Risk Dominance and Stochastic Potential

Stephen Morris[*]

Rafael Rob[*]

Hyun Song Shin[**]

December 1992

Revised, April 1993

Abstract

In this paper we propose an equilibrium selection mechanism which is based on the existence of higher order uncertainty, and on the way that players make interactive inferences about the game they play, and about the actions their opponents take in equilibrium. We show that the circumstances under which a unique equilibrium is selected depend on the structure of the games being played, and on the extent to which a belief at one state can be extended (through players' interrelated information structures) to beliefs at other states. We delineate these circumstances in a Bayesian framework in which games vary with the state of nature and in which players have more than two actions available to them. The results are illustrated both in the finite and the continuous information cases.

[*]Department of Economics, University of Pennsylvania.

[**]University College, Oxford.

Section I: Introduction

Of late there have been several attempts to go beyond the "refinement" literature and select among multiple strict Nash equilibria. Two basic approaches have emerged. One is the axiomatic approach (see Harsanyi and Selten (1988)) which selects an equilibrium according to a set of desiderata specified by the modeler. Another is the evolutive approach (see Foster and Young (1990), or Kandori, Mailath and Rob (1993)), which considers a 'perturbed' adjustment process, and selects its limit outcome. While these two approaches are quite different, they turn out to select the same (risk dominant equilibrium) in 2x2 games.

In this paper, we suggest a new approach which is based on the existence of higher order uncertainty, i.e. we consider a situation where each player is uncertain about the type of the other player, which may imply uncertainty about either payoffs, or actions in equilibrium, or both. In general, it might be expected that such higher order uncertainty would, if anything, enlarge the set of possible outcomes. However, in certain natural situations, such higher order uncertainty enables us to make a unique prediction of players' actions. The logic behind this is as follows. Assume (for a reason to be elaborated on below) that one player is known to take a certain action at a certain state (which may itself have a very low probability). Then this knowledge might imply a definitive response by the other player and this, in turn, implies how the original player responds to that knowledge, and so on. If this chain of reasoning results in a unique action profile, then this logic has elicited a specific prediction. The goal of this paper is to delineate the conditions under which this is indeed the case, i.e. to show for which class of games and for which class of information structures this reasoning leads unambiguously to a unique prediction.

As it turns out, when this is the case in a single 2x2 game, the uniqu

prediction is again the risk-dominant equilibrium. But, more generally, we sho

how this can be extended to a Bayesian framework in which games vary with th

'state of nature' and to games with more than two actions. To do so we introduc

two new concepts—stochastic potential and p-dominance—which enable us t

characterize when a unique equilibrium is selected. The stochastic potential i

a way of measuring the extent to which knowledge at one state can be extende

(via Bayes rule) to knowledge at other states given how knowledge interact

across players. p-dominance, on the other hand, is a property of the game itsel

and measures the resilience of equilibria as one varies the beliefs of a playe

over what other players may do. These concepts may have wider applicabilit

outside of the present context. In the present context the relationship betwee

them determines when a unique equilibrium is selected.

The logic of our argument is the same as that in Rubinstein's (1989

'electronic mail game' example, and this paper can be seen as delineating exactl

when, in general, that logic operates. Rubinstein's example showed how a sligh

departure from common knowledge alters discontinuously the outcome(s) that woul

have occurred under common knowledge. An ensuing paper by Monderer and Samet

(1989) argues that the discontinuity is an artifact of the way that 'almost

common knowledge' is defined, and that continuity can be restored by a

appropriate definition of 'closeness' on information structures. This paper uses

properties of 'belief operators', introduced by Monderer and Samet, to

characterize the situations under which the electronic mail game type argument

operates. We then demonstrate that this argument operates precisely in

situations where there is no 'almost common knowledge', in Monderer and Samet's

sense, anywhere in the state space. But we also argue that these situations

arise naturally, not only as in Rubinstein's setting, but also for example when

2

players observe some slightly noisy signals about the environment.

One such setting was illustrated by Carlsson and van Damme (1992), who showed that one form of higher order uncertainty, which they label "global uncertainty", leads to the risk dominant outcome being played. Apart from some minor technical differences and the fact that our concepts apply to a broader class of games, the result we report here agrees with theirs for the class of 2×2 games. The two formulations, therefore, illustrate that the sufficient conditions for this type of argument to work arise quite generally.

There are at least two further ideas in the literature to which the present paper relates. First is the reputation approach proposed by Kreps and Wilson (1982), and Milgrom and Roberts (1982). That approach developed the idea that a small degree of uncertainty alters the set of equilibria of games dramatically. In turn, this is related to the idea that a very small probability event may be 'leveraged' to have global implications. The argument in this paper is related, but the 'leveraging' occurs not through time but across overlapping information sets in the players' information partitions. Also, Kreps and Wilson were concerned with sustaining outcomes that are not equilibria to begin with (so their goal is <u>expansive</u>), whereas here we are interested in <u>reducing</u> the set of original equilibria.

The approach here is also related to the idea that players tend to (rationally) use simple (i.e. state independent) strategies although contingent ones are available, and are not more costly to use. Shin and Williamson (1992) showed that this was true in the context of co-ordination games. The present formulation identifies more generally the set of circumstances where this is expected to occur.

The remainder of the paper is organized as follows. In the next section, we provide a 'leading example' which illustrates some of the paper's key ideas.

In section 3, we lay out the general framework. In section 4, we use a 'belief operator' formulation to analyze the properties of the information structure. In section 5, we analyze properties of games which generalize the notion of risk-dominance for two action games to the idea of p-dominance for many action games. An action profile is p-dominant if the action of each player is the unique best response to any conjecture which assigns probability at least p to the corresponding action of the other player. Thus a strict Nash equilibrium is 1-dominant, a dominant action profile is 0-dominant and a risk-dominant equilibrium is 1/2-dominant. In section 6, two classes of "Separable Compatibility Games" are shown to always have a risk-dominant (ie, 1/2-dominant) Nash equilibrium. In section 7, we provide the main result which places joint restrictions on the games and information structures under which a unique outcome emerges. The existence of a globally risk-dominant equilibrium is required to generate a unique prediction, and for the class of separable compatibility games such an equilibrium always exists. In section 8, we provide an example where the information structure is generated by continuous signals with some local noise. In the limit, as noise tends to zero, the risk-dominant equilibrium is guaranteed to be played at every state. Section 9 concludes.

Section 2: Leading Example

Suppose two investors must decide simultaneously whether to invest in a project. Only if both invest is the project profitable. But there is some uncertainty about exactly how profitable the project is. Furthermore, there is uncertainty about what each investor knows about the project's profitability, which affects investors' incentives to invest not only directly (through their payoff-relevant knowledge) but also indirectly through the inferences they make about each others actions. A simple model incorporating these features is used

4

in this section to illustrate key ideas in the paper.

Suppose there is a set of possible states, $\Omega = \{0,1,2,3,\ldots,M,M+1\}$. The prior probability distribution $\{p(k)\}$ over Ω is parametrized by a constant $\alpha > 0$ which satisfies:

$$p(k+1)/p(k) = \alpha.$$

There are two players, 1 and 2, who are potential investors in the project. The success or otherwise of the project depends on the state as well as the actions of the players. At state 0, the project is a certain success. It is a dominant action for both players to invest in the project. For a constant $\beta > 0$, the payoff matrix for the game played at state 0 is:

$$G_0 = \begin{bmatrix} (0,0) & (0,\beta) \\ (\beta,0) & (\beta,\beta) \end{bmatrix}$$

Each player chooses from the set of actions $\{N,I\}$. N stands for "no investment" and I stands for "investment". For player 1, choosing the top row is to choose N, while for player 2, choosing the left column is to choose N. For state $k \in \{1,2,3,\ldots,M\}$, the project succeeds only if both players invest. If only one of the players invests, then the project fails, and the player who has invested loses the sunk cost of 1 in utility terms. If the project succeeds, both players collect a payoff of β^k at state k. The game played at state k has the following payoffs.

$$G_k = \begin{bmatrix} (0,0) & (0,-1) \\ (-1,0) & (\beta^k,\beta^k) \end{bmatrix}$$

Since $\beta > 0$, each of these games has two strict Nash equilibria, (N,N) and (I,I). Finally, at state $M+1$, the project is a certain failure. At this state,

any player who invests loses the sunk cost irrespective of the decision of the other player. N is therefore a dominant action for both players. The payoffs at state $M + 1$ are given as follows.

$$G_{M+1} = \begin{bmatrix} (0,0) & (0,-1) \\ (-1,0) & (-1,-1) \end{bmatrix}$$

We now turn to the information of the players. We shall assume that player 1 has the following information partition over Ω.

$$\{0\}, \{1,2\}, \{3,4\}, \ldots, \{k,k+1\}, \ldots, \{M-1,M\}, \{M+1\}.$$

Player 2's information partition is given by:

$$\{0,1\}, \{2,3\}, \{4,5\}, \ldots, \{k-1,k\}, \ldots, \{M-2,M-1\}, \{M,M+1\}.$$

Thus, at states 0 and $M + 1$, player 1 has strictly better information than player 2, but at all other states, neither player has superior information to the other. Denote by $\mathcal{P}^1$ the partition of player 1, and by $\mathcal{P}^2$ the partition of player 2. The strategy of player 1 in the Bayesian game is a function:

$$s^1 \colon \Omega \to \{N,I\}$$

which is measurable on $\mathcal{P}^1$, while player 2's strategy is a function $s^2 \colon \Omega \to \{N,I\}$ which is measurable on $\mathcal{P}^2$. Players have dominant strategies at the two end points of the state space, but the games at the intermediate states have multiple strict Nash equilibria. These are the games which are of economic interest.
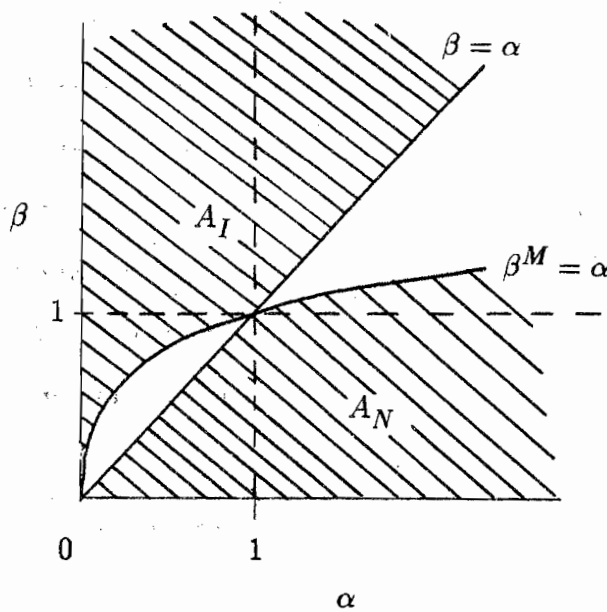
The outcome in the Bayesian game can be seen as the result of a 'tug of war' between the actions at the two end points of Ω. Which action wins out in this tug of war depends crucially on the relative magnitudes of the parameters α and β. We introduce the following sets. Let A_N and A_I be subsets of $\mathbb{R}^2_+$ defined as follows.

6

$$A_N = \{(\alpha,\beta) \,|\, \alpha > max(\beta,\beta^M)\}$$

$$A_I = \{(\alpha,\beta) \,|\, \alpha < min(\beta,\beta^M)\}.$$

These sets are illustrated in figure 1.

[Figure 1]

Let us denote by $\bar{\Omega}$ the set of states which excludes state 0 and $M + 1$. Then we have:

THEOREM 2.1. If $(\alpha,\beta) \in A_I \cup A_N$, there is a unique equilibrium (s^1, s^2) of the Bayesian game. Moreover,

$$(\alpha,\beta) \in A_N \Leftrightarrow s^1(\omega) = s^2(\omega) = N, \quad \forall \; \omega \in \bar{\Omega}. \tag{2.1}$$

$$(\alpha,\beta) \in A_I \Leftrightarrow s^1(\omega) = s^2(\omega) = I, \quad \forall \; \omega \in \bar{\Omega}. \tag{2.2}$$

In other words, in the shaded regions of figure 1, either both players invest at every intermediate state or both players choose not to invest at any of the intermediate states. We prove this theorem by showing that player 2's

7

equilibrium strategy is a constant function whenever $(\alpha, \beta) \in A_I \cup A_N$, and that it has the values stated above.

To begin with, suppose that $(\alpha, \beta) \in A_N$. We will show that $s^2(M+1) = N$, and show that if $s^2(k+1) = N$ then $s^2(k) = N$, for all k. At state $M+1$, player 2's information is given by the set $\{M, M+1\}$. At state $M+1$, player 1 chooses N, since N is the dominant action for 1, and $\{M+1\}$ is an element of 1's partition. Given the information $\{M, M+1\}$, player 2 can guarantee a payoff of zero by playing N. The maximum expected payoff from playing I is given by:

$$p(\{M\}|\{M, M+1\})(\beta^M) + p(\{M+1\}|\{M, M+1\})(-1). \tag{2.3}$$

This is negative if and only if $\alpha > \beta^M$. Given that $(\alpha, \beta) \in A_N$, this inequality holds. Hence, $s^2(M+1) = N$.

Now, suppose $s^2(k+1) = N$. If k and $k+1$ belong to the same element of 2's partition, then the measurability of s^2 on $\mathcal{P}^2$ implies that $s^2(k) = N$. Thus, suppose that k and $k+1$ belong to different elements of $\mathcal{P}^2$. Then, $\{k, k+1\}$ is an element of player 1's partition. Refer to figure 2.
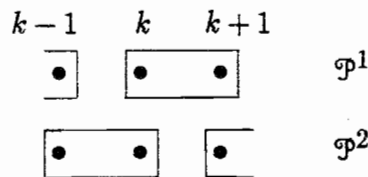
[Figure 2]

Consider 1's optimal action given the information $\{k, k+1\}$. 1 can guarantee a payoff of zero by choosing N. In contrast, the maximum expected payoff from playing I is given by:

8

$$p(\{k\}|\{k,k+1\})(\beta^k) + p(\{k+1\}|\{k,k+1\})(-1) \qquad (2.4)$$

This is negative if and only if $\alpha > \beta^k$. Given that $(\alpha,\beta) \in A_N$, this inequality holds. Hence, $s^1(k+1) = N$.

Now, consider 2's action at state k. At k, 2's information is $\{k-1,k\}$. 2 can guarantee a payoff of zero by playing N. The maximum expected payoff from playing I is:

$$p(\{k-1\}|\{k-1,k\})(\beta^{k-1}) + p(\{k\}|\{k-1,k\})(-1) \qquad (2.5)$$

which is negative if and only if $\alpha > \beta^{k-1}$. Given that $(\alpha,\beta) \in A_N$, this inequality holds. Hence, $s^2(k) = N$.

We have shown that s^2 is a constant function over Ω with value N. Player 1's best reply is to choose s^1 whose value is N at all states other than 0. Thus, we have shown:

$$(\alpha,\beta) \in A_N \Rightarrow s^1(\omega) = s^2(\omega) = N, \quad \forall \, \omega \in \overline{\Omega}$$

which is one half of (2.1). To see that the converse holds, we note from (2.3), (2.4) and (2.5) that if $s^1(\omega) = s^2(\omega) = N$ for all $\omega \in \{1,2,\ldots,M\}$, then $\alpha > \beta^k$ for all $k \in \{1,2,\ldots,M\}$. This implies that $(\alpha,\beta) \in A_N$.

Equation (2.2) can be proved in exactly analogous manner. If $(\alpha,\beta) \in A_I$, then s^2 can be shown to be a constant function, whose value is I everywhere. Again, the argument is inductive. Show first that $s^2(0) = I$, and then show that if $s^2(k) = I$, then $s^2(k+1) = I$. We leave the details to the reader. Notice also that this result is tight in the following sense. If $(\alpha,\beta) \notin A_I \cup A_N$, there exist (pure strategy) Nash Bayesian equilibria so that $s^1(\omega) = I$ and $s^1(\omega') = N$, for some $\omega,\omega' \in \Omega$. ∎

Section 3: The framework and a general possibility result

Theorem 2.1 illustrates the fact that a player's decision depends on what she <u>thinks the other player knows</u>. For instance, at $k = 1$ player 2 thinks that

player 1 may know $k = 0$ which would induce player 1 to invest, which makes it profitable for player 2 to invest as well (regardless of what player 1 might do if he knew $k = 1$). The interesting fact is that this type of 'local reasoning' at $k = 1$ has repercussions for actions of both players at other states (actually at most states). We shall refer to this form of reasoning as the 'infection argument'. In the example, we saw how this *infection* argument can generate surprisingly precise predictions about the play of a game with incomplete information where almost all of the complete information games have multiple Nash equilibria. But we also saw how changing the prior beliefs on the state space could reverse this prediction or might leave us with multiple equilibria. Thus when the prior probability of states declined sufficiently fast, the unique Bayesian Nash equilibrium of the incomplete information game had investment occurring almost everywhere, while if the prior probability of states increased sufficiently fast, the unique Bayesian Nash equilibrium of the incomplete information game had investment *not* occurring almost everywhere. If the prior probability varies at an intermediate rate, there exist multiple equilibria. Thus while the infection argument is powerful, the prediction implied by the infection argument is sensitive to the prior on states and the structure of information.

The main purpose of this paper is to show how natural properties of the information system do sometimes uniquely determine the outcome. But in this section, we want to give a formal statement of the sense in which "anything can happen", if we are allowed to manipulate the prior beliefs, for any given structure of payoffs and information partitions. This also helps make clear the logic of the infection argument itself. We will also introduce the general framework of multi-action two player games of incomplete information which is

10

used throughout the paper.

Private information is to be represented by *partitions* of the set of states of the world (as in Aumann (1976)), rather than by a cross product of sets of possible types of each player, as in the standard framework introduced by Harsanyi (1967/68). This modelling choice does not change any results. But discussion of iterated knowledge and belief is more naturally carried out in a partition framework.

An *information system* consists of $I = [\Omega, \{1,2\}, \{\mathcal{P}^i\}_{i=1,2}, \pi]$, where:-

Ω is a finite set of states of the world;

$\{1,2\}$ is the collection of players;

$\mathcal{P}^i$ is the partition of states of the world representing the information of player i;

π is a strictly positive prior probability distribution on Ω.

Each player is assumed to share the same prior π on Ω. We will write ω for a typical element of Ω. Then $\pi(\omega)$ is the probability of state ω. We will also write $P^i(\omega)$ for the element of i's partition, $\mathcal{P}^i$, containing state ω. Thus if the true state is ω, $P^i(\omega)$ is the set of states which player i thinks possible.

We will require a notion of *common knowledge* in this information system. An *event* is a subset of the state space Ω. An event E is said to be common knowledge at state ω if there exists an event F such that $\omega \in F \subset E$ and $P^i(\omega') \subset F^1$ for each player and all $\omega' \in F$. This notion of common knowledge can be shown to be equivalent to the intuitive iterated notion (see Aumann (1976)). We

---

[1]Throughout the paper, the symbol $\subset$ is used to represent weak *or* strict inclusion.

will be concerned with situations where there is no non-trivial common knowledge in the information system — that is, where the only event that is ever common knowledge is the universal set Ω. It can be shown that there is no non-trivial common knowledge if and only if the *meet* (finest common coarsening) of the two players' partitions is the trivial partition which consists only of the universal set Ω.

An *incomplete information game* consists of $G = [I, \{A^i\}_{i=1,2}, \{g^i\}_{i=1,2}]$, where:—

> I is an information system as described above;
>
> A^i is the finite set of actions available to player i;
>
> $g^i\colon A \times \Omega \to \mathbf{R}$ is player i's payoff function, with $A = A^1 \times A^2$.

Thus, in each state ω, there is a one shot (complete information) game associated with ω, with payoff to player i of $g^i(a^1, a^j; \omega)$ if player 1 chooses action $a^1 \in A^1$, and player 2 chooses action $a^2 \in A^2$. But in general players do not know the true state, so must choose strategies based only on their information. Thus a pure strategy for player i in the incomplete information game is a function $s^i\colon \Omega \to A^i$, measurable with respect to his partition. Write S^i for the set of such pure strategies. Now a pure strategy pair (s^1, s^2) is a *Bayesian Nash* equilibrium of the incomplete information game if, for $(i,j) = (1,2)$ and $(2,1)$ and for all $t^i \in S^i$,

$$\sum_{\omega \in \Omega} \pi(\omega) g^i[s^i(\omega), s^j(\omega); \omega] \geq \sum_{\omega \in \Omega} \pi(\omega) g^i[t^i(\omega), s^j(\omega); \omega].$$

Mixed strategy equilibria could be defined in the natural way. Results in

12

this pape
informati
games bei
action fo

Act
for (i,j)

In
which are
above cond

We a

THEOREM 3
informatio
equilibriu
state ω^*,
non-trivia
state spac
equilibriu

In o
adjust the
of every co

this paper will be characterizing the Bayesian Nash equilibria of incomplete information games in terms of the properties of the various complete information games being played at different states of the world. Thus, a^i is a *dominant* action for player i in the state ω game if

$$g^i(a^i, a^j; \omega) > g^i(b^i, a^j; \omega), \text{ for all } a^j \in A^j, b^i \in A^i. \quad (3.1)$$

Action pair (a^1, a^2) is a *strict Nash equilibrium* of the state ω game if, for $(i,j) = (1,2)$ and $(2,1)$,

$$g^i(a^i, a^j; \omega) > g^i(b^i, a^j; \omega), \text{ for all } b^i \in A^i. \quad (3.2)$$

In particular, we will often be interested in action pairs $a = (a^1, a^2)$ which are strict Nash equilibria in _every_ complete information game, so that the above condition holds for _every_ ω.

We are now in a position to state the main result of this section.

THEOREM 3.1. *Suppose that, for some information system and incomplete information game, [1] some action pair (a^1, a^2) is a strict pure strategy Nash equilibrium of the state ω game, for each ω; [2] for some player i, and for some state ω^*, action a^i is a dominant action in the state ω^* game; [3] there is no non-trivial common knowledge. Then for some strictly positive prior π on the state space Ω, playing (a^1, a^2) everywhere is the unique Bayesian Nash equilibrium of the incomplete information game.*

In other words, leaving everything other than the prior the same, we can adjust the prior to guarantee that any strategy pair which is a Nash equilibrium of every complete information game is always played in the unique Bayesian Nash

equilibrium of the incomplete information game.

PROOF. [1] Since (a^1, a^2) is a strict pure strategy equilibrium of every ω complete information game, there exists some p, $0 < p < 1$, such that for every state ω information game, and for all $b^1 \in A^1$, $b^j \in A^j$,

$$pg^i(a^i, a^j; \omega) + (1-p)g^i(a^i, b^j; \omega) > pg^i(b^i, a^j; \omega) + (1-p)g^i(b^i, b^j; \omega).$$

[2] Suppose, without loss of generality, that 1 is the player with a strictly dominant action on state ω^*. Then there exists a q, $0 < q < 1$, such that, for all $b^1 \in A^1$, $s^2 : \Omega \to A^2$, $\omega \neq \omega^*$,

$$qg^1(a^1, s^2(\omega^*); \omega^*) + (1-q)g^1(a^1, s^2(\omega); \omega) >$$
$$qg^1(b^1, s^2(\omega^*); \omega^*) + (1-q)g^1(b^1, s^2(\omega); \omega).$$

[3] Let N be the number of elements in Ω and choose an ε such that

$$0 < \varepsilon < \frac{1}{N}\min\left\{\frac{1-q}{q}, \frac{1-p}{p}\right\}.$$

[4] Let $L^1E = \{ \omega \in \Omega \mid P^1(\omega) \cap E \neq \emptyset \}$. Thus L^1E is the set of states where i allows that event E is possible. Let $A_0 = \{\omega^*\}$, $A_1 = P^1(\omega^*)$, define A_t inductively by $A_{t+1} = L^2A_t$ if t is odd, $A_{t+1} = L^1A_t$ if t is even. Since each L^1 is an increasing function ($L^1E \supset E$) and Ω is finite, $\{A_t\}$ is an increasing sequence which converges to some set A^*. But then we would have $L^iA^* = A^*$ for each i. Thus by non-trivial common knowledge, there is some integer T such that $A_t = \Omega$ for all $t \geq T$.

[5] Now we construct a prior using the sequence $\{A_t\}$. Let $\phi(\omega)$ be the minimum integer t such that $\omega \in A_t$. Then let

$$\pi(\omega) = \frac{e^{\phi(\omega)}}{\sum\limits_{\omega' \in \Omega} e^{\phi(\omega')}}.$$

14

Thus any

[6] We ar

strategy

$= \{ \omega \in \Omega \mid$

that $A_1 \subset$

informati

player 2

$A_t \subset \Omega^2$ i

event A_t.

Section

Wha

theorem

of each

probabili

response

states.

following

is some s

of player

1 assigns

at least

to be pla

entire st

we introd

stochasti

Thus any state not in A_t is at most ε times as likely as any state in A_t.

[6] We are now ready to make the infection argument. Suppose (s^1, s^2) is a pure strategy Bayesian Nash equilibrium of the incomplete information game. Let $\Omega^i = \{ \omega \in \Omega \mid s^i(\omega) = a^i \}$, for each i. By [2], [3] and the definition of A_1, we know that $A_1 \subset \Omega^1$, since 1 must assign probability at least q to state ω^*, given his information set $P^1(\omega^*) = A_1$. But for any t odd, $A_t \subset \Omega^1$ implies $A_{t+1} \subset \Omega^2$, since player 2 must assign probability at least p to event A_t. Similarly, t even and $A_t \subset \Omega^2$ implies $A_{t+1} \subset \Omega^1$, since player 1 must assign probability at least p to event A_t. By induction, $\Omega = \Omega^1 = \Omega^2$, and we are done. ∎

Section 4: Stochastic Potential

What *drives* the infection argument used in the example of section 2 and theorem 3.1? Suppose that (a^1, a^2) is a strict pure strategy Nash equilibrium of each complete information game. We need to have some lower bound p on the probability assigned by player j to player i playing a^1, to ensure that j's best response is a^j, regardless of what other actions are taken at other possible states. Once we have that lower bound, the trick is an iterative argument of the following form. Suppose player 1 chooses a^1 on some information set. Then there is some set where player 2 assigns probability at least p to that information set of player 1. Repeating the same argument, there is also some set where player 1 assigns probability at least p to the event where player 2 assigns probability at least p to the original information set of player 1 where he was hypothesized to be playing a^1. If we continue this iterative process, and converge to the entire state space Ω, then the infection argument has worked. In this section, we introduce the idea of the *stochastic potential* of an information system. The stochastic potential is the highest probability p such that starting with *any*

15

information set of either player, this iterative process (whereby players believe with probability p that the other player's information lies in some set) converges to the universal set.

In order to define the stochastic potential, it is useful to introduce the idea of belief operators on the state space. Such belief operators were introduced in a related context by Monderer and Samet (1989). Write (with some abuse of notation) the conditional probability of event E, given event F, as:-

$$\pi(E \mid F) = \frac{\pi(E \cap F)}{\pi(F)} = \frac{\sum_{\omega \in E \cap F} \pi(\omega)}{\sum_{\omega \in F} \pi(\omega)}$$

Now for a given information system, define player i's p–belief operator by:

$$B_p^i E = \{\omega \in \Omega \mid \pi(E \mid P^i(\omega)) \geq p\}$$

Thus $B_p^1 E$ is the set of states where player i believes, with probability at least p, that event E will occur. It is useful to compare this belief operator with natural knowledge and possibility operators. Define knowledge operator K^i and possibility operator L^i by

$$K^i E = \{\omega \in \Omega \mid P^i(\omega) \subset E\}$$

$$L^i E = \{\omega \in \Omega \mid P^i(\omega) \cap E \neq \emptyset\}$$

An event is known if that event is true at every state thought possible. An event is possible if the event contains some state thought possible. In our framework, the fact that π is positive everywhere implies that to believe with probability one is equivalent to knowledge (i.e. $K^i E = B_1^i E$ for all E), and to believe with arbitrarily small probability is equivalent to possibility (i.e. $L^i E$

$= B_\epsilon^i E$, for a

Now we

DEFINITION.

number p suc

Thus f

following fo

with probabi

that 1 think

believes wit

$p \ldots$ that 2

To se

potential, s

that $p > q$ in

p, it is bel

Now, by indu

$[B_q^1 B_q^2]^k L^1 E \cap$

increasing s

$\pi(E \mid P^1(\omega)) \geq$

holds.

It is

Consider the

16

$- B_\epsilon^1 E$, for all E, for some $\epsilon > 0$).

Now we can define the key belief property of an information system.

DEFINITION. The *stochastic potential* of an information system is the largest number p such that for every non-empty event E

$$[B_p^1 B_p^2]^k L^1 E \cap [B_p^2 B_p^1]^k L^2 E = \Omega, \text{ for some } k.$$

Thus for any non-empty event E, at any state ω, a statement of the following form is true: "1 believes with probability at least p that 2 believes with probability at least p that 1 believes with probability at least p..... that 1 thinks E is possible; *and* 2 believes with probability at least p that 1 believes with probability at least p that 2 believes with probability at least p... that 2 thinks E is possible".

To see that every information system has a well defined stochastic potential, so that there exists a "largest p" in the definition, notice first that $p > q$ implies $B_p^1 E \subset B_q^1 E$. If an event is believed with probability at least p, it is believed with probability at least q, if p is strictly greater than q. Now, by induction, we know that $p > q$ implies that $[B_p^1 B_p^2]^k L^1 E \cap [B_p^2 B_p^1]^k L^2 E \subset [B_q^1 B_q^2]^k L^1 E \cap [E_q^2 B_q^1]^k L^2 E$. But it is also the case that $B_p^1 E = \lim_{t \to \infty} B_{p(t)}^1 E$ for any increasing sequence, $p(t) \to p$, since $\pi(E | P^1(\omega)) \geq p(t)$ for all t implies $\pi(E | P^1(\omega)) \geq p$. So there must exist a largest p such that the above property holds.

It is useful to illustrate this requirement by means of an example. Consider the following information system:-

$$\Omega = \{1,2,3,4,5,6,7,8,9\}$$
$$\mathcal{P}^1 = (\{1,2,3\}, \{4,5,6\}, \{7,8,9\})$$
$$\mathcal{P}^2 = (\{1,4,7\}, \{2,5,8\}, \{3,6,9\})$$
$$\pi = \left\{\tfrac{1}{21}, \tfrac{1}{7}, \tfrac{1}{7}, \tfrac{1}{7}, \tfrac{1}{21}, \tfrac{1}{7}, \tfrac{1}{7}, \tfrac{1}{7}, \tfrac{1}{21}\right\}$$

Consider event $E = \{1\}$. Then $L^1E = \{1, 2, 3\}$ and $L^2E = \{1, 4, 7\}$. Suppose $p > 3/7$; then $B^1_p L^2E = B^2_p L^1E = \emptyset$. Thus the stochastic potential can be no more than 3/7. Suppose $1/7 < p \leq 3/7$. Then $B^1_p L^2E = \{4, 5, 6, 7, 8, 9\}$, $B^2_p L^1E = \{2, 3, 5, 6, 8, 9\}$ and $B^2_p B^1_p L^2E = B^1_p B^2_p L^1E = \Omega$. By the symmetry of the example, an analogous argument applies for any other singleton event. So the stochastic potential of this information system is 3/7. The idea of stochastic potential is closely related to Monderer and Samet's notion of common p–belief. An event is said to be common p–belief if and only if everyone believes with probability at least p that everyone believes ... etc.. that E is true. It can be shown that if p is greater than 1 minus the stochastic potential, then no event other than Ω is ever common p–belief.

We would like to characterize the possible stochastic potentials of information systems. The following property of the belief operator B^1_p when $p > 1/2$ will be useful. We write $\sim E$ for the complement of event E in Ω.

LEMMA 4.1. Suppose $p > 1/2$. Then $[B^j_p B^1_p]^k \sim E \subset \sim [B^j_p B^1_p]^k E$, for all events E, $k \geq 1$.

PROOF. $\omega \in B^1_p F \cap B^1_p \sim F$ implies $\pi[F|P^1(\omega)] > 1/2$ and $\pi[\sim F|P^1(\omega)] = 1 - \pi[F|P^1(\omega)] > 1/2$, which gives a contradiction. So $B^1_p F \cap B^1_p \sim F = \emptyset$, for all events F. Therefore $B^1_p \sim F \subset \sim B^1_p F$ for all players and events. Iterated application of this gives the result. ∎

LEMMA 4.2. [i] The stochastic potential of any information system is less than

18

or equal to 1/2 if and only if there is some private information [i.e. $P^i(\omega) \neq \Omega$, for some i, $\omega \in \Omega$]. [ii] The stochastic potential of an information system is strictly greater than zero if and only if there is no non-trivial common knowledge.

PROOF. [i] Suppose $p > 1/2$. If player i has private information, then there is an event E such that $E = L^i E$ and $\sim E = \sim L^i E$. But now by lemma 4.1, $[B^i_p B^j_p]^k L^i \sim E = [B^i_p B^j_p]^k \sim E \subset \sim [B^i_p B^j_p]^k E = \sim [B^i_p B^j_p]^k L^i E$. So one of $[B^i_p B^j_p]^k L^i \sim E$ and $[B^i_p B^j_p]^k L^i E$ is not equal to Ω, for all k, so p is not the stochastic potential. If all players have no private information, then $L^1 E = L^2 E = \Omega$, for all non-empty E, so stochastic potential is 1. [ii] Suppose there is non-trivial common knowledge; then there exists an event $F \neq \Omega$ such that $P^i(\omega) \subset F$ for all $\omega \in F$. Now if $p > 0$, $[B^1_p B^2_p]^k L^1 F = [B^2_p B^1_p]^k L^2 F = F \neq \Omega$ for all k, so p is not the stochastic potential. If there is *no* non-trivial common knowledge, let $p = \min_{\omega \in \Omega} \pi(\omega)$. Now $B^i_p E = L^i E$, for all events E. But $[L^i L^j]^k E = \Omega$ for some k by no non-trivial common knowledge. ∎

We can check that 1/2 is not only an upper bound on the stochastic potential, but is also attained for some information systems. For example, the information system in the leading example of section 2 has a stochastic potential of $min \left\{ \dfrac{1}{1+\alpha}, \dfrac{\alpha}{1+\alpha} \right\}$; so that if there is a uniform prior (ie, $\alpha = 1$), the stochastic potential is 1/2.

Information systems with the structure of overlapping information sets of the leading example arise naturally in many classes of problems where there is a lack of common knowledge. Examples include the co-ordinated attack problem of the computer science literature [Halpern (1986)], the electronic mail game of Rubinstein (1989) and the hidden envelopes trading problem of Geanakoplos (1992).

In section 8, we show that if uncertainty is generated by some symmetric

19

and local noise in an infinite state space, the stochastic potential tends to 1/2 as the noise tends to zero.

Section 5: p-Dominance

The infection argument depends on the relation between the stochastic potential, introduced in the last section, and properties of the payoffs of the games being played. This section introduces those critical properties. In order to ensure that a pair of actions (a^1, a^2) is going to infect the state space, we require a probability p such that a^1 is a best response for player i to a^j, just as long as player i assigns at least probability p to j choosing action a^j (regardless of what probability he assigns to other possible actions).

DEFINITION. Action pair (a^1, a^2) is _p-dominant_ in the state ω game if, for (i,j) = (1,2) and (2,1), and for all $b^1 \in A^1$ and $b^2 \in A^2$,

$$pg^i(a^i, a^j; \omega) + (1-p) g^i(a^i, b^j; \omega) > pg^i(b^i, a^j; \omega) + (1-p) g^i(b^i, b^j; \omega). \quad (5.1)$$

This definition is equivalent to requiring (a^1, a^2) to be p-dominant in every two-by-two game generated by restricting the multi-action game to actions sets $B^1 = \{a^1, b^1\}$, $B^2 = \{a^2, b^2\}$ for some $b^1 \in A^1$, $b^2 \in A^2$.

Notice that if (a^1, a^2) is p-dominant for some p, then (a^1, a^2) is q-dominant for any $p \leq q \leq 1$. This suggests the following definition of the critical level at which a given action pair becomes p-dominant.

DEFINITION. If action pair (a^1, a^2) is _p-dominant_ for some $0 < p \leq 1$ in the state ω game, we say that the _critical dominance level_ of (a^1, a^2) is the infinum of the set of p's for which (a^1, a^2) is p-dominant.

Notice that since p-dominance is defined by means of a strict inequality, the action pair (a^1, a^2) is p-dominant if and only if p is strictly greater than the critical dominance level of (a^1, a^2).

We also note in passing that the notion of p-dominance unifies a number of standard concepts: (a^1, a^2) is a strict Nash equilibrium if (a^1, a^2) is 1-dominant and actions a^1 and a^2 are dominant actions if (a^1, a^2) is 0-dominant (these concepts are defined on p.13 in equations (3.1) and (3.2)).

Finally, notice that the notion of p-dominance introduced here is closely related to Harsanyi and Selten's (1988) notion of risk dominance. To see why, observe that action pair (a^1, a^2) is *1/2-dominant* in the state ω complete information game if, for $(i,j) = (1,2)$ and $(2,1)$, and for all $b^1 \in A^1$ and $b^2 \in A^2$,

$$\tfrac{1}{2} g^i(a^i, a^j; \omega) + \tfrac{1}{2} g^i(a^i, b^j; \omega) > \tfrac{1}{2} g^i(b^i, a^j; \omega) + \tfrac{1}{2} g^i(b^i, b^j; \omega).$$

This exactly coincides with Harsanyi and Selten's risk-dominance in 2 by 2 games. For many action games, the notion of 1/2-dominance is more stringent, as it makes comparisons between *every* action pair, not just between Nash equilibrium action pairs.

The key step in the infection argument uses an alternative characterization of p-dominance which is stated in the next lemma.

LEMMA 5.1. Action pair (a^1, a^2) is p-dominant if and only if for $(i,j) = (1,2)$ and $(2,1)$, for every probability distribution λ on A^j such that $\lambda(a^j) \geq p$, and every $b^i \in A^i$,

$$\sum_{b^j \in A^j} \lambda(b^j) g^i(a^i, b^j; \omega) > \sum_{b^j \in A^j} \lambda(b^j) g^i(b^i, a^j; \omega). \qquad (5.2)$$

This lemma tells us that if (a^1, a^2) is p–dominant, a^i is the *unique* best
[re]sponse for player i as long as he believes that the other player will play a^j
[wi]th probability at least p. If $p = 0$, the lemma states that an action is a best
[re]sponse to any conjecture over the other player's actions.

[PR]OOF   Notice first that if $p=1$, (5.1) and (5.2) are trivially equivalent. So
[we] restrict ourselves to $p < 1$. Now suppose (5.2) is true. Then in particular
[(5.]2) is true for $\lambda(a^j) = p$, $\lambda(b^j) = 1-p$, for any $b^j \in A^j$. This gives (5.1).
[Con]versely, suppose (5.1) is true. For any probability distribution λ on A^j with
$[\lambda(a^j) \geq p$, (5.2) implies:-

$$[1] \quad p(1-\lambda(a^j)) g_i(a^i, a^j; \omega) + (1-p) \sum_{b^j \ne a^j} \lambda(b^j) g^i(a^i, b^j; \omega) >$$
$$p(1-\lambda(a^j)) g_i(b^i, a^j; \omega) + (1-p) \sum_{b^j \ne a^j} \lambda(b^j) g^i(b^i, b^j; \omega)$$
$$[2] \quad (\lambda(a^j)-p) g^i(a^i, a^j; \omega) \geq (\lambda(a^j)-p) g^i(b^i, a^j; \omega)$$

[Sum]ming these constraints gives

$$(1-p) \sum_{b^j \in A^j} \lambda(b^j) g^i(a^i, b^j; \omega) > (1-p) \sum_{b^j \in A^j} \lambda(b^j) g^i(b^i, b^j; \omega),$$

[and] thus (5.2).   ∎

An example will illustrate the character of p–dominance and the relation
[to] risk–dominance. Consider the following symmetric game.

22

player 2

		L	C	R
	T	7,7	0,0	0,0
player 1	M	0,0	2,2	7,0
	D	0,0	0,7	8,8

This game has three strict Nash equilibria: (T,L), (M,C) and (D,R). Risk dominance in Harsanyi and Selten (1988) involves pairwise comparisons of these strict Nash equilibria: (T,L) risk dominates (M,C), (M,C) risk dominates (D,R) and (D,R) risk dominates (T,L). The notion of p-dominance we introduced is global, in the sense that each strict Nash equilibrium action pair is being compared with every other pair of actions. In our terminology, the critical dominance level of (T,L) is 8/15: thus (T,L) is p-dominant for every p strictly greater than 8/15; (D,C) is p-dominant for any p greater than 7/9; (D,R) is p-dominant for any p greater than 2/3. This example illustrates two properties of p-dominance proved in theorem 5.1 below: if there is more than one p-dominant action pair, then p is more than 1/2; and, generically for games with pure strategy equilibria, there is some p for which there is a unique action pair which is p-dominant. In the example, for any $8/15 < p \leq 2/3$, (T,L) is the unique p-dominant action pair. For $p \leq 1/2$ the relationship is cyclical, and we cannot single out an action pair which p dominates every other action pair. On the other hand, for certain classes of games it is possible to find a unique p-dominant action pair for $p \leq 1/2$. Such games will be shown in the next section.

THEOREM 5.1. [i] If $p \leq 1/2$, there is at most one p–dominant action pair.

[ii] Generically in payoffs, every game has *either* only strictly mixed Nash equilibria *or* exactly one p–dominant Nash equilibrium action pair (a^1, a^2), for some p.

PROOF [i] Suppose $p \leq 1/2$, (a^1, a^2) and (b^1, b^2) are p–dominant and $a^i \neq b^i$. Then we have the contradictory implications:–

$$\tfrac{1}{2}g^i(a^i, a^j; \omega) + \tfrac{1}{2}g^i(a^i, b^j; \omega) > \tfrac{1}{2}g^i(b^i, a^j; \omega) + \tfrac{1}{2}g^i(b^i, b^j; \omega),$$

$$\tfrac{1}{2}g^i(b^i, a^j; \omega) + \tfrac{1}{2}g^i(b^i, b^j; \omega) > \tfrac{1}{2}g^i(a^i, a^j; \omega) + \tfrac{1}{2}g^i(a^i, b^j; \omega).$$

[ii] If there is a pure strategy Nash equilibrium, then generically there is at least one strict Nash equilibrium. Generically, the critical p–dominance of every strict Nash equilibrium will be different. For example, if (a^1, a^2) is a strict Nash equilibrium, adding ε to each player's payoff from action pair (a^1, a^2) will always decrease the critical dominance level of (a^1, a^2) and will weakly increase the critical dominance level of every other strict Nash equilibrium pair. ∎

Section 6: Separable Compatibility Games

Focusing on games with pure–strategy Nash equilibria, theorem 5.1 shows that we can find a small enough p for which only one Nash equilibrium is p–dominant. On the other hand, the example before theorem 5.1 shows that this p may exceed 1/2. In this section we provide examples of games with multiple equilibria, but where one of the equilibria is p–dominant for $p < 1/2$. In the next section we show that such examples are key to the nonvacuousness of the

24

infection argument in general multi-action games.

Consider then the following 'partnership game'. There are two players who play a symmetric game in which they both pick 'effort levels'. The payoff to each player increases in his own and his partner's effort, but decreases the further apart these effort levels are (perhaps, because the two players have to spend some time rectifying the difference between their resulting abilities). Another scenario which fits this description is where each player chooses a technology which may not be fully compatible with the technology chosen by her opponent. In such circumstances players are penalized the further apart their technologies are. The payoff function corresponding to such scenarios can be written as

$$g^i(a^i, a^j) = u(a^i, a^j) - d(a^i, a^j),$$

where a^i, $a^j \in A$, a common action set, u is a (symmetric) gross payoff and $d(\cdot, \cdot)$ is a distance function[2] which reflects the cost of incompatibility. As a special case of this we might have $A = \mathbf{R}_+$, $u(a^i, a^j) = \lambda a^i + (1-\lambda)a^j - p^i$, for some $\lambda \in (0,1)$ and some $p^i > 0$ (representing the cost of action i) and $d(a^i, a^j) = d |a^i - a^j|$, where $|\cdot|$ represents "absolute value", and d is a positive constant. For d sufficiently large all diagonal elements, (a^i, a^i) for all $a^i \in A$, are Nash equilibria.

Assume now that there exists a dominant action —say a^*— for the gross payoff u (this action need not be dominant for g). Then this action (played by both players) 1/2 dominates every other action pair in the g game. To see that note that

---

[2]$d(\cdot, \cdot)$ satisfies $d(a,a') = d(a',a)$, $d(a,a) = 0$ and $d(a,a'') \leq d(a,a') + d(a',a'')$.

$$g(a^*,a^*) + g(a^*,a') > g(a,a^*) + g(a,a')$$

and only if

$$u(a^*,a^*) - d(a^*,a^*) + u(a^*,a') - d(a^*,a')$$

$$> u(a,a^*) - d(a,a^*) + u(a,a') - d(a,a'),$$

ch holds because a^* is dominant in u (so $u(a^*,a^*) > u(a,a^*)$ and $u(a^*,a') >$,$a')$), and because of the triangle inequality as applied to d. In the ticular example the dominant action is the one for which $\lambda a^i - p^i$ is maximal r $a^i \in A$.

This example exploits the symmetry of the distance function d, selecting equilibrium solely on the basis of asymmetry in the gross payoff function, u , on the basis of the tradeoff between effort and cost alone). Analogously, can create examples with a symmetric u function, and select an equilibrium on basis of asymmetries in the distance function (ie, on the basis of ompatibility considerations alone). Consider the case where $A = \mathbb{R}_+$, $u(a, a')$ (a constant), and $d\ (a,\ a') = c_{a'}|a - a'|$.

Then if we let a^* be the effort level for which c_a is maximal (ie, this is effort level from which it is most costly to deviate), we have:

$$g(a^*,\ a^*) + g(a^*,\ a') > g(a,\ a^*) + g(a,\ a')$$

and only if

$$c_{a'}|a' - a^*| < c_{a'}|a^* - a| + c_{a'}|a - a'|,$$

ch holds because $c_a^* > c_{a'}$ and because of the triangle inequality.

Finally note that these examples are robust in that slight perturbations their payoff functions (g) leave us with games in which one equilibrium is p-inant with $p < 1/2$.

26

Section 7:

We are

incomplete i

characteriza

infection ar

where some

stochastic p

some state w

case, we kno

infection ar

On the

a strongly do

makes a *conc*

equilibrium v

everywhere.

'reversible',

Therefore, st

stuck' somewh

THEOREM 7.1.

potential of

i at some sta

playing (a^1,

incomplete in

PROOF. [i]

equilibrium o

Section 7: Stochastic Potential, p-Dominance and Risk Dominance

We are now in a position to state the main theorem of the paper for finite incomplete information games. There are a number of different ways in which our characterizations of p—dominance and stochastic potential can be used, via the infection argument, to pin down outcomes. Theorem 7.1 considers a situation where some action pair (a^1, a^2) is p—dominant everywhere, where p is the stochastic potential of the information system, and where, in addition, there is some state where a^1 is known to be strongly dominant for some player. In this case, we know that action a^1 will be taken somewhere on the state space, so the infection argument is allowed to operate.

On the other hand, the corollary which follows makes no assumption about a strongly dominant action somewhere on the state space. Instead, the corollary makes a *conditional* claim. Suppose there is a pure strategy Bayesian Nash equilibrium where the p—dominant action pair are played. Then it must be played everywhere. It is important to note, though, that this statement is not 'reversible', ie, it can be made only about <u>one</u> Bayesian—Nash equilibrium. Therefore, starting from any other equilibrium the infection argument may 'get stuck' somewhere.

THEOREM 7.1.   [i] *Suppose that [1] the information system has stochastic potential of p; [2] (a^1, a^2) is p-dominant at every state; and [3] some player i at some state ω^* knows the event "a^1 is a strongly dominant action".   Then playing (a^1, a^2) everywhere is the unique Bayesian Nash equilibrium of the incomplete information game.*

PROOF.   [i] Suppose (σ^1, σ^2) is a (possibly mixed strategy) Bayesian—Nash equilibrium of the game, where $\sigma^1: \Omega \rightarrow \Delta(A^1)$ and σ^1 measurable with respect to

Let $\Omega^i = \{ \omega \in \Omega \mid \sigma^i(\omega)(a^i) = 1 \}$. Thus Ω^i is the set of states where i plays

the pure strategy a^i in the given equilibrium. Now by [2], $B^i_p \Omega^j \subset \Omega^i$. Let $E =$

(ω^*), i.e the event where a^i is strictly dominant everywhere for some player

By construction, $L^i E = E$. Clearly, $E \subset \Omega^i$. Therefore $B^j_p E \subset B^j_p \Omega^i \subset \Omega^j$. Thus

induction $[B^i_p B^j_p]^k E \subset \Omega^i$ and $B^j_p [B^i_p B^j_p]^k E \subset \Omega^j$ for all k. But by [1] $[B^i_p B^j_p]^k E$

$B^j_p [B^i_p B^j_p]^k E = \Omega$ for some k. So $\Omega^1 = \Omega^2 = \Omega$. ∎

COROLLARY 1. *Suppose that [1] the information system has stochastic potential*

p; [2] (a^1, a^2) is p-dominant at every state. Suppose that s is a pure

strategy Bayesian Nash equilibrium of the incomplete information game, with $s^i(\omega)$

a^i for some i and some ω. Then playing (a^1, a^2) everywhere is the unique

yesian Nash equilibrium of the incomplete information game.

The following observation may help clarify the difference between theorem

1 and corollary 1. Observe that the (Bayesian) game described in the theorem

in fact dominance solvable; that is, iterated deletion of strategies which are

rongly dominated in the incomplete information game leaves playing the

ecified action pair (a^1, a^2) everywhere as the unique iteratively undominated

nd thus rationalizable) strategy pair. In the game described in the corollary,

erated deletion will typically have no bite. Thus it could be that there are

o action pairs which are strict pure strategy Nash equilibria at every state,

which case we know that there are two simple Bayesian Nash equilibria where

ch of these action pairs is played everywhere. Nonetheless, if one of these

tion pairs is p–dominant for $p < 1/2$, the conditional statement of the

rollary can be made in reference to it, but not in reference to the other

uilibrium. As a special case of this consider the classes of games in section

28

6. Then we h

for $p < 1/2$.

(and no matt

played everyw

More precisel

COROLLARY 2:

state, and tl

$s^i(\omega) = a^i$ fo

Notice

implies that

strategies a

player takes

if one of ri

dominant act

there are mo

Section 8:

Theore

above zero.

maximum poss

is chosen.

If the

1), then the

well known s

rather speci

6. Then we have shown that there exists a unique action pair which is p-dominant for $p < 1/2$. Therefore if we know that this action pair was played <u>somewhere</u> (and no matter how unlikely the contingency was), then we can infer that it is played everywhere. An analogous claim about any other equilibrium is impossible. More precisely we have the following (which is a special case of corollary 1).

COROLLARY 2: *Suppose that (a^1, a^2) is a risk-dominant action profile at every state, and that the stochastic potential is sufficiently close to 1/2. Then if $s^i(\omega) = a^i$ for some i and some ω then (a^1, a^2) is played everywhere.*

Notice that if each player has only two possible actions, corollary 2 implies that in any pure strategy Bayesian Nash equilibrium, only "simple" strategies are played. That is, there may be a pure strategy BNE where each player takes the action which is not in the risk-dominant action profile. But if one of risk-dominant actions is ever taken, each player takes the risk-dominant action everywhere by corollary 2. This argument need not apply when there are more than two actions.

Section 8: A noisy information system

Theorem 7.1 has power only if the stochastic potential is significantly above zero. It has most power if the stochastic potential is at or close to its maximum possible value, 1/2, so that we can guarantee that a risk dominant action is chosen. How likely or unlikely are we to get close to 1/2?

If there is a uniform prior in the leading example of section 2 (i.e, $\alpha = 1$), then the stochastic potential is actually equal to 1/2. Although there are well known stories generating information systems with that structure, it is rather special. In a finite state space, a stochastic potential of 1/2 must be

ged to be a rather extreme property.

On the other hand, with continuous signals (and thus an infinite state
ce), there are circumstances in which a stochastic potential of 1/2 arises
y naturally. Suppose that there is some true "circumstance" describing the
ld, which is distributed along a continuum. But one or both players observe
oisy version of the signal — that is, they observe some signal which is
metrically distributed around the true circumstance, and whose support is
irely within ε of the true circumstance. What probability does one player
ign to the other player observing a signal to the left of his? In this
tion, it is shown (in lemma 8.4) that, for any $q < 1/2$, we can choose ε
ficiently small such that this probability is greater than q.

This would imply that, if we extended the definition of stochastic
ential to infinite state spaces, we could find, for any $q < 1/2$, some $\varepsilon > 0$
that if the support of the noise is all within ε of the true circumstance,
stochastic potential is more than q. Therefore as $\varepsilon \to 0$ we can generate a
chastic potential which is as close to 1/2 as we wish. However, defining
ief operators on infinite state spaces raises some somewhat tangential (for
purposes) issues concerning conditional probabilities on zero probability
ts. Rather than deal with these issues in general, we provide a direct proof
t the noisy information system described here leads, via the infection
ument, to a globally risk dominant action to be played everywhere.

The framework of this section is similar to that of Carlsson and van Damme
92), and yields the same qualitative conclusions. One difference is that
lsson and van Damme exploit continuity between the state space and payoffs to
ive a stronger characterization of when the risk dominant outcome is played.
the other hand, our framework does not depend on that continuity.

Let us introduce the set of circumstances as the circle with unit radius,
oted by **C**. Nature chooses a circumstance according to a continuous density

$p(\cdot)$, where p

There a
true circumst
state space Ω
surface of a
the combinati
message y. W
number a, the
x by rotating
n. Ω may be
understanding
the top border

We shall foll
realization is
realization is
function $X{:}\Omega \to$
function $Y{:}\Omega \to$

30

$p(\cdot)$, where $p(x) > 0$, for all $x \in C$.

There are two players — player 1 and player 2. Player 1 can observe the true circumstance perfectly, but player 2 has some noise in his signal. The state space Ω is the product $C \times C$. Geometrically, Ω is a torus — that is, the surface of a doughnut whose cross section is C. The state $\omega = (x,y)$ represents the combination in which the true circumstance is x and player 2 observes the message y. We adopt the notational convention in which, for $x \in C$ and real number a, the expression $x + a$ denotes the point in C which may be reached from x by rotating clockwise by a radians. Hence, $x = x + 2\pi n$ for any x and integer n. Ω may be represented as in figure 3 as a square with side 2π, with the understanding that the left border is connected with the right border, and that the top border is connected to the bottom border.
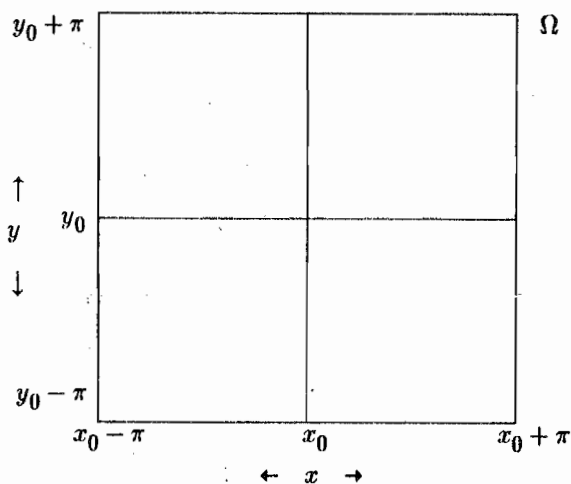
[Figure 3]

We shall follow the convention of denoting by X the random variable whose realization is the true circumstance, and denoting by Y the random variable whose realization is the message received by player 2. In other words, X is the function $X:\Omega \rightarrow C$ such that $X(\omega)$ is the first component of ω, and Y is the function $Y:\Omega \rightarrow C$ such that $Y(\omega)$ is the second component of ω.

The joint density over Ω is parametrized by number ε for which $0 < \varepsilon < \pi$,

is given by:

$$f(x,y,\varepsilon) \equiv p(x)\eta(x,y,\varepsilon)$$

$y,\varepsilon)$ represents the noise in 2's signal which is independent of the choice

he true circumstance. $\eta(x,y,\varepsilon)$ is continuous and satisfies the following

e properties.

) $$\int_{x-a}^{x+a} \eta(x,y,\varepsilon)dy = 1 \Leftrightarrow a \geq \varepsilon$$

) $$\eta(x,x+a,\varepsilon) = \eta(x,x-a,\varepsilon), \forall x, \forall a.$$

) $$\eta(x,x+a,\varepsilon) = \eta(x',x'+a,\varepsilon), \forall x, \forall x', \forall a.$$

first condition states that, conditional on x being the true circumstance,

noise has support $[x-\varepsilon, x+\varepsilon]$. (A2) states that the noise is symmetric around

true circumstance, and (A3) states that shape of the noise is invariant to

true circumstance. We can represent the support of the density I in terms

he shaded diagonal strip and the two shaded corner pieces as represented in

re 4.

[Figure 4]

32
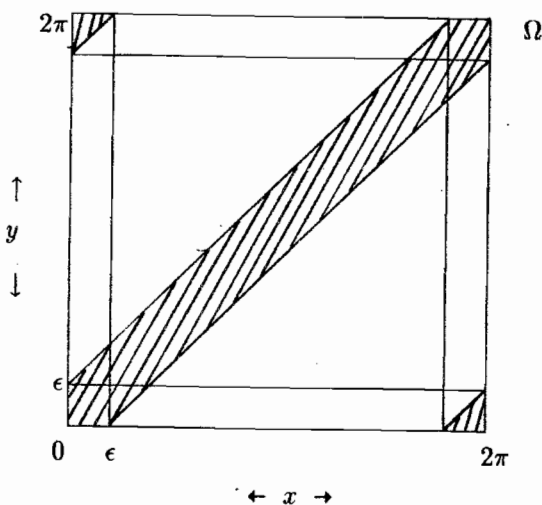
We then have:

LEMMA 8.1 (A2)

That is, for an

PROOF.

The following is

(A4)

We shall be inte

the probability

of y_0 given that

given by $\eta(x,y,\varepsilon$

$$Prob_\varepsilon$$

The following pr

We then have:

LEMMA 8.1 (A2) and (A3) imply that η is symmetric in the fist two components. That is, for any $x, y, \in C$,

$$\eta(x,y,\varepsilon) = \eta(y,x,\varepsilon).$$

PROOF.

$$\begin{aligned}
\eta(x,y,\varepsilon) &= \eta(x+(y-x),y+(y-x),\varepsilon) \quad by \ (A3), \\
&= \eta(y,y-(y-x),\varepsilon) \quad by \ (A2), \\
&= \eta(y,x,\varepsilon). \qquad \blacksquare
\end{aligned}$$

The following is an immediate consequence of this lemma and (A2).

$$(A4) \qquad \int_{y}^{x+\pi} \eta(x,y,\varepsilon)\,dx = \int_{y-\pi}^{y} \eta(x,y,\varepsilon)\,dx = \frac{1}{2}, \quad \forall \ y \in C.$$

We shall be interested in the following probabilities. We denote by:

$$Prob_{\varepsilon}(y_0 - \pi \le X \le y_0 | Y = y_0)$$

the probability that the true circumstance lies in the half circle to the left of y_0 given that the player 2 has received the message y_0, when the noise is given by $\eta(x,y,\varepsilon)$. In other words,

$$Prob_{\varepsilon}(y_0 - \pi \le X \le y_0 | Y = y_0) = \int_{y_0-\pi}^{y_0} f(x,y_0,\varepsilon)\,dx \Big/ \int_{y_0-\pi}^{y_0+\pi} f(x,y_0,\varepsilon)\,dx$$

The following probabilities are defined analogously.

$$Prob_{\varepsilon}(y_0 + \pi \ge X \ge y_0 | Y = y_0)$$

$$Prob_{\varepsilon}(x_0 - \pi \le Y \le x_0 | X = x_0)$$

$$Prob_{\varepsilon}(x_0 + \pi \geq Y \geq x_0 | X = x_0).$$

e shall appeal to the following features of these probabilities. The proofs are
resented in the appendix.

EMMA 8.2.

$$Prob_{\varepsilon}(x_0 - \pi \leq Y \leq x_0 | X = x_0) = \frac{1}{2}, \ \forall x_0, \forall \varepsilon, \qquad (8.1)$$

$$Prob_{\varepsilon}(y_0 - \pi \leq X \leq y_0 | Y = y_0) \to \frac{1}{2}, \ as \ \varepsilon \to 0, \ \forall y_0. \qquad (8.2)$$

EMMA 8.3. For any $q < \frac{1}{2}$ and $\varepsilon > 0$, there is a $\delta > 0$ such that, for all x_0
C we have:

$$Prob_{\varepsilon}(x_0 - \pi \leq Y \leq x_0 | X = x_0 + \delta) > q \qquad (8.3)$$

$$Prob_{\varepsilon}(x_0 - \pi \geq Y \geq x_0 | X = x_0 - \delta) > q. \qquad (8.4)$$

EMMA 8.4. For any $q < \frac{1}{2}$, there is an $\varepsilon^* > 0$ such that, for all $y_0 \in C$ and all
$< \varepsilon^*$,

$$Prob_{\varepsilon}(y_0 - \pi \leq X \leq y_0 | Y = y_0) > q \qquad (8.5)$$

$$Prob_{\varepsilon}(y_0 - \pi \geq X \geq y_0 | Y = y_0) > q. \qquad (8.6)$$

quilibrium in Bayesian Games.

With these preliminary results, we can address the question of what actions

will be pla

sets A^1 and

is a functi

strategy s^2

to an actio

level of no

Bayesian ga

We sha

equilibrium.

critical do

probability

conversely,

2's unique b

THEOREM 8.1.

are numbers ε

any equilibri

$\in (b,c)$, ther

34

will be played in equilibrium. The two players face a game G in which the action sets A^1 and A^2 are finite. The <u>strategy</u> of player 1 is denoted by s^1, where s^1 is a function which maps each circumstance $x \in \mathbf{C}$ to an action in A^1. Player 2's strategy s^2 is a function which maps each realization $y \in \mathbf{C}$ of player 2's message to an action in A^2. If strategies are best replies to each other given some level of noise ε, we shall say that these strategies are an <u>equilibrium of the Bayesian game with noise ε.</u>

We shall be interested in the case in which the game G has a $\frac{1}{2}$-dominant equilibrium. In other words, there is some action pair (a^1, a^2) and some critical dominance level $q < \frac{1}{2}$ for which, as long as player 2 places probability greater than q on a^2, player 1's unique best reply is to play a^1, and conversely, as long as player 1 places probability greater than q on a^1, player 2's unique best reply is to play a^2. We have the following theorem.

THEOREM 8.1. Suppose G has a $\frac{1}{2}$-dominant equilibrium (a^1, a^2). Then, there are numbers $\varepsilon^* > 0$ and $\delta > 0$ such that, for any non-empty open interval (b,c) and any equilibrium of the Bayesian game with noise $\varepsilon < \varepsilon^*$, if $s^1(x) = a^1$ for all $x \in (b,c)$, then $s^1(x) = a^1$ for all $x \in (b - \delta, c + \delta)$.

$$b' = b - \delta$$
$$c' = c + \delta$$
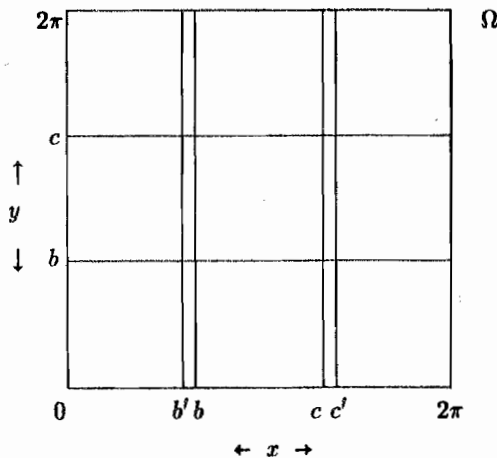
[Figure 5]

OOF. Let $q < \frac{1}{2}$ be the critical dominance level for the $\frac{1}{2}$-dominant

uilibrium (a^1, a^2) of G. Consider player 2's inference problem when faces with

e message $b \in C$. Refer to figure 5. By lemma 8.4, there is some benchmark ε^*

0 such that, for any $\varepsilon < \varepsilon^*$ and any $\omega \in \Omega$, player 2 places probability greater

an q on the event $\{\omega | b \leq X(\omega) \leq c\}$ for a sufficiently small value of the

nchmark ε^*. Thus, for all $\varepsilon < \varepsilon^*$,

$$Prob_\varepsilon(b \leq X \leq c | Y = b) > q. \qquad (8.7)$$

an exactly analogous argument for the message $c \in C$, we have:

$$Prob_\varepsilon(b \leq X \leq c | Y = c) > q. \qquad (8.8)$$

r any message $y_0 \in (b,c)$, if $\varepsilon < \varepsilon^*$ player 2 places probability greater than

on the event $\{\omega | b \leq X(\omega) \leq c\}$. Hence, together with (8.7) and (8.8), for all

$< \varepsilon^*$, we have:

$$Prob_\varepsilon(b \leq X \leq c | Y = y_0) > q, \quad \forall y_0 \in [b,c]. \qquad (8.9)$$

hypothesis, the action pair (a^1, a^2) is a $\frac{1}{2}$-dominant equilibrium of G, and

is the critical dominance level of this equilibrium. Also, by hypothesis,

ayer 1 plays a^1 on the event $\{\omega | X(\omega) \in (b,c)\}$. Thus, by (8.9), player 2 plays

on the event $\{\omega | X(\omega) \in [b,c]\}$.

We now turn to player 1's inference in the neighborhood of the states in

$|X(\omega) - b\}$ and $\{\omega | X(\omega) - c\}$ given that player 2 plays a^2 on the event $\{\omega | Y(\omega)$

$[b,c]\}$. By lemma 8.3, for any level of noise ε, there is a number $\delta > 0$ which

uniform across the states such that, given the message $x_0 - \delta$, player 1 places

obability greater than q on the event $\{\omega | x_0 \leq Y(\omega) \leq x_0 + \pi\}$. In particular,

r states ω for which $X(\omega) - b - \delta$, player 1 places probability greater than q

the event $\{\omega | b \leq Y(\omega) \leq b + \pi\}$. Hence,

36

$$Prob_\varepsilon(b \le Y \le c \,|\, X = b - \delta) > q.$$

Similarly, by lemma 8.3, there is some $\delta > 0$ which is uniform across the states such that, when faced with the message $c + \delta$, 1 places probability greater than q to the event $\{\omega \,|\, b < Y(\omega) \le c\}$. That is,

$$Prob_\varepsilon(b \le Y \le c \,|\, X = c + \delta) > q.$$

Since player 2 plays a^2 on the event $\{\omega \,|\, b < Y(\omega) \le c\}$, and since (a^1, a^2) is a $\frac{1}{2}$ dominant equilibrium with critical value q, player 1 plays a^1 on $\{\omega \,|\, X(\omega) = b - \delta\}$ and on $\{\omega \,|\, X(\omega) = c + \delta\}$. Moreover, player 1 plays a^1 given every realization of X between $b - \delta$ and $c + \delta$. This implies the statement of theorem 8.1. ∎

COROLLARY. Suppose G has a $\frac{1}{2}$-dominant equilibrium (a^1, a^2). Then, if a^1 is played on any non-empty interval of the set of circumstances, then (a^1, a^2) is played at every state.

PROOF. Since ε^* and δ are defined independently of the interval (b, c), theorem 1 can be applied to $(b-\delta, c+\delta)$ to yield the conclusion that a^1 is played by 1 given any message in $(b-2\delta, c+2\delta)$. By iteration, a^1 is played by 1 everywhere on Ω. Hence, (a^1, a^2) is played everywhere on Ω. ∎

Section 9: Conclusion

This paper introduces the idea that the existence of higher order uncertainty can actually pin down outcomes in games with multiple equilibria. The essence of the idea is that a 'small grain of knowledge' can be extrapolated to have implications at states that are completely different from the state where

37

the knowledge was originally available. For instance, in the example of section 2 an investor's decision is sometimes dictated by the fact that the other investor knows — in some set of circumstances — something which he does not know, although he also knows that such knowledge is certainly not available in the _present_ set of circumstances. Therefore it is sometimes possible to determine the action taken in one state on the basis of which action is taken at another state.

This idea is already suggested in Rubinstein's electronic mail game example, where the risk dominant action pair in a 2×2 game turns out to be the selected equilibrium. Here we develop this idea in a broader context, showing that the appropriate generalization of risk dominance is p-dominance, and that whether a unique equilibrium is indeed selected depends on the specification of the information structure (in particular, its stochastic potential), and on the level of p-dominance of the different action profiles. These two concepts—p-dominance and the stochastic potential—might be of independent interest apart from how (and whether) they determine which equilibrium is selected.

We close by indicating a few directions for future research. First, our focus here has been on generating sufficient conditions for the infection argument to work. It would be of interest to generate necessary conditions as well, delineating thereby the exact circumstances when the infection argument has power. Second we have analyzed systematically only the finite information structure case, where a large stochastic potential (close to 1/2) is rather unusual. On the other hand, the example in section 8 shows that it is quite easy to generate a stochastic potential close to 1/2 in an infinite state space. It remains an open question whether this is an artifact of the 'geometric' nature of the example or whether this feature arises more generally in a continuous framework. We hope to address these issues in future work.

Aumann F

Carlsson
Forthcor

Foster
Populati

Geanakop

Halpern
reasonir
Kaufman

Harsanyi
MIT Pres

Kandori
games.

Kreps D,
253–279

Milgrom
Theory

Mondere
Games E

Rubinst
common

Shin H,
CARESS

REFERENCES

Aumann R (1976) Agreeing to disagree. *Annals of Statistics* 4: 1236–1239.

Carlsson H, van Damme E (1992) Global games and equilibrium selection. Forthcoming in *Econometrica*.

Foster D, Young P (1990) Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38: 219–232.

Geanakoplos J (1992) Common knowledge. *J Econ Perspectives* 6: 53–82.

Halpern J (1986) Reasoning about knowledge: An overview. *Theoretical aspects of reasoning about knowledge: proceedings of the 1986 conference.* Los Altos: Morgan Kaufman.

Harsanyi J, Selten R (1988) *A general theory of equilibrium in games.* Cambridge: MIT Press.

Kandori M, Mailath G, Rob R (1993) Learning, mutation and long–run equilibria in games. *Econometrica* 61: 21–58.

Kreps D, Wilson R (1982) Reputation and imperfect information. *J Econ Theory* 27: 253–279.

Milgrom P, Roberts J (1982) Predation, reputation and entry deterrence. *J Econ Theory* 27: 280–312.

Monderer D, Samet D (1989) Approximating common knowledge with common beliefs. *Games Econ Behavior* 1: 170–190.

Rubinstein A (1989) The electronic mail game: Strategic behavior under almost common knowledge. *Amer Econ Rev* 79: 385–391.

Shin H, Williamson T (1992) How much common belief is necessary for a convention? CARESS Working Paper #92-32, University of Pennsylvania.

Appendix

OF OF LEMMA 8.2. For (8.1), note that:

$$b_\varepsilon(x_0 - \pi \leq Y \leq x_0 | X = x_0) = \int_{x_0-\pi}^{x_0} f(x_0, y, \varepsilon)dy \Big/ \int_{x_0-\pi}^{x_0+\pi} f(x_0, y, \varepsilon)dy$$

$$= p(x_0) \int_{x_0-\pi}^{x_0} \eta(x_0, y, \varepsilon)dy \Big/ p(x_0) \int_{x_0-\pi}^{x_0+\pi} \eta(x_0, y, \varepsilon)dy$$

$$= 1/2, \quad by \quad (A2).$$

(8.2), we start with two definitions. For some point $y_0 \in C$ and $\varepsilon > 0$, let:

$$M(y_0, \varepsilon) = max\{p(x) \,|\, y_0 - \varepsilon \leq x \leq y_0 + \varepsilon\}$$

$$m(y_0, \varepsilon) = min\{p(x) \,|\, y_0 - \varepsilon \leq x \leq y_0 + \varepsilon\}.$$

numbers are well-defined since p is continuous, and the closed interval is
ct. Also, since $p(x) > 0$ for all x, both $m(y_0, \varepsilon)$ and $M(y_0, \varepsilon)$ are positive.

$$Prob_\varepsilon(y_0 - \pi \leq X \leq y_0 | Y = y_0) = \int_{y_0-\pi}^{y_0} p(x)\eta(x, y_0, \varepsilon)dx \Big/ \int_{y_0-\pi}^{y_0+\pi} p(x)\eta(x, y_0, \varepsilon)dx$$

$$\leq M(y_0, \varepsilon) \int_{y_0-\pi}^{y_0} \eta(x, y_0, \varepsilon)dx \Big/ m(y_0, \varepsilon) \int_{y_0-\pi}^{y_0+\pi} \eta(x, y_0, \varepsilon)dx$$

$$= \frac{1}{2}\left[\frac{M(y_0, \varepsilon)}{m(y_0, \varepsilon)}\right], \quad by \quad (A4).$$

similar argument,

$$Prob_\varepsilon(y_0 - \pi \leq X \leq y_0 | Y = y_0) \geq \frac{1}{2}\left[\frac{m(y_0, \varepsilon)}{M(y_0, \varepsilon)}\right].$$

e,

Since p is c
to zero. Her

PROOF OF L

$$Prob_\varepsilon(x_0 - \pi$$

From lemma 8.
$\delta > 0$ for whi
this δ is uni

PROOF OF LE
value is well
compactness o
positive numb

40

$$\frac{1}{2}\left[\frac{m(y_0,\varepsilon)}{M(y_0,\varepsilon)}\right] \leq Prob_\varepsilon(y_0 - \pi \leq X \leq y_0 | Y = y_0) \leq \frac{1}{2}\left[\frac{M(y_0,\varepsilon)}{m(y_0,\varepsilon)}\right]. \qquad (**)$$

Since p is continuous, the expressions in square brackets tend to 1 as ε tends to zero. Hence, the probability in question tends to a half as ε tends to zero. ∎

PROOF OF LEMMA 8.3.

$$
\begin{aligned}
Prob_\varepsilon(x_0 - \pi \leq Y \leq x_0 | X = x_0 + \delta) &= \int_{x_0-\pi}^{x_0} f(x_0+\delta,y,\varepsilon)\,dy \Big/ \int_{x_0-\pi}^{x_0+\pi} f(x_0+\delta,y,\varepsilon)\,dy \\
&= \int_{x_0-\pi}^{x_0} \eta(x_0+\delta,y,\varepsilon)\,dy \Big/ \int_{x_0-\pi}^{x_0+\pi} \eta(x_0+\delta,y,\varepsilon)\,dy \\
&= \int_{x_0-\pi}^{x_0} \eta(x_0+\delta,y,\varepsilon)\,dy.
\end{aligned}
$$

From lemma 8.2, this expression tends to $1/2$ as δ tends to zero. Thus, there is $\delta > 0$ for which the expression is strictly greater than q. Moreover, by (A3), this δ is uniform across all $x_0 \in \mathbf{C}$. By symmetry, (8.4) also holds. ∎

PROOF OF LEMMA 8.4. Denote by p_{min} the minimum value of $p(x)$ over $\mathbf{C}$. This value is well-defined since p is continuous and $\mathbf{C}$ is compact. Moreover, the compactness of $\mathbf{C}$ implies that p is underlined{uniformly} continuous over $\mathbf{C}$. Consider the positive number $(1-2q)p_{min}$. By uniform continuity, there is $\varepsilon^* > 0$ such that:

$$|x - x'| < \varepsilon^* \Rightarrow |p(x) - p(x')| < (1 - 2q)p_{\min}, \quad \forall x, x',$$

Hence for this value of ϵ^*, $M(y_0, \varepsilon^*) - m(y_0, \varepsilon^*) < (1-2q)p_{\min}, \quad \forall y_0,$

$$\Rightarrow M(y_0, \varepsilon^*) - m(y_0, \varepsilon^*) < (1 - 2q)M(y_0, \varepsilon^*), \quad \forall y_0,$$

$$\Rightarrow 2qM(y_0, \varepsilon^*) - m(y_0, \varepsilon^*) < 0, \quad \forall y_0,$$

$$\Rightarrow \frac{1}{2}\left[\frac{m(y_0, \varepsilon^*)}{M(y_0, \varepsilon^*)}\right] > q, \quad \forall y_0.$$

together with (*), for all $\varepsilon < \varepsilon^*$,

$$Prob_\varepsilon(y_0 - \pi \le X \le y_0 | Y = y_0) \ge \frac{1}{2}\left[\frac{m(y_0, \varepsilon)}{M(y_0, \varepsilon)}\right] \ge \frac{1}{2}\left[\frac{m(y_0, \varepsilon^*)}{M(y_0, \varepsilon^*)}\right] > q, \quad \forall y_0.$$

demonstrates (8.5). By symmetry, (8.6) also holds.