus has a belief map $\mu^i_\varepsilon : \Sigma_{-i} \to M(\Sigma_{-i})$, where $M(\Sigma_{-i})$ is the subset of probability measures on Σ_{-i} for which the respective σ^R_j, for $j \neq i$, are independently distributed. The belief map μ^i_ε associates a belief measure to each focus σ^f_{-i}; $\mu_\varepsilon = (\mu^1_\varepsilon, \dots, \mu^n_\varepsilon)$ denotes a belief map profile.

Definition 3 specifies that the belief measures $\mu^i_{j\varepsilon}(\sigma^f_j)$ center in a weak sense around the focus σ^f_j by requiring that, as $\varepsilon \in \mathbb{R}_+$ goes to zero, $\mu^i_{j\varepsilon}(\sigma^f_j)$ weakly converges to the measure that puts all mass on σ_j^f.[4] The definition also requires continuity of $\mu^i_{j\varepsilon}(\sigma_j^f)$ with respect to σ_j^f, and that any subset of Σ_j has positive measure under $\mu^i_{j\varepsilon}(\sigma_j^f)$ if and only 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 B_j be the Borel sets of Σ_j. For each $\sigma_j^f \in \Sigma_j$, each $i \in N$, and each $j \neq i$, denote by $\langle \Sigma_j; B_j; \mu^i_{j\varepsilon}(\sigma_j^f)\rangle$ a collection of probability spaces parameterized by $\varepsilon \in R_+$ having the properties that (i) for any $\varepsilon \in R_{++}$ and any $\sigma_j^f \in \Sigma_j$, $\mu^i_{j\varepsilon}(\sigma_j^f)$ is absolutely continuous with respect to Lebesgue measure on Σ_j and $\mu^i_{j\varepsilon}(\sigma_j^f)(A) > 0$ whenever $A \in B_j$ has positive Lebesgue measure; (ii) for any $\varepsilon \in R_{++}$ and any $A \in B_j$, $\mu^i_{j\varepsilon}(\sigma_j^f)(A)$ is continuous with respect to $\sigma_j^f \in \Sigma_j$; (iii) as $\varepsilon \to 0$, $\mu^i_{j\varepsilon}(\sigma_j^f)$ weakly converges to $1_{\sigma_j^f}$, where $1_{\sigma_j^f}$ is the singular probability measure that puts all mass on σ_j^f.

For $\sigma^f \in \Sigma$, denote by $1_{\sigma^f_{-i}}$ the limit of $\mu^i_\varepsilon(\sigma^f_{-i})$ as $\varepsilon \to 0$ and let $1_{\sigma^f} = (1_{\sigma^f_{-1}}; \ldots; 1_{\sigma^f_{-n}})$.

A random belief game is a pair consisting of a finite game Γ and a belief map profile μ_ε. This is stated formally as Definition 4.

Definition 4 A random belief game is given by $\langle \Gamma; \mu_\varepsilon\rangle$ where $\Gamma = \langle N; \Sigma; (u_i)\rangle$ satisfies Definition 1 and μ_ε is a belief map profile satisfying (i) and (ii) of Definition 3.

2.3 Random Belief Equilibrium

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

[4] On weak convergence, see, for example, Parthasarathy [7]. 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 $\int f d\mu_\varepsilon \to \int f d\mu$ for all $f \in C(X)$. The topology generated by weak convergence is called the weak topology. The Prohorov metric p on $M(X)$ also generates the weak topology. For $\mu; \nu \in M(X)$, the Prohorov metric is

$$p(\mu;\nu) = \inf\{\delta > 0 \mid \mu(B) \le \nu(B^\delta) + \delta \text{ and } \nu(B) \le \mu(B^\delta) + \delta \text{ for all } B \in \mathcal{B}\}$$

where $B^\delta = \{x \in X \mid d(x;B) < \delta\}$.

reply mapping, $b_i : \Sigma_{-i} \to \Sigma_i$, is single-valued everywhere except on a set of measure zero, then integrating the best reply with respect to the belief measure $\mu^i_\varepsilon(\sigma^f_{-i})$ yields the expected best reply $\psi_{\mu^i_\varepsilon}(\sigma^f_{-i})$ of the player. Namely,

$$\psi_{\mu^i_\varepsilon}(\sigma^f_{-i}) = \int_{\Sigma_{-i}} b_i(\sigma^R_{-i}) d\mu^i_\varepsilon(\sigma^f_{-i})$$

In such a case, a random belief equilibrium is a pair $\langle \mu_\varepsilon, \sigma' \rangle$ such that

$$\sigma' = (\sigma'_1, \ldots, \sigma'_n) \in \psi_{\mu_\varepsilon}(\sigma') = [\psi_{\mu^1_\varepsilon}(\sigma'_{-1}), \ldots, \psi_{\mu^n_\varepsilon}(\sigma'_{-n})];$$

that is, for each player $i \in N$, σ'_i is an expected best reply to the belief measure with focus σ'_{-i}. However, if b_i is multi-valued on a set of positive measure, then we need to be clear about how to define espected best replies. We now proceed to provide such a precise specification by using the inverse of the conventional best reply mapping.

For the game Γ denote the inverse of the best reply correspondence of player i by

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

$D_i(s_i) \subset \Sigma_{-i}$ 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

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

The set $d_i(T_i) \subset \Sigma_{-i}$ consists of mixed strategy profiles of player i's opponents to which precisely the 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 Σ_{-i}; the sets $\{d_i(T_i)\}$ and $\{D_i(s_i)\}$ are related by

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

Recall that any strategy that mixes over T_i and puts zero probability on strategies in $S_i \backslash T_i$ 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 $\Psi_{\mu^i_\varepsilon}(\sigma_{-i})$. For $T_i \subset S_i$, $T_i \neq \emptyset$, the set $\Lambda_{T_i} \subset \Sigma_i$, given by

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

is the set of all mixed strategies of player i that place zero probability on any strategy in $S_i \backslash T_i$. 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 $\lambda_i = \{\lambda_{T_i}\}_{T_i \in 2^{S_i}}$ 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^i_ε, 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 $\mu^i_\varepsilon(\sigma_{-i})$:

$$\psi^{\lambda_i}_{\mu^i_\varepsilon(\sigma_{-i})}(s_i) = \sum_{\substack{T_i \subset S_i \\ T_i \neq \emptyset}} \lambda_{T_i}(s_i)\, \mu^i_\varepsilon(\sigma_{-i})(d_i(T_i)) \text{ for all } s_i \in S_i \tag{1}$$

and

$$\psi^{\lambda_i}_{\mu^i_\varepsilon(\sigma_{-i})} = (\psi^{\lambda_i}_{\mu^i_\varepsilon(\sigma_{-i})}(1), \ldots, \psi^{\lambda_i}_{\mu^i_\varepsilon(\sigma_{-i})}(m_i)) \tag{2}$$

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

$$\Psi_{\mu^i_\varepsilon}(\sigma_{-i}) = \{\psi^{\lambda_i}_{\mu^i_\varepsilon(\sigma_{-i})} \mid \lambda_i \in \Lambda_i\} \tag{3}$$

Thus, for each $i \in N$ equations (1) to (3) define an expected best reply correspondence $\Psi_{\mu^i_\varepsilon} : \Sigma_{-i} \to \Sigma_i$; this correspondence maps each focus $\sigma^f_{-i} \in \Sigma_{-i}$ into the set of expected best replies to $\mu^i_\varepsilon(\sigma_{-i})$. Let $\Psi_{\mu_\varepsilon} : \Sigma \to \Sigma$ be the correspondence defined by $\Psi_{\mu_\varepsilon}(\sigma) = \times_{i \in N} \Psi_{\mu^i_\varepsilon}(\sigma_{-i}) = [\Psi_{\mu^1_\varepsilon}(\sigma_{-1}); \ldots; \Psi_{\mu^n_\varepsilon}(\sigma_{-n})]$.

If the set $d_i(T_i)$ has zero Lebesgue measure, then any belief measure $\mu^i_\varepsilon(\sigma^f_{-i})$ that satisfies Definition 3 must attach zero probability mass to the set of strategies that are completely mixed over precisely T_i. Thus, if $d_i(T_i)$ has zero measure whenever T_i contains more than one strategy then the expected best reply $\beta_{\mu^i_\varepsilon}(\sigma_{-i})$ is necessarily a singleton and will be written as a single-valued function $\psi_{\mu^i_\varepsilon}(\sigma_{-i})$. Lemma 5 in Appendix A 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. 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 $s_{-i} \in S_{-i}$.[5] If s_i and s_i' are not indifferent, then we say that they are distinct strategies for player i.

An assessment is a pair $\langle \mu_\varepsilon; \sigma \rangle$ where μ_ε is the players' belief map profile and $\sigma \in \Sigma$ is a strategy profile. 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) and (ii) of Definition 3, $\langle \mu_\varepsilon; \sigma(\varepsilon) \rangle$ be an assessment, and β_{μ_ε} be the expected best reply correspondence given by equations (1) to (3). If $\sigma(\varepsilon) \in \beta_{\mu_\varepsilon}(\sigma(\varepsilon))$, then $\langle \mu_\varepsilon; \sigma(\varepsilon) \rangle$ is a random belief equilibrium or RBE.

Borrowing the term assessment from Kreps and Wilson's [4] sequential equilibrium is deliberate, because both sequential equilibrium and RBE consist of the players' beliefs and their strategy profile, both require some sort of consistency between beliefs and the strategy profile, and both require optimality of the latter with respect to the former. At an RBE, $\langle \mu_\varepsilon; \sigma(\varepsilon) \rangle$, for each $i \in N$ the strategy $\sigma_i(\varepsilon)$ is an expected best reply with the expectation taken over Σ_{-i} using the belief measure $\mu^i_\varepsilon[\sigma_{-i}(\varepsilon)]$.

[5] Note 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.

2.4 Ultraperfect Equilibrium

An ultraperfect equilibrium is a limit point of strategy profiles associated with random belief equilibria as $\varepsilon \to 0$. Let $M = \{\mu_\varepsilon\}_{\varepsilon \to 0}$ be a sequence of belief map profiles $\{\mu_{\varepsilon_h}\}_{h=1}^{\infty}$ with the property that $\lim_{h \to \infty} \varepsilon_h = 0$ and denote by $\langle \Gamma; M \rangle = \{\langle \Gamma; \mu_\varepsilon \rangle\}_{\varepsilon \to 0}$ 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 $\langle \mu_\varepsilon; \sigma(\varepsilon) \rangle$ be a random belief equilibrium of $\langle \Gamma; \mu_\varepsilon \rangle$. A cluster point $\langle 1_{\sigma^0}; \sigma^0 \rangle$ of the sequence of random belief equilibria $\{\langle \mu_\varepsilon; \sigma(\varepsilon) \rangle\}_{\varepsilon \to 0}$ is an M-ultraperfect assessment. A strategy profile σ^0 is an ultraperfect equilibrium if $\langle 1_{\sigma^0}; \sigma^0 \rangle$ is an M-ultraperfect assessment for some M.

Below we will see that any ultraperfect equilibrium is a perfect Nash equilibrium of Γ. Note that σ^0 is ultraperfect as long as one can find a sequence of belief map profiles M with respect to which $\langle 1_{\sigma^0}; \sigma^0 \rangle$ is an M-ultraperfect assessment. This is reminiscent of the specification of Selten's [9] perfect equilibrium; recall that σ^0 is a perfect Nash equilibrium of Γ as long as one can find a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \to 0}$ that converges to σ^0 and such that $\sigma(\varepsilon)$ is an $\varepsilon-$perfect equilibrium of Γ for all values of ε. The strategy profile $\sigma(\varepsilon)$ is an $\varepsilon-$perfect equilibrium of Γ if it is completely mixed and satisfies the property that if $u_i(s_i^0; \sigma_{-i}(\varepsilon)) < u_i(s_i; \sigma_{-i}(\varepsilon))$ then $\sigma_i(\varepsilon)(s_i^0) \leq \varepsilon$ for all i and $s_i^0; s_i \in S_i$. In an $\varepsilon-$perfect equilibrium 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 $\mu^i_{j\varepsilon}(\sigma_j)$ associated with an RBE are random; that is, trembles occur in each player's beliefs about her opponents' strategies.

3 Existence and Characterization

Section 3.1 establishes the existence of both random belief and ultraperfect equilibria. Section 3.2 first shows that every ultraperfect equilibrium is a perfect Nash equilibrium, then provides a complete characterization of ultraperfect equilibrium. Section 3.3 shows that perfect equilibria need not be ultraperfect and that ultraperfection is neither necessary nor sufficient for properness. Section ?? shows that all essential equilibria are ultraperfect.

3.1 Existence of Random Belief and Ultraperfect Equilibra

Lemma 1 establishes that any random belief game $\langle\Gamma;\mu_\varepsilon\rangle$ has a random belief equilibrium. Theorem 1 then shows that any such game has an ultraperfect equilibrium.

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

Proof. An equilibrium is an assessment $\langle\mu_\varepsilon;\sigma(\varepsilon)\rangle$ that consists of a belief map profile μ_ε and a strategy profile $\sigma(\varepsilon)$ that is a fixed point of $\Phi_{\mu_\varepsilon}:\Sigma\to\Sigma$. It is easy to verify that Φ_{μ_ε} satisfies the conditions of the Kakutani fixed point theorem and hence that there exists $\sigma(\varepsilon)\in\Phi_{\mu_\varepsilon}(\sigma(\varepsilon))$. First, Σ is a compact and convex set. Second, by Equations (1)-(3), $\Phi_{\mu_\varepsilon}(\sigma)$ is a non-empty and convex set for all $\sigma\in\Sigma$. Finally, from Definition 3, for any $d_i(T_i)\subset\Sigma_{-i}$, $\mu^i_\varepsilon(\sigma_{-i})(d_i(T_i))$ is continuous with respect to $\sigma_{-i}\in\Sigma_{-i}$ and thus Φ_{μ_ε} is upper hemicontinuous. ■

To prove existence of an ultraperfect equilibrium we construct a sequence of random belief equilibria $\{\mu_\varepsilon;\sigma(\varepsilon)\}_{\varepsilon\to 0}$ to show that the associated sequence of strategy profiles $\{\sigma(\varepsilon)\}_{\varepsilon\to 0}$ has a cluster point σ^0. This cluster point is an ultraperfect equilibrium.

Theorem 1 Every game Γ satisfying Definitions 1 and 3 has an ultraperfect equilibrium.

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

3.2 Characterization of Ultraperfect Equilibrium

We begin by proving that all ultraperfect equilibria are perfect. Then we show that an ultraperfect equilibrium will put zero probability on any pure strategy s_i that is a best reply only on a measure zero subset of Σ_{-i} (i.e., $\mu^{-i}_{\varepsilon}(\sigma_{-i})[D_i(s_i)] = 0$ for all $\sigma_{-i} \in \Sigma_{-i}$). This explains why not all perfect (or proper) equilibria are ultraperfect: a strategy s_i could be in the support of a perfect (proper) equilibrium strategy profile and, at the same time, be a best reply only on a set of measure zero.

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

Theorem 2 Every ultraperfect equilibrium is a perfect Nash equilibrium.

Proof. Let $m_0 = \max_{i \in N} m_i$. For any ultraperfect equilibrium σ^0, consider a sequence $\{\langle \mu_\varepsilon; \sigma(\varepsilon)\rangle\}_{\varepsilon \in E}$ of RBE's where $E = \{\varepsilon_\ell\}_{\ell=1,2,\ldots}$, $\varepsilon_1 < 1/m_0$, $\varepsilon_\ell > \varepsilon_{\ell+1}$ for all ℓ, $\varepsilon_\ell \to 0$ as $\ell \to \infty$, and $\sigma(\varepsilon) \to \sigma^0$. Denote by $C_i(\sigma_i)$ the carrier of σ_i and by $k_i(\varepsilon)$ the cardinality of $S_i \setminus C_i(\sigma_i(\varepsilon))$.[7] Define the sequence $\{\sigma^*(\varepsilon)\}_{\varepsilon \in E}$ by $\sigma^*_i(\varepsilon)(s_i) = (1 - \varepsilon k_i(\varepsilon))\sigma_i(\varepsilon)(s_i)$ if $s_i \in C_i(\sigma_i(\varepsilon))$ and $\sigma^*_i(\varepsilon)(s_i) = \varepsilon$ if $s_i \notin C_i(\sigma_i(\varepsilon))$ for all $i \in N$. Thus $\sigma^*(\varepsilon)$ is a profile of completely mixed strategies.

---

[6] This 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.

[7] The 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.

For any s_i that is not a best reply to σ^0_{-i}, there must be an open neighborhood of σ^0_{-i} within which s_i is never a best reply. More formally, since the sets $D_i(s_i)$ are closed, we can find r^* such that for $r < r^*$ the ball $B_r(\sigma^0_{-i}) \subset \Sigma_{-i}$ of radius $r > 0$ around σ^0_{-i} satisfies $B_r(\sigma^0_{-i}) \cap D_i(s_i) = \emptyset$ for all s_i such that $\sigma^0_{-i} \notin D_i(s_i)$. Let the integer ℓ^* be such that $r^* > \varepsilon_{\ell^*}$. For any $\xi > 0$ there exists a finite integer $\ell(\xi) > \ell^*$ such that for $\ell \geq \ell(\xi)$, $\sigma_{-i}(\varepsilon_\ell) \in B_r(\sigma^0_{-i})$ and, by weak convergence of $\mu^i_{\varepsilon_\ell}(\sigma_{-i})$ to $1_{\sigma_{-i}}$, $\mu^i_{\varepsilon_\ell}(\sigma_{-i}(\varepsilon_\ell))(B_r(\sigma^0_{-i})) > 1 - \xi$. This implies that

$$\sigma_i(\varepsilon_\ell)(s_i) < \xi \text{ for any } s_i \text{ satisfying } B_r(\sigma^0_{-i}) \cap D_i(s_i) = \emptyset \text{ and for all } \ell \geq \ell(\xi) \quad (4)$$

Thus, in particular, pure strategies that are not best replies against $\sigma_{-i}(\varepsilon_{\ell(\xi)})$ have probability of at most ξ in $\sigma_i(\varepsilon_{\ell(\xi)})$. This implies that $\sigma^*_i(\varepsilon_{\ell(\xi)})(s_i) < \max(\xi; \varepsilon_{\ell(\xi)})$ for any s_i such that $\sigma^0_{-i} \notin D_i(s_i)$. Furthermore, there exists ξ^0 such that, if $\xi < \xi^0$, then $\sigma^*_{-i}(\varepsilon_{\ell(\xi)}) \in B_r(\sigma^0_{-i})$ for all $r < r^*$. Hence, pure strategies that are not best replies against $\sigma^*_{-i}(\varepsilon_{\ell(\xi)})$ receive probability of at most $\phi(\xi) = \max(\xi; \varepsilon_{\ell(\xi)})$ in $\sigma^*_i(\varepsilon_{\ell(\xi)})$; that is, $\sigma^*(\varepsilon_{\ell(\xi)})$ is a $\phi(\xi)$-perfect equilibrium of Γ. By letting $\xi \to 0$, $\phi(\xi)$ converges to zero and $\sigma^*(\varepsilon_{\ell(\xi)})$ converges to σ^0. This implies that σ^0 is a limit point of a sequence of $\phi(\xi)$-perfect equilibria of Γ, with $\phi(\xi) \to 0$, and thus it is a perfect equilibrium of Γ. ■

Let $S^0_i \subset S_i$ be the set of pure strategies of player i for which $D_i(s_i)$ has positive Lebesgue measure and let $S^0 = \times_{i \in N} S^0_i$. Clearly, the support of any RBE, and hence of any ultraperfect equilibrium σ^0 must be contained in S^0.

Weakly dominated strategies never belong to S^0_i. 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, furthermore, s_i is never a best reply to any mixed strategy profile $\sigma_{-i} \in \Sigma_{-i}$ 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 an ultraperfect equilibrium can be weakly dominated.

	L	R
T	1; 1; 1	1; 0; 1
B	1; 1; 1	0; 0; 1

ℓ

	L	R
T	1; 1; 0	0; 0; 0
B	0; 1; 0	1; 0; 0

r

Table 2: Positive Measure Does Not Imply Perfect

Similarly, if s_i is a convex combination of strategies in a set $S_i^K \subset S_i$ containing at least two distinct strategies, then $D_i(s_i) \subset \cap_{s_i^k \in S_i^K} D_i(s_i^k)$ and by Lemma 5 the latter must have Lebesgue measure zero. This implies that two games with the same reduced normal form have the same set of ultraperfect equilibria.[8]

That each strategy in the support of a Nash equilibrium σ is a best reply to a set of positive Lebesgue measure is not sufficient for σ to be ultraperfect. This illustrated by the game in Table 2 where $(B; L; \ell)$ is a Nash equilibrium using only strategies that are best replies to a set of positive Lebesgue measure, but the equilibrium is not perfect (hence not ultraperfect). 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 $(B; L; \ell)$ is represented as $(0; 1; 1)$. What drives the example is that the set

$$D_1(B) = \{(p_2; p_3) \in [0; 0.5]^2 | p_3 \le (1 - 2p_2)/(2 - 3p_2)\} \cup \{(1; 1)\}$$

has positive measure and 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.

As shown by the next theorem, for a Nash equilibrium σ^* to be ultraperfect, 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 must have positive Lebesgue measure. This

[8] The reduced normal form of a game Γ is the game that results from the semi-reduced normal form (see fn.5) 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 $j \in N$.

condition fails for the Nash equilibrium $(0, 1, 1)$ of Table 2 because $(1, 1)$ is an isolated element of $D_1(B)$. At the same time, $D_1(B)$ has positive measure, because B is a best reply to very small values of (p_2, p_3).

Theorem 3 Let σ^0 be a Nash equilibrium of Γ and, for $i \in N$, let $B_r(\sigma^0_{-i})$ be an open ball of radius r around σ^0_{-i}. Then σ^0 is ultraperfect if and only if there exists $r' > 0$ such that for all $0 < r < r'$ and for all i, $D_i(s_i) \cap B_r(\sigma^0_{-i})$ has positive Lebesgue measure for all $s_i \in C_i(\sigma^0_i)$.

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 $s_i \in C_i(\sigma^0_i)$, $D_i(s_i) \cap B_{r''}(\sigma^0_{-i})$ has zero Lebesgue measure. Then, it is also true that $D_i(s_i) \cap B_r(\sigma^0_{-i})$ has zero Lebesgue measure for all $r < r''$. We will show that σ^0 cannot be an ultraperfect equilibrium. Suppose, to the contrary, that σ^0 is ultraperfect. Then there must exist a sequence of RBE's $\{\langle \mu_\varepsilon, \sigma(\varepsilon)\rangle\}$ 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) $\sigma_{-i}(\varepsilon) \in B_r(\sigma^0_{-i})$ for all $r < r''$. However, as ε converges to zero, $\mu^i_\varepsilon(\sigma_{-i}(\varepsilon))(B_r(\sigma^0_{-i}))$ converges to one and hence $\mu^i_\varepsilon(\sigma_{-i}(\varepsilon))(D_i(s_i))$ 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 $\sigma^0_i(s_i) > 0$ and σ^0 cannot be an ultraperfect 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, $D_i(s_i) \cap B_r(\sigma^0_{-i})$ has positive Lebesgue measure for all s_i such that $\sigma^0_i(s_i) > 0$. We will show that σ^0 is an ultraperfect equilibrium by constructing a sequence $\{\sigma^0(\varepsilon)\}$ of RBE's converging to σ^0. Let ω^i be the Lebesgue measure over Σ_{-i} and

$$S^0_{i\varepsilon}(\sigma_{-i}) = \{s_i \in S^0_i : \omega_i[D_i(s_i) \cap B_r(\sigma_{-i})] \neq 0 \text{ for all } 0 < r < \varepsilon\}$$

$S^0_{i\varepsilon}(\sigma_{-i})$ is the set of player i's strategies that, for all $0 < r < \varepsilon$, are best replies to a subset of $B_r(\sigma_{-i})$ with positive Lebesgue measure. For all $s_i \in S^0_i$ and $A \subset \Sigma_{-i}$, let

$$\omega^i_D(s_i)(A) = \frac{\omega^i[A \cap D_i(s_i)]}{\omega^i[D_i(s_i)]}$$

and for all $s_i \in S^0_{i\varepsilon}(\sigma_{-i})$, $\varepsilon > 0$, $\sigma_{-i} \in \Sigma_{-i}$, and $A \subset \Sigma_{-i}$, let

$$\omega^i_\varepsilon(s_i;\sigma_{-i})(A) = \frac{\omega^i[A \cap D_i(s_i) \cap B_\varepsilon(\sigma_{-i})]}{\omega^i[D_i(s_i) \cap B_\varepsilon(\sigma_{-i})]}$$

Thus $\omega^i_D(s_i)$ is the uniform probability measure with support $D_i(s_i)$ and $\omega^i_\varepsilon(s_i;\sigma_{-i})$ is the uniform probability measure with support $D_i(s_i) \cap B_\varepsilon(\sigma_{-i})$. Finally, let

$$\eta_{i\varepsilon}(s_i;\sigma_{-i}) = \sigma^0_i(s_i) \text{ if } s_i \in S^0_{i\varepsilon}(\sigma_{-i}) \cap C_i(\sigma^0_i) \text{ and}$$

$$\eta_{i\varepsilon}(s_i;\sigma_{-i}) = \frac{\varepsilon}{2m^{C}_{i\varepsilon}(\sigma_{-i};\sigma^0_i) + 2} \text{ if } s_i \in S^0_{i\varepsilon}(\sigma_{-i}) \setminus C_i(\sigma^0_i)$$

where $m^{C_i(\sigma^0_i)}_{\varepsilon;\sigma_{-i}}$ is the cardinality of $S^0_{i\varepsilon}(\sigma_{-i}) \setminus C_i(\sigma^0_i)$ and m^0_i is the cardinality of S^0_i. For each player i let μ^i_ε be the belief map defined by:

$$\mu^i_\varepsilon(\sigma_{-i})(A) = \frac{\varepsilon}{2(1+\varepsilon)m^0_i}\left[\frac{2 + m^{C_i(\sigma^0_i)}_{\varepsilon;\sigma_{-i}}}{1 + m^{C_i(\sigma^0_i)}_{\varepsilon;\sigma_{-i}}}\right]\sum_{s_i \in S^0_i}\omega^i_D(s_i)(A) + \frac{1}{1+\varepsilon}\sum_{s_i \in S^0_{i\varepsilon}(\sigma_{-i})}\eta_{i\varepsilon}(s_i;\sigma_{-i})\,\omega^i_\varepsilon(s_i;\sigma_{-i})(A) \tag{5}$$

Note that μ^i_ε satisfies Definition 3. There exists ε' such that $C_i(\sigma^0_i) \subset S^0_{i\varepsilon}(\sigma^0_{-i})$ for all $\varepsilon < \varepsilon'$, since, by assumption, for all $0 < r < r'$, $D_i(s_i) \cap B_r(\sigma^0_{-i})$ has positive Lebesgue measure for all $s_i \in C_i(\sigma^0_i)$. Let $\varepsilon'' < \varepsilon'$ be such that, for any $\varepsilon < \varepsilon''$, if $\sigma_{-i} \in B_\varepsilon(\sigma^0_{-i})$ then $S^0_{i\varepsilon}(\sigma_{-i}) = S^0_{i\varepsilon}(\sigma^0_{-i})$. When the focus is $\sigma_{-i} \in B_\varepsilon(\sigma^0_{-i})$, the expected best reply by player i is given by $\psi_{\mu^i(\sigma_{-i})} = \sigma^0_i(\varepsilon)$; however, $\sigma^0_i(\varepsilon)(s_i) = \mu^i_\varepsilon(\sigma_{-i})[D_i(s_i)]$. When $A = D_i(s_i)$ equation (5) is simplified, because $\omega^i_D(s_i)[D_i(s_i)] = 1$ for all $s_i \in S^0_i$ and

$\mu^{-}_{\varepsilon}(s_i;\sigma_{-i})[D_i(s_i)] = 1$ for all $s_i \in S^0_{i\varepsilon}(\sigma_{-i})$. Hence

$$
\begin{aligned}
\sigma_i^0(\varepsilon)(s_i) &= \frac{\varepsilon}{2(1+\varepsilon)m_i^0}\left[\frac{2+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}{1+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}\right] + \frac{1}{1+\varepsilon}\sigma_i^0(s_i) && \text{if } s_i \in C_i(\sigma_i^0)\\
&= \frac{\varepsilon}{2(1+\varepsilon)m_i^0}\left[\frac{2+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}{1+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}\right] + \frac{\varepsilon}{2(1+\varepsilon)\left[1+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}\right]} && \text{if } s_i \in S^0_{i\varepsilon}(\sigma^0_{-i})\setminus C_i(\sigma_i^0)\\
&= \frac{\varepsilon}{2(1+\varepsilon)m_i^0}\left[\frac{2+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}{1+m^{C_i(\sigma_i^0)}_{\varepsilon;\sigma_{-i}}}\right] && \text{if } s_i \in S_i^0\setminus S^0_{i\varepsilon}(\sigma^0_{-i})\\
&= 0 && \text{if } s_i \in S_i\setminus S_i^0
\end{aligned}
$$

The distance $\Delta(\sigma^0(\varepsilon);\sigma)$ between $\sigma^0(\varepsilon)$ and σ^0 is less than ε, which implies that $\sigma^0_{-i}(\varepsilon) \in B_\varepsilon(\sigma^0_{-i})$. Thus, $\sigma_i^0(\varepsilon)$ is also the expected best reply when the focus is $\sigma^0_{-i}(\varepsilon)$ and $\langle \mu_\varepsilon; \sigma^0(\varepsilon)\rangle$ is an RBE for $\varepsilon < \varepsilon''$. Since $\sigma^0(\varepsilon)$ converges to σ^0 as ε goes to zero, σ^0 is an ultraperfect equilibrium. ■

3.3 Ultraperfect and Proper Equilibria

Recall that σ^0 is a proper Nash equilibrium of Γ if there is a sequence $\{\sigma(\varepsilon)\}_{\varepsilon\to 0}$ that converges to σ^0 and such that $\sigma(\varepsilon)$ is an $\varepsilon-$proper equilibrium of Γ for all values of ε. The strategy profile $\sigma(\varepsilon)$ is an $\varepsilon-$proper equilibrium of Γ if it is completely mixed and satisfies the property that if $u_i(s_i';\sigma_{-i}(\varepsilon)) < u_i(s_i;\sigma_{-i}(\varepsilon))$ then $\sigma_i(\varepsilon)(s_i') \leq \varepsilon\sigma_i(\varepsilon)(s_i)$ for all i and $s_i';s_i \in S_i$. In each game Γ, proper equilibria exist and are perfect, but perfect equilibria need not be proper (see Myerson [6]).

Lemma 2 Not all proper (or perfect) equilibria of games Γ are ultraperfect.

Proof. Consider the game in Table 3: In this game $[(p;1-2p;p);(0.5;0.5)]$ is a proper (hence, perfect) equilibrium for all $p \in (0;0.5)$; however, only $[(0.5;0;0.5);(0.5;0.5)]$ is

	L	R
T	2; 0	0; 1
M	1; 1	1; 1
B	0; 1	2; 0

Table 3: Proper Does Not Imply Ultraperfect

	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 4: Ultraperfect Does Not Imply Weakly Proper

an accumulation point of a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \to 0}$ and thus an ultraperfect equilibrium. ■

In the game in Table 3 M is a best reply to a set of zero Lebesgue measure; 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. Therefore, the expected best reply function of player 1 always puts probability of zero on M.

Completely mixed equilibria satisfy many of the known refinements, e.g., they are strictly perfect and strictly proper (see van Damme [10]). It is interesting to note that although the game in Table 3 has several completely mixed equilibria, none of them are ultraperfect.

The strategy profile σ^0 is a weakly proper equilibrium of Γ if there exists a sequence $\{\sigma(\varepsilon)\}_{\varepsilon \to 0}$ of completely mixed strategies that converges to σ^0 and such that (i) σ^0_i is a best reply against $\sigma_{-i}(\varepsilon)$ for all players i and all values of ε; (ii) if $u_i(s'_i; \sigma^0_{-i}) < u_i(s_i; \sigma^0_{-i})$ then $\sigma_i(\varepsilon)(s'_i) \le \varepsilon\sigma_i(\varepsilon)(s_i)$ for all i and $s'_i, s_i \in S_i$. Proper equilibria are weakly proper and weakly proper equilibria are perfect, but perfect equilibria need not be weakly proper and weakly proper equilibria need not be proper.

Lemma 3 Not all ultraperfect equilibria are weakly proper.

Proof. Consider the game in Table 4: The pure strategy profile $(M; C)$ is perfect

	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 5: Ultraperfect Equilibrium is not Invariant to Adding Dominated Strategies

and ultraperfect, but not weakly proper. ■

The pure strategy equilibrium (T, L) is the only stable set of the game in Table 4. Thus, ultraperfect equilibria need not satisfy Kohlberg and Mertens' [3] definition of strategic stability.

The game in Table 5 is obtained by adding the dominated strategy R to the game in Table 3. By Theorem 3, in this game $[(p, 1-2p, p), (0.5, 0.5, 0)]$ is an ultraperfect 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 ultraperfect equilibria may expand . However, we conjecture that it will never contract. Note that the set of perfect, or proper, equilibria is not invariant either to adding or deleting dominated strategies.

3.4 Essential Equilibria are Ultraperfect

Recall that we have defined the distance between two strategy profiles σ and σ' as $\Delta(\sigma, \sigma') = \max_{\substack{i \in N \\ s_i \in S_i}} |\sigma_i(s_i) - \sigma_i'(s_i)|$. Define the payoff distance between two games $\Gamma = \langle N, \Sigma, (u_i) \rangle$ and $\Gamma' = \langle N, \Sigma, (u_i') \rangle$ as the maximal payoff difference between them: $\Delta(\Gamma, \Gamma') = \max_{\substack{i \in N \\ \sigma \in \Sigma}} |u_i(\sigma) - u_i'(\sigma)|$. 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 $\delta > 0$ such that every game with payoff distance less than δ from Γ has a Nash equilibrium with distance less than ε from σ^0. Van Damme [10] has shown that every essential equilibrium is strictly perfect (and hence perfect); we will now show that essential equilibria are also ultraperfect.

Lemma 4 Every essential equilibrium is ultraperfect.

Proof. Let σ^0 be an essential Nash equilibrium of $\Gamma = \langle N; S; (u_i)\rangle$. We will show that there exists $r' > 0$ such that for all $0 < r < r'$ and for all i, $D_i(s_i) \cap B_r(\sigma^0_{-i})$ has positive Lebesgue measure for all s_i such that $\sigma^0_i(s_i) > 0$. Then, by Theorem 3, σ^0 is ultraperfect. Suppose, to the contrary, that for all $r > 0$ there exist $i^* \in N$, $s^*_{i^*} \in S_{i^*}$ and $r'' < r'$ such that $\sigma^0_{i^*}(s^*_{i^*}) > 0$ and $D_i(s^*_{i^*}) \cap B_{r''}(\sigma^0_{-i^*})$ has zero Lebesgue measure. Then $D_i(s^*_{i^*}) \cap B_\varepsilon(\sigma^0_{-i^*})$ also has zero Lebesgue measure for all $\varepsilon < r''$. Define u^δ_i as follows:

$$
\begin{aligned}
u^\delta_i(s_i; \sigma_{-i}) &= u_i(s_i; \sigma_{-i}) - \delta \quad \text{for} \quad i = i^* \quad \text{and} \quad s_{i^*} = s^*_{i^*} \\
&= u_i(s_i; \sigma_{-i}) \quad \text{otherwise}
\end{aligned}
$$

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

Essential equilibria need not exist, which, together with Theorem 1, implies that not all ultraperfect equilibria are essential. Wu and Jiang [11] 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.

4 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 the Nash equilibrium. That condition, exact and correct beliefs, seems to us much too stringent.

We examine limiting equilibrium behavior as beliefs converge to the measure that puts all mass on the opponents' equilibrium choices, $\varepsilon \to 0$ and find that a limit equilibrium, called an ultraperfect equilibrium, exists, refines perfect equilibrium, and coarsens essential equilibrium. It is neither a refinement nor a coarsening of proper equilibrium. Ultraperfect equilibria utilize only strategies that are best replies to sets of positive measure, implying an appealing robustness of ultraperfection as compared, particularly, with perfect equilibrium.

Appendix A

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 $s_{-i} \in S_{-i}$. Denote by $e_i^k \in \Sigma_i$ the vector whose k-th element equals one and whose other elements are zero and let $e_L^k = (e_i^k)_{i \in L}$.

Lemma 5 Let $\langle \Gamma; \mu_\varepsilon \rangle$ be a random belief game satisfying Definition 1 and parts (i) and (ii) 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 no player has indifferent strategies, then the expected best reply mapping is single-valued, β_{μ_ε} reduces to a single-valued function, $\psi_{\mu_\varepsilon} : \Sigma \to \Sigma$.

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 δ_0 and all $\delta \in (-\delta_0; \delta_0)$ the strategy profile $\sigma_{-1}' = \delta e_{-1}^1 + (1 - \delta)\sigma_{-1} \in D_1(1) \cap D_1(2)$. For $L \subset N \setminus \{1\}$ and $L^C = N \setminus \{1\} \setminus L$ define

$$u_1(e_1^k; e_{L^C}^1; \sigma_L) \equiv \sum_{s_L \in S_L} u_1(e_1^k; e_{L^C}^1; s_L)\sigma_L(s_L) \tag{6}$$

Equation (6) 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. Using equation (6), the payoff to player 1 when she uses pure strategy k, and the remaining players use the mixed profile σ_{-1}' is

$$u_1(e_1^k; \sigma_{-1}') = \sum_{L \subset N \setminus \{1\}} \delta^{n-|L|-1}(1-\delta)^{|L|} u_1(e_1^k; e_{L^C}^1; \sigma_L) \tag{7}$$

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

	1	2
1	2; 0	3; 1
2	1; 1	2; 1
3	1; 1	2; 0

Table 6: Duplicate Strategies

of equation (7), implies

$$u_1(e_1^1; e_{-1}^1) - u_1(e_1^2; e_{-1}^1) = \sum_{\substack{L \subset N \setminus \{1\} \\ L \neq N \setminus \{1\}}} \left(\frac{1-\delta}{\delta} \right)^{|L|} \left[u_1(e_1^2; e_{L^c}^1; \sigma_L) - u_1(e_1^1; e_{L^c}^1; \sigma_L) \right] \tag{8}$$

The left hand side of equation (8) is non-zero by assumption; however, the right hand side cannot equal the left hand side for all $\delta \in (-\delta_0; \delta_0)$. 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 $[(1-\delta)/\delta]^{|L|}$, the sum of terms sharing this coefficient equals zero. Second, if (i) and (ii) do not hold then the right hand side changes in value as δ varies and it must go to infinity in absolute value as $\delta \to 0$. 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 statement in the lemma is not true. In the game in Table 6 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.

References

[1] Bernheim, B. Douglas. (1984). “Rationalizable Strategic Behavior,” Econometrica 52, 1007–1028.

[2] Harsanyi, John C. (1973). “Game with Randomly Disturbed Payo¤s: A New Rationale for Mixed-Strategy Equilibrium Points,” International Journal of Game Theory 2, 1–23.

[3] Kohlberg, Elon and Mertens, Jean-François. (1986). “On the Strategic Stability of Equilibria,” Econometrica 54, 1003–1037.

[4] Kreps, David M. and Wilson, Robert. (1982). “Sequential Equilibrium,” Econometrica 50, 863–894.

[5] McKelvey, Richard D. and Palfrey, Thomas R. (1995). “Quantal Response Equilibria for Normal Form Games,” Games and Economic Behavior 10, 6–38.

[6] Myerson, Roger. (1978). “Re…nements of the Nash Equilibrium Concept,” International Journal of Game Theory 7, 73–80.

[7] Parthasarathy, K. R. (1967). Probability Measures on Metric Spaces. New York: Academic Press.

[8] Pearce, David G. (1984). “Rationalizable Strategic Behavior and the Problem of Perfection,” Econometrica 52, 1029–1050.

[9] Selten, Reinhard. (1975). “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” International Journal of Game Theory 4, 25–55.

[10] van Damme, Eric. (1991). Stability and Perfection of Nash Equilibria, 2nd ed. Berlin: Springer-Verlag.

[11] Wu, Wen-Tsün and Jiang, Jia-He. (1962). "Essential Equilibrium Points n-Person Non-Cooperative Games," Scientia Sinica 11, 1307–1322.