

p-Dominance and Belief Potential

Author(s): Stephen Morris, Rafael Rob and Hyun Song Shin

Source: *Econometrica*, Vol. 63, No. 1 (Jan., 1995), pp. 145-157

Published by: [The Econometric Society](#)

Stable URL: <http://www.jstor.org/stable/2951700>

Accessed: 29-05-2015 02:30 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The *Econometric Society* is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*.

<http://www.jstor.org>

p-DOMINANCE AND BELIEF POTENTIAL

BY STEPHEN MORRIS, RAFAEL ROB, AND HYUN SONG SHIN¹

This paper elucidates the logic behind recent papers which show that a unique equilibrium is selected in the presence of higher order uncertainty, i.e., when players lack common knowledge. We introduce two new concepts: belief potential of the information system and *p*-dominance of Nash-equilibria of the game, and show that a Nash-equilibrium is uniquely selected whenever its *p*-dominance is below the belief potential. This criterion applies to many-action games, not merely 2×2 games. It also applies to games without dominant strategies, where the set of equilibria is shown to be smaller and simpler than might be initially conjectured. Finally, the new concepts help understand the circumstances under which the set of equilibria varies with the amount of common knowledge among players.

KEYWORDS: Common knowledge, higher-order uncertainty, infection argument, equilibrium selection, risk-dominance, stochastic potential.

1. INTRODUCTION

WHEN PAYOFFS IN A GAME are not common knowledge, the outcome depends not only on players' beliefs about payoffs, but also on their beliefs about others' beliefs about payoffs, on their beliefs about others' beliefs about their own beliefs, and so on ad infinitum. Initially, it might be expected that such a hierarchy of beliefs would, if anything, enlarge the set of equilibrium outcomes. This paper argues, however, that in certain games it *reduces* the set of equilibria to a unique outcome. The logic behind this is as follows. Suppose that one player is known to take a certain action at some information set (which may itself have a very small probability). Then this knowledge might imply a unique best response by some other player at information sets where the first information set is thought possible. This, in turn, implies how the original player responds to that knowledge, at a yet larger information set, and so on. If this chain of reasoning results in a unique action profile, then we have elicited a unique equilibrium.

Infection arguments of this kind have been used in a number of recent papers. Rubinstein's (1989) electronic mail game example showed how outcomes of a game with common knowledge of payoffs may be very different from outcomes of the game with a "small" departure from common knowledge. Carlsson and van Damme (1993) showed how a natural representation of "global uncertainty" implies that the risk-dominant action is always played in any particular realization of payoffs in a two-person, two-action game. Shin and Williamson (1992) showed that in a class of coordination games with incomplete information, the

¹ We thank seminar participants at CORE, SITE, Churchill College (Cambridge), Nuffield College (Oxford), Warwick University, Universidad Autònoma de Barcelona, Universidad Carlos III de Madrid, the 1993 summer workshop at Gerzensee, Switzerland and the GREQE Marseille Conference on "Epistemic Logic and the Theory of Games." Special thanks are due to Dieter Balkenborg, Chris Harris, Atsushi Kajii, George Mailath, Meg Meyer, Dov Samet, and Fernando Vega-Redondo.

only pure-strategy equilibria involve each player choosing a “simple” strategy—a strategy which is constant across states—even when state-contingent strategies are Pareto-improving. Each of these papers uses a version of the infection argument.

Our primary purpose here is to delineate, for two-person finite games of incomplete information, exactly which properties of payoffs and the information structure allow such an argument to operate. When is it the case that some player choosing a certain action somewhere on the state space implies that each player chooses the same action everywhere on the state space? We give a *global* version of the infection argument, a version which is independent of where the infection starts. In Section 3, we introduce the notion of *belief potential*, which is a property of the information system and measures the extent to which information sets overlap. Roughly, the belief potential is the largest probability p such that, for any information set of either player, some statement of the form “1 believes with probability at least p that 2 believes with probability at least p that 1 believes... that the true state is in the original information set” is true at every state. In Section 4, we introduce the notion of p -dominance for a strict Nash equilibrium. An action pair is p -dominant if each action is a best response to any conjecture placing at least probability p on the other player taking his action in the pair. Our results show that the relationship between the belief potential of the information system and the level of p -dominance of Nash-equilibria of the game determine whether the infection argument operates.

These new concepts allow us to evaluate and extend earlier arguments using the infection argument. Carlsson and van Damme (1993) showed how a certain kind of global uncertainty implies that a risk-dominant equilibrium is selected in 2×2 games. Interestingly, this prediction coincides with other attempts to select among strict Nash equilibria by other methods—i.e., the axiomatic approach of Harsanyi and Selten (1988) and the evolutive approach of Kandori, Mailath, and Rob (1993) and Young (1993). Although these earlier attempts provide clear answers for 2×2 games, the mode of their argument is specific to this class, with the consequence that generalizations to larger classes of games are not obvious. However, by appeal to the notions of belief potential and p -dominance, we are able to construct an argument for *many-action* games. In particular, we show that if some action pair is p -dominant at every state of the world, one of those actions is a dominant strategy for some player at some information set, and the belief potential exceeds p , then that action pair is played everywhere. Moreover, this result is based on iterative deletion of dominated strategies, and not on the more restrictive equilibrium reasoning. Conversely, we show that for any p greater than the belief potential, it is possible to construct a game where some action pair is p -dominant at every state of the world, one of those actions is a dominant strategy for some player at some information set, and yet a pure-strategy Bayesian Nash equilibrium involves other actions being played somewhere on the state space.

Another application of our concepts is to the continuity of equilibria with respect to the amount of common knowledge that players are assumed to

possess. Rubinstein (1989) argued that his example shows that a slight departure from common knowledge (in the sense of a large but finite number of levels of iterated knowledge) alters discontinuously the outcome(s) that would 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 an appropriate definition of “closeness” on information structures.² In particular, an event is said to be *common p -belief* if everyone believes it with probability at least p , everyone believes with probability p that everyone believes it with probability p , and so on ad infinitum. If payoffs are common p -belief, for p close to one, then equilibrium outcomes are close to those under common knowledge. We show that the infection argument operates *exactly* when there is *no* nontrivial event which is common p -belief for p close to one. But we also argue that this situation arises naturally in the finite (or even countably infinite) state-space case. Therefore, in such cases there is no contradiction between the Rubinstein and the Monderer-Samet approaches, although the issue acquires greater interest in the infinite state-space case.

Carlsson and van Damme (1993) and our main result rely on the action of some player at some information set starting the infection argument. But it is also possible to make a conditional statement: if a certain action is played in some Bayesian Nash equilibrium, then it is played everywhere in that equilibrium. This means that, in a two-action game, if a certain action is p -dominant everywhere, then the only possible pure-strategy Bayesian Nash equilibria are to always play the p -dominant action pair, or to always play the other action pair. In other words, we rule out the possibility of equilibria where a player chooses one action at a certain subset of states and then switches to another action at other states. This result illustrates that the infection argument can reduce the set of equilibria even when there are no dominant strategies. Shin and Williamson’s (1992) more general argument for such “simple” strategies uses similar logic but relies on *continuous* action spaces.

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 define the belief potential of an information system. In Section 4, we characterize the p -dominance of strict Nash equilibria. Section 5 connects these concepts and gives the main results.

2. LEADING EXAMPLE

Consider the following incomplete information game. There is a state space consisting of $2N$ states, $\Omega = \{1, 2, \dots, 2N\}$, with each state equally likely. There are two players: 1 observes information partition, $\mathcal{P}_1 = (\{1\}, \{2, 3\}, \dots, \{2N - 2, 2N - 1\}, \{2N\})$, while 2 observes partition, $\mathcal{P}_2 = (\{1, 2\}, \{3, 4\}, \dots, \{2N - 1, 2N\})$.

² Stinchcombe (1989) introduced a related notion of almost common knowledge, which is in between Rubinstein’s and Monderer-Samet’s.

Information systems with this structure of overlapping information sets arise naturally in many classes of problems where there is a lack of common knowledge. Examples include the coordinated 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).

Now suppose that at each information set, each player has a choice of two actions, L or R ; if the true state is ω , then payoffs are:

		Player 2	
		L	R
Player 1	L	2, 2	0, 0
	R	0, 0	$x_1(\omega), x_2(\omega)$

where $x_i: \Omega \rightarrow (-\infty, 2)$, and $x_i(\omega)$ is constant on player i 's information set for $i = 1, 2$. Now suppose that in addition, $x_i(\omega) < 0$ for some i and ω . For $x_i(\omega)$ positive, there are two strict Nash equilibria of this 2×2 game. However, for the incomplete-information game as a whole, playing " L " *everywhere* is the only strategy which survives iterated deletion of dominated strategies. To see why, observe that player i must play " L " at the information set which contains the ω for which $x_i(\omega) < 0$. Suppose that i 's information at ω is $\{\omega, \omega + 1\}$ and that $\omega + 1 < 2N$ (an analogous argument holds if i 's information at ω is $\{\omega - 1, \omega\}$). But now, the other player, j , has information sets $\{\omega - 1, \omega\}$ and $\{\omega + 1, \omega + 2\}$ overlapping that first information set. At those information sets, she assigns probability at least $1/2$ to player i choosing " L ." But then, by the assumption that $x_j(\omega) < 2$, for all $\omega \in \Omega$, " L " must be the best response for player j on those information sets. But now an inductive argument shows this must be true for all states. Note that for this argument it does not matter at which ω we have $x_i(\omega) < 0$, only that such an ω exists. Hence, the infection argument is *global*, not local.

We can also see in this example that an infection argument makes interesting predictions even when there does not exist a dominant strategy. Suppose now that $x_i(\omega) \geq 0$, for all $\omega \in \Omega$ and $i = 1, 2$. Consider a pure-strategy Bayesian Nash Equilibrium (BNE) where " L " is played by some player at some information set. By *exactly* the argument above, we see that " L " must be played everywhere in that equilibrium. Thus, there exist exactly two pure-strategy BNE: both always play L and both always play R .

We want to explore what makes this infection argument operate in general. In this example, two properties were key. First, action L was always the unique best response for any player on any information set where he assigned probability $1/2$ to the other player choosing L . Second, if we fix any information set E of any player, some statement of the form " 1 believes with probability $1/2$ that 2 believes with probability $1/2$ that 1 believes... that the true state is in E " is always true. The crucial fact is that the probability in the second statement is no

smaller than the probability in the first (the fact that the probabilities are both $1/2$ is immaterial—any probabilities in $(0, 1)$ will do). The next sections generalize this idea by introducing properties of the information system and properties of the payoffs, and showing that the relationship between them gives a sufficient condition for the infection argument to operate.

3. THE INFORMATION SYSTEM AND THE BELIEF POTENTIAL

In this section, we show how a single number, the *belief potential*, can be thought of as measuring the susceptibility of the information system to infection. First, we introduce notation. An *information system* is the structure $[\Omega, \{1, 2\}, \{\mathcal{P}_i\}_{i=1,2}, \pi]$, where Ω is a finite set of states of the world; $\{1, 2\}$ is the set of players; \mathcal{P}_i is the partition of states of the world representing the information of player i ; and 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 write \mathcal{F}_i for the field generated by i 's partition, i.e., the set of unions and intersections of events in \mathcal{P}_i . We assume that there is some nontrivial information so that for some i , $\mathcal{P}_i \neq \{\Omega\}$.

The belief potential is defined using the idea of *belief operators* on the state space. Such belief operators were introduced in a related context by Monderer and Samet (1989). Write the conditional probability of event E , given event F , as $\pi[E|F] = \sum_{\omega \in E \cap F} \pi(\omega) / \sum_{\omega \in F} \pi(\omega)$. Now define player 1's p -belief operator, $B_1^p: 2^\Omega \rightarrow 2^\Omega$, by

$$(3.1) \quad B_1^p E = \{\omega \in \Omega | \pi(E|P_1(\omega)) \geq p\}.$$

Thus $B_1^p E$ is the set of states where player 1 attaches probability at least p to the event E . With B_2^p defined analogously, $B_2^p B_1^p E$ is the set of states where 2 believes with probability p that 1 believes with probability p that event E will occur. We will be interested in the set of states where either this is the case, or E is true. Thus define the operator $H_1^p(\cdot)$ as $H_1^p E \equiv B_1^p B_2^p E \cup E$. Notice that, for any given prior π , there exists an $\varepsilon > 0$ such that, for all $p \leq \varepsilon$, $B_1^p E = \{\omega \in \Omega | P_1(\omega) \cap E \neq \emptyset\}$. Therefore, if there is no nontrivial common knowledge and p is sufficiently small, H_1^p is a strictly increasing function on nontrivial events in \mathcal{F}_i i.e., $E' \supset E \Rightarrow H_1^p(E') \supset H_1^p(E)$ (\subset means "weak inclusion"). Thus progressive application of H_1^p will eventually cover the whole state space. We will be interested in the largest p for which this is true.

DEFINITION: The *belief potential*, $\sigma_1(E)$, of an event E is the largest number p such that, for some $k \geq 1$, $[H_1^p]^k E = \Omega$.

To see that every event has a well defined belief potential, notice first that $p > q$ implies $B_1^p E \subset B_1^q E$: 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 $[H_1^p]^k E \subset [H_1^q]^k E$. But since belief operators are defined in terms of a weak inequality, $\omega \in B_i^q E$ for all $q < p$ implies $\omega \in B_i^p E$. So there exists a largest p such that the above property holds.

It would be possible to derive results using this local (i.e., event-specific) notion of belief potential. But our main result, like the example in the previous section, does not depend on a particular event starting the infection argument. The belief potential of the information system is the largest p such that the infection argument works for *every* nontrivial event measurable with respect to some individual's partition.

DEFINITION: The *belief potential*, σ , of an information system is the minimum belief potential of any nonempty measurable event in the system:

$$(3.2) \quad \sigma \equiv \min_{i \in \{1, 2\}, E \in \mathcal{F}_i \setminus \emptyset} \sigma_i(E).$$

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

$$\Omega = \{1, 2, 3, 4, 5, 6, 7, 8, 9\},$$

$$\mathcal{P}_1 = (\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8, 9\}),$$

$$\mathcal{P}_2 = (\{1, 4, 7\}, \{2, 5, 8\}, \{3, 6, 9\}),$$

$$\pi = \left\{ \frac{1}{21}, \frac{1}{7}, \frac{1}{7}, \frac{1}{7}, \frac{1}{21}, \frac{1}{7}, \frac{1}{7}, \frac{1}{7}, \frac{1}{21} \right\}.$$

Consider event $E = \{1, 2, 3\}$. If $p > 3/7$; then $B_2^p E = \emptyset$. Thus the belief potential of E can be no more than $3/7$. Suppose $1/7 < p \leq 3/7$. Then $B_2^p E = \{2, 3, 5, 6, 8, 9\}$ and $B_1^p B_2^p E = \Omega$. Thus $\sigma_1(E) = 3/7$. By the symmetry of the example, the belief potential of every other information set is also $3/7$; larger events measurable on some individual's partition have higher belief potential, and so the (global) belief potential of this information system is $3/7$.

As we will see in Section 5, the notion of belief potential is precisely what is needed in characterizing those situations in which the "infection argument" operates.³ It can also be shown that the belief potential is the mirror image of the notion of common belief, as defined by Monderer and Samet (1989). The following lemma is important in developing our arguments, and has some independent interest. All proofs for this section are presented in the Appendix.

LEMMA 3.1: Suppose σ is the belief potential of the information system. Then, the following three statements are equivalent:

- (i) $p > \sigma$.
- (ii) There is an event E such that $B_j^p E \subset E$ for all j and $B_i^p E \notin \{\emptyset, \Omega\}$ for some i .
- (iii) There is an event $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$ such that, for $j \neq i$, $B_i^p B_j^p E_i \subset E_i$.

³ Sorin (1993) introduces a related measure of the global impact of an event on an information system.

Monderer and Samet (1989, p. 177) show that the notion of common belief can be given a “fixed point” characterization in terms of “ p -evident” events. An event E is said to be p -evident if $E \subset B_j^p E$ for all j . Event A is common p -belief at ω if and only if there exists a p -evident event E such that $\omega \in E \subset B_j^p A$, for all j . Our claim that the belief potential is the mirror image of common belief is embodied in the following result.⁴

THEOREM 3.1: *Suppose σ is the belief potential of the information system. Then $\sigma \leq 1 - p$ if and only if there is a p -evident event E for which $B_i^p E \not\subseteq \{\emptyset, \Omega\}$ for some i .*

Thus, for a given belief potential, we can guarantee the existence of a $(1 - \sigma)$ -evident event which is believed somewhere, but not everywhere, by some individual. Conversely, if we find a nontrivial p -evident event with the largest possible p , then $\sigma = 1 - p$. Anticipating our main result concerning the infection argument in Section 5, we can say that Theorem 3.1 provides a link between common beliefs and the infection argument.

It is a consequence of Theorem 3.1 that the belief potential of an information system is zero if and only if there exist nontrivial subsets of the state space which are 1-evident, which in this context implies they are common knowledge whenever they are true. This shows when the lower bound of the belief potential is attained. We can also provide an upper bound.

THEOREM 3.2: *The belief potential of any information system is less than or equal to $1/2$.*

Theorem 3.2 is tight in the sense that $1/2$ is not only an upper bound on the belief potential, but is also *attained* for some information systems. The information system in the example of Section 2 has a belief potential of $1/2$.

4. p -DOMINANCE

Having set out the key property of the information system, we now turn to payoffs. 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 in the previous section; A_i is the finite set of actions available to player i ; and $g_i: A \times \Omega \rightarrow \mathbb{R}$ is player i 's payoff function, where $A = A_1 \times A_2$. Assume players know their own payoff function, i.e., $g_i(a; \omega)$ is measurable with respect to \mathcal{P}_i .

Thus, in each state ω , there is a one-shot (complete-information) game, with payoff function $g_i(a_i, a_j; \omega)$. A pure strategy for player i in the incomplete-information game is a function $s_i: \Omega \rightarrow A_i$, measurable with respect to his partition. Write S_i for the set of such pure strategies. Now a pure-strategy pair

⁴ A related treatment can be found in Monderer and Samet (1990, p. 15–16) although they don't explicitly link common beliefs and the infection argument.

(s_1, s_2) is a *Bayesian Nash equilibrium* of the incomplete-information game if, for $i \neq j$ and all $t_i \in S_i$,

$$\sum_{\omega \in \Omega} \pi(\omega) g_i[s_i(\omega), s_j(\omega); \omega] \geq \sum_{\omega \in \Omega} \pi(\omega) g_i[t_i(\omega), s_j(\omega); \omega].$$

For a pair of actions (a_1, a_2) to infect the state space, we require the notion that a_i 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 every probability distribution λ on A_j such that $\lambda(a_j) \geq p$ and all $b_i \in A_i$,

$$\sum_{b_j \in A_j} \lambda(b_j) g_i(a_i, b_j; \omega) > \sum_{b_j \in A_j} \lambda(b_j) g_i(b_i, b_j; \omega).$$

If (a_1, a_2) is *p-dominant*, a_i is the *unique* best response for player i as long as he believes that the other player will play a_j with probability at least p . 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 action sets $B_1 \equiv \{a_1, b_1\}$, $B_2 \equiv \{a_2, b_2\}$ 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 $p \leq q \leq 1$.

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. Also notice that the notion of *p-dominance* introduced here generalizes 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 $i \neq j$ and for all $b_1 \in A_1$ and $b_2 \in A_2$,

$$\frac{1}{2} g_i(a_i, a_j; \omega) + \frac{1}{2} g_i(a_i, b_j; \omega) > \frac{1}{2} g_i(b_i, a_j; \omega) + \frac{1}{2} g_i(b_i, b_j; \omega).$$

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

An example will illustrate the character of *p-dominance* and the relation to risk-dominance. Consider the following symmetric game.

		Player 2		
		L	C	R
Player 1	T	7, 7	0, 0	0, 0
	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) . Thus, the relationship may very well be *cyclical*. The notion of p -dominance we introduced is *global*, in the sense that each strict Nash equilibrium action pair is being compared with all other actions. In our terminology, (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 $2/3$. So, for any $8/15 < p \leq 2/3$, (T, L) is the *unique* p -dominant action pair. Therefore, a suitable choice of p can eliminate the cycle, and pick a unique p -dominant action pair; but in this example a p greater than $1/2$ is required. More generally, Lemma 4.1 below shows the following: 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. The proof to this lemma is in the Appendix.

LEMMA 4.1: (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 .*

The main result of the next section is applicable only if there exists an action pair which is p -dominant for some p less than or equal to $1/2$ (since the belief potential is no larger than $1/2$, as Theorem 3.2 shows). Lemma 4.1 shows that there is *at most* one such action pair. But how likely is it that there *exists* such a pair? Note first that for generic symmetric 2 by 2 games with pure equilibria, there is *always* a risk-dominant, and thus a $1/2$ -dominant equilibrium. But for many-action games there need not exist $1/2$ -dominant equilibria, as the example above shows. On the other hand, an earlier version of our paper (Morris, Rob, and Shin (1993)) shows that $1/2$ -dominant equilibria can occur in some contexts, including a partnership game with many actions.

5. BELIEF POTENTIAL, p -DOMINANCE, AND RISK DOMINANCE

We are now in a position to state the main theorem of the paper.

THEOREM 5.1: *Suppose that (1) the information system has belief potential of σ , (2) (a_1, a_2) is σ -dominant at every state; and (3) some player i at some state ω^* knows the event “ a_i is a strongly dominant action.” Then playing (a_1, a_2) everywhere is the unique rationalizable strategy profile of the incomplete-information game.*

PROOF: Let R_i be the set of rationalizable strategies of i . Let $\Omega_i = \{\omega \in \Omega \mid s_i(\omega) = a_i, \text{ for all } s_i \in R_i\}$. Thus Ω_i is the set of states where i plays

action a_i in every rationalizable strategy. Let $E = P_i(\omega^*)$, i.e., the event where a_i is strictly dominant for player i ; clearly $E \subset \Omega_i$. Now by (2), $B_i^\sigma \Omega_i \subset \Omega_j$ and $B_i^\sigma \Omega_j \subset \Omega_i$. Thus $H_i^\sigma E = B_i^\sigma B_j^\sigma E \cup E \subset \Omega_i$ and by induction $[H_i^\sigma]^k E \subset \Omega_i$ for all k . But by (1) $[H_i^\sigma]^k E = \Omega$, so $\Omega_1 = \Omega_2 = \Omega$. Q.E.D.

This result should be viewed in the light of Lemma 4.1, namely that there can be at most one p -dominant action pair if $p \leq 1/2$. Since the belief potential cannot exceed $1/2$, Theorem 5.1 has an impact when some action pair is p -dominant for a small p . The smaller p is, the greater will be the impact of Theorem 5.1.

Theorem 5.1 is also tight in the following sense. Suppose an information system has belief potential σ . Then for any $p > \sigma$, it is possible to specify payoffs such that some action pair (a_1, a_2) is p -dominant at every state, some player i at some state ω^* knows the event “ a_i is a strongly dominant action,” but there is a pure-strategy Bayesian Nash equilibrium where a_i is not played everywhere. Let us give the construction, since it helps illustrate the significance of the belief potential. Choose α such that $p > \alpha > \sigma$; then (by clause (iii) of Lemma 3.1), there exists for some player (say, 1), an event $E_1 \in \mathcal{F}_1$, $E_1 \neq \emptyset$ or Ω , such that $B_1^\alpha B_2^\alpha E_1 \subset E_1$. Now suppose that at states in E_1 , payoffs are

		Player 2	
		L	R
Player 1	L	$1 - \alpha, 1 - \alpha$	$1 - \alpha, 0$
	R	$0, 0$	$0, \alpha$

while at states in $\sim E_1$, payoffs are

		Player 2	
		L	R
Player 1	L	$1 - \alpha, 1 - \alpha$	$0, 0$
	R	$0, 0$	α, α

Note that (L, L) is p -dominant at every state and L is a dominant action for 1 on event E_1 . This game has a pure strategy Bayesian Nash equilibrium where $s_1(\omega) = L$ if $\omega \in E_1$, $s_1(\omega) = R$ if $\omega \in \sim E_1$; $s_2(\omega) = L$ if $\omega \in B_2^\alpha E_1$ and $s_2(\omega) = R$ if $\omega \in \sim B_2^\alpha E_1$. Player 2's strategy is optimal, since he will play L exactly when he assigns probability greater than α to player 1 choosing L , i.e., on event $B_2^\alpha E_1$. At all states in $\sim E_1$, 1 assigns probability strictly less than α to 2 playing L (since $B_1^\alpha B_2^\alpha E_1 \subset E_1$), and so plays R .

Theorem 5.1 used the infection argument and the existence of a strongly dominant action to give a precise prediction of play in a game of incomplete

information. The corollary which follows makes no assumption about a strongly dominant action somewhere on the state space. Therefore, we are not able to predict uniquely the outcome. Still, we can use the infection argument to make a *conditional* claim.

COROLLARY: *Suppose that (1) the information system has belief potential of σ ; (2) (a_1, a_2) is σ -dominant at every state. Suppose that $(s_1(\cdot), s_2(\cdot))$ 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 $s_i(\omega) = a_i$, for both i and all ω .*

In the game described in the Corollary (unlike in Theorem 5.1), iterated deletion of dominated strategies will typically have no bite. Thus it could be that there are two action pairs which are strict pure-strategy Nash equilibria at every state, in which case we know that there are two simple Bayesian Nash equilibria where each of these action pairs is played everywhere. Nonetheless, if one of these action pairs is σ -dominant (where σ is the belief potential), the conditional statement of the Corollary can be made in reference to it, but not in reference to the other equilibrium. However, if each player has only two possible actions, the corollary implies that in *any* pure strategy BNE, 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 σ -dominant action profile. But if one of the σ -dominant actions is ever taken, each player takes the σ -dominant action everywhere. Put differently, there are no equilibria where a player “alternates” between actions. This result need not apply, however, when there are more than two actions.

*Dept. of Economics, University of Pennsylvania, Philadelphia, PA 19104, U.S.A.,
Dept. of Economics, University of Pennsylvania, Philadelphia, PA 19104, U.S.A.,
and
University College, University of Oxford, Oxford OX14BH, U.K.*

Manuscript received May, 1993; final revision received August, 1994.

APPENDIX

In the proofs, we appeal to the following properties of belief operators, which are stated without proof:

- (0) $E \subset F$ implies $B_i^p E \subset B_i^p F$.
- (1) If $p \geq q$, then $B_i^p E \subset B_i^q E$.
- (2) If $\omega \in B_i^q E$ for all $q < p$, then $\omega \in B_i^p E$.
- (3) If $E \in \mathcal{F}_i$, then $B_i^p[E \cup F] = E \cup B_i^p F$.
- (4) For a given information system and any $p > 0$, there exists $q > 1 - p$ such that $B_i^q E \subset E \Rightarrow \sim E \subset B_i^p \sim E$.
- (5) If $q > 1 - p$, then $E \subset B_i^p E \Rightarrow B_i^q \sim E \subset \sim E$.
- (6) If there are two individuals and $\mathcal{P}_i \neq \{\Omega\}$ for some i , then there exists an event E with $E \subset B_j^{0.5} E$, for all j , and $B_i^{0.5} E \notin \{\emptyset, \Omega\}$ for some i .

PROOF OF LEMMA 3.1: [(iii) \Rightarrow (i)] Suppose $B_i^p B_j^p E_i \subset E_i$ for some $j \neq i$, $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$. Then $H_i^p E_i = E_i$ for some $i \in \{1, 2\}$, $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$

$$\begin{aligned} &\Rightarrow [H_i^p]^k E_i \neq \Omega \text{ for all } k \geq 1, \text{ for some } i \in \{1, 2\}, E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\} \\ &\Rightarrow \sigma_i(E_i) < p \text{ for some } i \in \{1, 2\}, E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\} \\ &\Rightarrow \sigma < p. \end{aligned}$$

[(i) \Rightarrow (iii)] Conversely, suppose there does not exist $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$ such that $B_i^p B_j^p E_i \subset E_i$, where $j \neq i$. Then H_i^p is strictly increasing for all i , $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$

$$\begin{aligned} &\Rightarrow [H_i^p]^k E_i = \Omega \text{ for some } k \geq 1, \text{ for all } i, E_i \in \mathcal{F}_i \setminus \{\emptyset\} \\ &\Rightarrow \sigma_i(E_i) \geq p \text{ for all } i, E_i \in \mathcal{F}_i \setminus \emptyset \\ &\Rightarrow \sigma \geq p. \end{aligned}$$

[(ii) \Rightarrow (iii)] Suppose $B_j^p E \subset E$ for all j and $B_i^p E \not\subset \{\emptyset, \Omega\}$ for some i . Let $E_i = B_i^p E \not\subset \{\emptyset, \Omega\}$; now for $j \neq i$ (using fact 0), $B_j^p B_i^p E_i = B_j^p B_i^p B_i^p E \subset B_j^p B_i^p E \subset B_i^p E = E_i$.

[(iii) \Rightarrow (ii)] Suppose $B_i^p B_j^p E_i \subset E_i$ for some $j \neq i$, $E_i \in \mathcal{F}_i \setminus \{\emptyset, \Omega\}$. Now let $E = E_i \cup B_j^p E_i$. Then (by fact 3 above) $B_j^p E = E_i \cup B_j^p B_j^p E_i = E_i \subset E$; $B_i^p E = B_i^p E_i \subset E$; and $B_i^p E = E_i \not\subset \{\emptyset, \Omega\}$. Thus E satisfies the property of clause (ii) of the Lemma.

PROOF OF THEOREM 3.1: Suppose $E \subset B_j^p E$ for all j and $B_i^p E \not\subset \{\emptyset, \Omega\}$ for some i . Then (1) (by fact 5) $B_j^q E \subset E$ for all $q > 1 - p$, for all j ; (2) $B_i^q E \not\subset E$, for all $q > 1 - p$ (since $q > 1 - p$ and $B_j^q E \subset E = \Omega \Rightarrow B_i^q E = \emptyset$, a contradiction); (3) $B_i^q E \not\subset E \neq \emptyset$, for all $q \in (1 - p, q')$, for some $q' > 1 - p$ (since $B_j^q E \subset E = \emptyset$ for all $q > 1 - p \Rightarrow B_j^q E = \Omega$, a contradiction). Now by Lemma 3.1 we have $q > \sigma$ for all $q \in (1 - p, q')$, i.e. $\sigma \leq 1 - p$.

Suppose $p < 1 - \sigma$. Choose some $q > 1 - p$ such that fact 4 is satisfied. Since $\sigma < 1 - p < q$, there is (by Lemma 3.1) an event E with $B_j^q E \subset E$ for all j and $B_i^q E \not\subset \{\emptyset, \Omega\}$ for some i . Now (1) (by fact 4) $\sim E \subset B_j^p \sim E$, for all j ; (2) $B_i^p \sim E \neq \emptyset$ (since $B_i^p \sim E = \emptyset \Rightarrow \sim E = \emptyset \Rightarrow E = \Omega \Rightarrow B_i^q E = \Omega$, a contradiction); (3) $B_i^p E \neq \Omega$ (since $q > 1 - p$ and $B_i^p \sim E = \Omega \Rightarrow B_i^q E = \emptyset$, a contradiction). So there is a nontrivial p -evident event whenever $p < 1 - \sigma$. By fact (2), there is also a nontrivial p -evident event when $p = 1 - \sigma$.

PROOF OF THEOREM 3.2: By fact 6 above, there exists a nontrivial event which is $1/2$ -evident. Thus by Theorem 3.1, $\sigma \leq 1 - 1/2 = 1/2$.

PROOF OF LEMMA 4.1: (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:

$$\begin{aligned} \frac{1}{2}g_i(a_i, a_j; \omega) + \frac{1}{2}g_i(a_i, b_j; \omega) &> \frac{1}{2}g_i(b_i, a_j; \omega) + \frac{1}{2}g_i(b_i, b_j; \omega), \\ \frac{1}{2}g_i(b_i, a_j; \omega) + \frac{1}{2}g_i(b_i, b_j; \omega) &> \frac{1}{2}g_i(a_i, a_j; \omega) + \frac{1}{2}g_i(a_i, b_j; \omega). \end{aligned}$$

(ii) If there is a pure strategy Nash equilibrium, then generically there is at least one strict Nash equilibrium. Generically, every strict Nash equilibrium will be p -dominant for a different set of p . 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 enlarge the set of p for which (a_1, a_2) is p -dominant and weakly reduce the set of p for which every other strict Nash equilibrium pair is p -dominant.

REFERENCES

- CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989–1018.
 GEANAKOPOLOS, J. (1992): "Common Knowledge," *Journal of Economic 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., AND R. SELTEN (1988): "A General Theory of Equilibrium in Games," Cambridge: MIT Press.

- KANDORI, M., G. MAILATH, AND R. ROB (1993): "Learning, Mutation and Long-run Equilibria in Games," *Econometrica*, 61, 29–56.
- MONDERER, D., AND D. SAMET (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170–190.
- (1990): "Proximity of Information in Games with Incomplete Information," Israel Institute of Business Research, Working Paper #63/90.
- MORRIS, S., R. ROB, AND H. SHIN (1993): "Risk Dominance and Stochastic Potential," CARESS Working Paper #93-11, University of Pennsylvania.
- RUBINSTEIN, A. (1989): "The Electronic Mail Game: Strategic Behavior under almost Common Knowledge," *American Economic Review*, 79, 385–391.
- SHIN, H., AND T. WILLIAMSON (1992): "How Much Common Belief Is Necessary for a Convention?" CARESS Working Paper #92-32, University of Pennsylvania.
- SORIN, S. (1993): "On the Impact of an Event," Mimeo, Ecole Normale Supérieure.
- STINCHCOMBE, M. (1989): "Approximating Common Knowledge," Mimeo, University of California, San Diego.
- YOUNG, P. (1993): "The Evolution of Conventions," *Econometrica*, 61, 57–84.