

Global Games and Equilibrium Selection

Author(s): Hans Carlsson and Eric van Damme

Source: *Econometrica*, Vol. 61, No. 5 (Sep., 1993), pp. 989-1018

Published by: [The Econometric Society](#)

Stable URL: <http://www.jstor.org/stable/2951491>

Accessed: 07/12/2010 12:58

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=econosoc>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

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*.

GLOBAL GAMES AND EQUILIBRIUM SELECTION¹

BY HANS CARLSSON AND ERIC VAN DAMME

A global game is an incomplete information game where the actual payoff structure is determined by a random draw from a given class of games and where each player makes a noisy observation of the selected game. For 2×2 games, it is shown that, when the noise vanishes, iterated elimination of dominated strategies in the global game forces the players to conform to Harsanyi and Selten's risk dominance criterion.

KEYWORDS: Global games, equilibrium selection, risk dominance, payoff uncertainty, incomplete information, common knowledge, iterated dominance.

1. INTRODUCTION

ANY MODEL, BY ITS NATURE, is based on assumptions that schematize and simplify the phenomena under investigation. The basic assumption underlying main-stream game-theoretic models is that the rules of the game, including its payoff structure and the rationality of the players, are common knowledge. There seems to be almost general agreement that game theory's agents are excessively rational and well-informed in comparison with their real-life counterparts. One way of assessing the role of this kind of assumption is to compare the model with perturbed variants that are based on slightly modified assumptions. In this manner, Harsanyi's (1973) games with randomly disturbed payoffs and Selten's (1975) concept of trembling-hand perfection perturb certain aspects of game theory's information and rationality assumptions. Analyzing such richer models may yield considerable benefits: Harsanyi's approach produces a plausible justification and interpretation of mixed strategy equilibria while Selten's approach frequently leads to a drastic reduction in the number of possible solutions.

The present paper pursues this line of research by analyzing an incomplete information model—to be called a *global game*—which is based on a perturbation of the players' payoff information in 2×2 games. The game to be played is determined by a random draw from some subclass of all 2×2 games. Each player observes the selected game with some noise and then chooses one of his two available actions. If the initial subclass of games is large enough and contains games with different equilibrium structures, iterated elimination of dominated strategies in the incomplete information game yields a surprising result: When the 2×2 game actually selected by Nature is one with two strict Nash equilibria, iterated dominance forces the players to coordinate on the equilibrium which is risk-dominant in the sense of Harsanyi and Selten (1988),

¹This paper is a combination and substantial generalization of Carlsson (1989) and Carlsson and van Damme (1989). Some basic ideas on global games and their relation to risk dominance originate from a note written by Carlsson in 1988. The authors thank Reinhard Selten, Lars-Gunnar Svensson, Jörgen Weibull, and various seminar audiences for helpful comments. The constructive criticism from an editor and several referees considerably improved the paper's quality. Carlsson gratefully acknowledges financial support from the Swedish Council for Research in the Humanities and Social Sciences and the Jan Wallander Foundation.

provided that the amount of noise in the players' observations is sufficiently small.

The result is surprising because strict equilibria have proved to be immune against most perturbations and refinement criteria; for instance, any strict equilibrium survives strategic stability à la Kohlberg and Mertens (1986). More specifically, our result goes against Harsanyi's (1973) conclusion that, generically, any Nash equilibrium in a normal form game can be approximated by equilibria of games with randomly disturbed payoffs. The explanation of this apparent contradiction is that we use a quite different payoff perturbation: In Harsanyi's set-up the payoff matrices of different players are independent and each player only learns his own matrix. Hence, the players' observations are uncorrelated so a player's beliefs do not depend on his information but are common knowledge. By contrast, in our model, both players make noisy observations of the entire game. As a result, different players' observations are correlated and a player's (first and higher order) beliefs depend on his observation. For a more extensive comparison between the two models, we refer to Section 7.2, where we also briefly raise the question regarding their relative appropriateness, and to Appendix B where we construct a hybrid model that contains our model and Harsanyi's as special cases.

Our result is driven by the fact that, in a global game, the uncertainty forces the players to take account of the entire class of a priori possible games which may be large even if the amount of noise is small. When the class contains games with different equilibrium structures, players have to switch actions at some points in their observation spaces. Equilibrium selection in accordance with risk dominance results from the conditions which optimally chosen switching points must satisfy in the limit, as noise vanishes. Hence, there is a certain parallel between our approach and the axiomatic models which, starting with Nash's (1950, 1953) bargaining analyses, derive determinate solutions to individual games by imposing consistency requirements on classes of games. An important difference, though, is that in the present model the consistency conditions—instead of simply being postulated—are derived from more basic assumptions about the players' information in a full-fledged noncooperative game. This feature also distinguishes our approach from existing equilibrium selection theories—such as Harsanyi and Selten's (1988) tracing procedure—where the solution to a game typically results from the application of some more or less ad hoc scheme of expectation formation. Hence our approach can be viewed as an attempt to extend Nash's (1951) program of providing noncooperative foundations for axiomatic solution concepts.

The various links between this paper and the above cited works will be given a more extensive discussion later. At present we will study a simple example which catches the basic intuition for our result. This will be a suitable context for providing an explicit definition of the notion of risk dominance for 2×2 games.

Let $g(x)$ be the 2×2 game shown in Figure 1, and consider the class of games $\{g(x)\}_{x \in \mathbb{R}}$. Note that $g(x)$ has strict dominance solution $\alpha = (\alpha_1, \alpha_2)$ for

	α_2	β_2
α_1	x, x	$x, 0$
β_1	$0, x$	$4, 4$

FIGURE 1.—Game $g(x)$.

$x > 4$ and $\beta = (\beta_1, \beta_2)$ for $x < 0$. If $x \in (0, 4)$, both α and β are strict Nash equilibria. Then, according to Harsanyi and Selten (1988), α risk-dominates β if α is associated with the largest product of deviation losses and vice versa. These products are given by x^2 for α and $(4 - x)^2$ for β , so α is risk-dominant in $g(x)$ if $x \in (2, 4)$ while β is risk-dominant for $x \in (0, 2)$. To see the intuition for the risk dominance criterion, note that each player will choose α or β according as he believes that the probability of the opponent playing β is smaller or larger than $x/4$. Hence, if $x < 2$, playing β is less risky than playing α in the sense that, for each player, β is justified by a larger range of beliefs about the opponent's action.

For $g(x)$ with $x \in (2, 4)$, playing α can be motivated by an analogous argument. In these games, however, there is a conflict between risk dominance and payoff dominance: Although α is risk-dominant both players prefer the outcome $(4, 4)$ associated with β . Many people feel it is obvious that β should be selected in this case. Harsanyi and Selten, too, consider that payoff dominance should be given precedence over risk dominance. We will now show how, in our approach, rational players are *forced* to pick the risk-dominant equilibrium even when the other equilibrium is Pareto-preferred.

Consider the incomplete information game where x is the realization of a random variable X which is uniform on some interval $[\underline{x}, \bar{x}]$ such that $\underline{x} < 0$ and $\bar{x} > 4$; given x , each player i makes an observation corresponding to a random variable X_i which is uniform on $[x - \varepsilon, x + \varepsilon]$ for some $\varepsilon > 0$, the two players' observation errors $X_1 - x$ and $X_2 - x$ being independent; having made their observations, the players choose actions simultaneously and get payoffs corresponding to $g(x)$. It is understood that the structure of the class of games and the joint distribution of X , X_1 and X_2 are common knowledge.

It is easily seen that player i 's posterior of X will be uniform on $[x_i - \varepsilon, x_i + \varepsilon]$ if he observes $x_i \in [\underline{x} + \varepsilon, \bar{x} - \varepsilon]$, so his conditionally expected payoff from choosing α_i will simply be x_i . Moreover, for $x_i \in [\underline{x} + \varepsilon, \bar{x} - \varepsilon]$, the conditional distribution of the opponent's observation X_j will be symmetric around x_i and have support $[x_i - 2\varepsilon, x_i + 2\varepsilon]$. Hence, $\text{Prob}\{X_j < x_i | x_i\} = \text{Prob}\{X_j > x_i | x_i\} = 1/2$.

Now assume $\varepsilon < -\underline{x}/2$ and suppose i observes $x_i < 0$. Then his conditionally expected payoff from choosing α_i is certainly negative and, thus, smaller than the payoff associated with β_i . Hence β_i is conditionally (strictly) dominant for i when he observes $x_i < 0$. It should be clear that iterated dominance arguments allow us to get further. For instance, if player 2 is restricted to playing β_2 for observations $x_2 < 0$, then player 1, observing $x_1 = 0$, must assign at least

probability $1/2$ to β_2 . Consequently, 1's conditionally expected payoff from playing β_1 will be at least 2, so α_1 (which yields 0) can be excluded by iterated dominance for $x_1 = 0$. Let x_i^* be the smallest observation for which β_i cannot be established by iterated dominance. By symmetry, obviously, $x_1^* = x_2^* = x^*$. Iterated dominance requires player i to play β_i for any $x_i < x^*$, so if j observes x^* he will assign at least probability $1/2$ to i 's choosing β_i and, thus, j 's expected payoff from choosing β_j will be at least 2. Since j 's expected payoff from choosing α_j equals x^* , we must have $x^* \geq 2$, for otherwise iterated dominance would require j to play β_j when he observes x^* .

Proceeding in the same way for large values of the observations, we see that, for $\varepsilon < (\bar{x} - 4)/2$, α_i is dominant for each player if $x_i > 4$. Letting x^{**} denote the lower bound on the iterated dominance region for α , we find that, if j observes x^{**} , his expected payoff from choosing β_j will be at most 2 given that i conforms to iterated dominance. Since x^{**} equals j 's expected payoff from choosing α_j , we conclude that $x^{**} \leq 2$. Combining this with the above and the obvious fact that $x^* \leq x^{**}$, we get

$$x^* = x^{**} = 2.$$

In words, iterated dominance in the global game forces equilibrium selection in accordance with risk dominance in the sense that each player i should play the risk-dominant equilibrium of $g(x_i)$ for any $x_i \in (0, 2) \cup (2, 4)$. As the reader may have noticed, this does not imply that players always coordinate on the risk-dominant equilibrium of the actual game $g(x)$. However, such coordination will be ensured if, for given x , one requires ε to be small enough (pick $\varepsilon < |x - 2|$).

The main purpose of the paper is to show that the above result holds in a fairly general way for 2×2 games. In particular the result will be shown not to depend on the symmetries exploited in the example nor on the exact shapes of the probability distributions nor on the underlying class of games being one-dimensional. On the other hand, from the previous argument it is clear that we do need some assumption on the class of games. The result depends critically on the existence—within this class—of a subclass of dominance solvable games that serve as take-offs for the iterated dominance argument, and, thus, exert a kind of remote influence on the games with multiple equilibria.

To illustrate the last point and to get a more detailed intuition for our result, let us assume that $x = 3$ so that Nature has selected the game $g(3)$ where β is Pareto-dominant but α is risk-dominant. Assuming that ε is small, say $\varepsilon < 1/4$, both players will know these facts about α and β from their observations $x_1, x_2 \in [3 - \varepsilon, 3 + \varepsilon]$. Each player will even know that his opponent knows that β is a Pareto-dominant equilibrium (observing x_i , i knows that the conditional support of X given x_j belongs to $[x_i - 3\varepsilon, x_i + 3\varepsilon]$). Why, then, cannot the players coordinate on β ? The reason is that their choices at x_1 and x_2 have to be part of a consistent plan for all possible observations. By the strict dominance of α_i for $x_i > 4$, each player would have to switch from α_i to β_i at some observation $x_i^* \in [3, 4]$. However, given any switching point x_j^* for the oppo-

ment, each player would like to switch slightly below x_j^* as long as $x_j^* > 2$. Hence the result.

The key to understanding the phenomenon is the realization that there is a sharp separation between *knowledge* and *common knowledge* in a global game. Even though, when their observations are close to 3, the players know that β is the Pareto-dominant equilibrium and even though, for small ε , each knows that the other knows this, this information is not common knowledge for any positive ε . As a matter of fact, the only information that is common knowledge is that some game $g(x)$ with $x \in [\underline{x}, \bar{x}]$ has to be played. This lack of common knowledge enables remote areas (viz. those where α or β is the dominance solution) to influence the behavior when the observations are close to 3. This phenomenon also supplies a deeper motivation for using the term *global games*: When determining rational behavior it is not enough to look at the equilibrium structure that is known to prevail; it must also be ensured that the chosen action is part of a consistent plan for all situations that could have occurred within the underlying class of games.

The remainder of the paper is organized as follows: Section 2 introduces the notation that will be used for 2×2 games and provides a convenient alternative characterization of the risk dominance relation for such games. In Section 3 we formally define global games and state our main result, which is proved in Section 4. Section 5 discusses the role of our assumptions, while Section 6 discusses the result, in particular its relationship to Harsanyi and Selten's justification of risk dominance and to Nash's works. Section 7 deals with related literature on games with payoff uncertainty and Section 8 concludes.

2. 2×2 GAMES AND RISK DOMINANCE

This section introduces notation and definitions to be used in the subsequent analysis. A 2×2 game is a two-person normal form game where each player $i \in \{1, 2\}$ can choose between two pure actions α_i and β_i . Let G be the set of all 2×2 games. Clearly, since such a game is fully described by its eight payoff entries, G can be identified with \mathbb{R}^8 . For $g \in G$, a mixed action for player i is a probability distribution s_i on $\{\alpha_i, \beta_i\}$. We identify s_i with the probability which s_i assigns to α_i . If $s = (s_1, s_2)$ is an action pair, we write $g_i(s)$ for player i 's expected payoff when s is played in g . Moreover, we write g_i^α (resp. g_i^β) for the loss that i incurs by deviating unilaterally from the action pair $\alpha \equiv (\alpha_1, \alpha_2)$ (resp. $\beta \equiv (\beta_1, \beta_2)$). Hence for $i, j \in \{1, 2\}, i \neq j$,

$$g_i^\alpha = g_i(\alpha) - g_i(\beta_i, \alpha_j),$$

$$g_i^\beta = g_i(\beta) - g_i(\alpha_i, \beta_j).$$

For $\gamma \in \{\alpha, \beta\}$, G^γ is the set of games where γ is a strict Nash equilibrium:

$$G^\gamma = \{g \in G : g_i^\gamma > 0 \text{ for } i = 1, 2\},$$

D_i^γ denotes the subset of G^γ where γ_i is a strictly dominant action for player i :

$$D_i^\alpha = \{g \in G^\alpha : g_i^\beta < 0\},$$

$$D_i^\beta = \{g \in G^\beta : g_i^\alpha < 0\},$$

and D^γ denotes the set of games for which γ is the iterated dominance solution:

$$D^\gamma = D_1^\gamma \cup D_2^\gamma.$$

Generically, games in G have one or three equilibria. In the latter case, there is one (weak) equilibrium in mixed strategies as well as two (strict) pure strategy equilibria. We may, without loss of generality, focus on the case where both α and β are strict equilibria. For $g \in G^\alpha \cap G^\beta$, α is said to risk-dominate β if the product of deviation losses is larger at α than at β , i.e. if

$$(2.1) \quad g_1^\alpha g_2^\alpha > g_1^\beta g_2^\beta,$$

while β risk-dominates α if the reverse inequality holds.

The basic intuition behind risk dominance stands out more clearly in the following alternative characterization, which we will use in the sequel. Let $V_i(s_j)$ denote player i 's net gain from playing α_i rather than β_i when j plays α_j with probability s_j :

$$V_i(s_j) = s_j g_i^\alpha - (1 - s_j) g_i^\beta.$$

For $g \in G^\alpha \cap G^\beta$, let \bar{s}_j denote the unique solution to $V_i(s_j) = 0$, i.e. the probability which j should attach to α_j to make i indifferent between his two actions:

$$(2.2) \quad \bar{s}_j = \frac{g_i^\beta}{g_i^\alpha + g_i^\beta}.$$

(Note that (\bar{s}_1, \bar{s}_2) is the mixed strategy equilibrium of g .) Clearly, i should choose α_i or β_i according as he considers the actual s_j to be larger or smaller than \bar{s}_j . Hence a small \bar{s}_j implies a large range of beliefs which justify choosing α_i . This provides a rationale for using risk dominance as an equilibrium selection criterion, for (2.1) is equivalent to

$$(2.3) \quad \bar{s}_1 + \bar{s}_2 < 1.$$

We write R^α (resp. R^β) for the set of games where α (resp. β) either is the unique strict equilibrium or risk-dominates β (resp. α):

$$R^\alpha = \{g \in G^\alpha : \text{if } g_1^\beta, g_2^\beta > 0, \text{ then } \bar{s}_1 + \bar{s}_2 < 1\},$$

$$R^\beta = \{g \in G^\beta : \text{if } g_1^\alpha, g_2^\alpha > 0, \text{ then } \bar{s}_1 + \bar{s}_2 > 1\}.$$

It is convenient to introduce the notation G_i^+ for the set of games for which the denominator in (2.2) is positive and to extend the definition of \bar{s}_j from $G^\alpha \cap G^\beta$ to G_i^+ . Note that $D_i^\gamma \subseteq R^\gamma$ and that the sets D_i^γ , R^γ , and G_i^+ are open

in \mathbb{R}^8 ($i \in \{1, 2\}$, $\gamma \in \{\alpha, \beta\}$) and that

$$(2.4) \quad R^\alpha \cap G_1^+ \cap G_2^+ = \{g \in G^\alpha \cap G_1^+ \cap G_2^+ : \bar{s}_1 + \bar{s}_2 < 1\}.$$

3. GLOBAL GAMES AND ITERATED DOMINANCE

We now picture the players in a situation where it is common knowledge that some game in G will be played but the players do not know which one. Initially, they have common prior beliefs represented by a probability distribution with support on some subclass of G . However, before choosing his action, each player gets additional (private) information in the form of a fuzzy observation of the actual game to be played. The resulting incomplete information game—to be called a *global game*—may thus be described by the following steps:

1. Nature selects a game from G .
2. Each player observes g with some noise.
3. Players choose actions simultaneously.
4. Payoffs are determined by g and the players' choices.

Since G may be identified with \mathbb{R}^8 , a simple way of modeling the above situation would be to let the players make observations directly in \mathbb{R}^8 ; each player would then observe the payoffs of the actual game plus some error terms. Here, however, we will use an alternative formulation where the selected game is observed indirectly through some parameter space which is mapped on G . Therefore, let Θ be a parameter space such that to any x in Θ corresponds a game $g(x)$ in G . We assume:

(A1) Θ is an open subset of \mathbb{R}^m for some m , $g: \Theta \rightarrow G$ is continuously differentiable and the partial derivatives $\partial g^k / \partial x^l$ ($k = 1, \dots, 8$; $l = 1, \dots, m$) are bounded on Θ .

The advantage with this parametric formulation is that it allows a more flexible modeling of various classes of games. For instance, it makes it possible to represent the class of games corresponding to Figure 1 as a one-dimensional space. The case where players observe payoffs directly may be modeled by letting $\Theta \subseteq \mathbb{R}^8$ and $g(x) = x$ for all $x \in \Theta$.

We let the players' prior be described by a random variable X that takes values in Θ . Moreover, Player i 's observation is described by a random variable X_i^ε which is defined by

$$X_i^\varepsilon = X + \varepsilon \cdot E_i, \quad i = 1, 2,$$

where E_i is a random variable that takes values in \mathbb{R}^m and $\varepsilon > 0$ is a scale parameter. In what follows, Θ , g , X , E_1 , and E_2 will remain fixed and we write Γ^ε for the global game parameterized by ε . We are particularly interested in the case where the observations are almost correct so we will focus on Γ^ε with ε

close to zero. We make the following assumptions:

(A2) X admits a density h which is strictly positive, continuously differentiable, and bounded on Θ .

(A3) The vector (E_1, E_2) is independent of X and admits a continuous density φ .

(A4) The support of each E_i is contained in a ball with radius 1 around 0 in \mathbb{R}^m .

A strategy for i is a measurable function s_i that assigns a probability mixture on $\{\alpha_i, \beta_i\}$ to each observation x_i of player i in Γ^ε . We let Θ_i^ε be the set of all possible observations for i , S_i^ε denotes the set of strategies for i , and we write $s_i(x_i)$ for the probability which s_i assigns to the action α_i at x_i .

Let $V_i^\varepsilon(s_j|x_i)$ denote the expected net gain for player i from playing α_i rather than β_i when he observes x_i and player j uses strategy s_j in Γ^ε . If $V_i^\varepsilon(s_j|x_i) > 0$ (resp. $V_i^\varepsilon(s_j|x_i) < 0$) for all $s_j \in S_j^\varepsilon$, then α_i (resp. β_i) is (conditionally) dominant at x_i . Hence, we may define the following process of iterated elimination of strategies which prescribe dominated actions for some observations:

$$\begin{aligned}
 S_i^{\varepsilon,0} &= S_i^\varepsilon, & A_i^{\varepsilon,0} &= \phi, & B_i^{\varepsilon,0} &= \phi, \\
 S_i^{\varepsilon,n} &= \{s_i \in S_i^\varepsilon : s_i(x_i) = 1 \text{ if } x_i \in A_i^{\varepsilon,n} \text{ and } s_i(x_i) = 0 \text{ if } x_i \in B_i^{\varepsilon,n}\}, \\
 A_i^{\varepsilon,n+1} &= \{x_i \in \Theta_i^\varepsilon : V_i^\varepsilon(s_j|x_i) > 0 \text{ for all } s_j \in S_j^{\varepsilon,n}\}, \\
 B_i^{\varepsilon,n+1} &= \{x_i \in \Theta_i^\varepsilon : V_i^\varepsilon(s_j|x_i) < 0 \text{ for all } s_j \in S_j^{\varepsilon,n}\}.
 \end{aligned}$$

The following properties, which hold for all i and n , are immediate from the definitions:

$$\begin{aligned}
 A_i^{\varepsilon,n} \cap B_i^{\varepsilon,n} &= \phi, & S_i^{\varepsilon,n} &\neq \phi. \\
 A_i^{\varepsilon,n+1} \supseteq A_i^{\varepsilon,n}, & & B_i^{\varepsilon,n+1} \supseteq B_i^{\varepsilon,n}, & & S_i^{\varepsilon,n+1} \subseteq S_i^{\varepsilon,n}.
 \end{aligned}$$

Let us define

$$A_i^\varepsilon = \bigcup_{n=0}^\infty A_i^{\varepsilon,n}, \quad B_i^\varepsilon = \bigcup_{n=0}^\infty B_i^{\varepsilon,n}.$$

We say that α (resp. β) is iteratively dominant at x in Γ^ε if $x \in A_1^\varepsilon \cap A_2^\varepsilon$ (resp. $x \in B_1^\varepsilon \cap B_2^\varepsilon$). The main result of the paper is the following theorem.

THEOREM: *Let $\gamma \in \{\alpha, \beta\}$. If x lies on a continuous curve C such that $C \subseteq \Theta$, $g(C) \subseteq R^\gamma$, and $g(C) \cap D^\gamma \neq \phi$, then γ is iteratively dominant at x in Γ^ε if ε is sufficiently small.*

For the nontrivial case where $g(x)$ has two strict equilibria, the Theorem says that iterated dominance in the global game forces each player to select the risk-dominant equilibrium of the game corresponding to his observation provided that ε is sufficiently small. The next section provides the essential steps of

the proof of the Theorem, but some technical details will be deferred to the Appendix. The role of the assumptions underlying the Theorem will be discussed in Section 5.

4. PROOF OF THE THEOREM

The Theorem consists of two completely symmetric statements according as γ equals α or β . To economize on notation, we will formulate lemmas and proofs only for the former case. We first provide an outline of the main steps of the proof.

4.1. Outline

The proof consists of three parts: (i) an investigation of the beliefs of the players conditional on their observations, (ii) derivation of properties of the sets A_i^ε , and (iii) the main part which involves an argument by contradiction.

(i) We first show that the posterior beliefs are approximately symmetric in the sense that, if ε is small, the likelihood that player 1 assigns to player 2 making an observation x_2 when his own observation is x_1 is approximately equal to the likelihood that 2 assigns to 1 observing x_1 when his information is x_2 . This symmetry holds exactly if the prior is uniform and, for general priors, it holds approximately if the observation errors are small since in that case the prior is almost constant throughout the region of possible realizations. This property implies that the players have almost the same beliefs about the difference of their observations, which for the one-dimensional case entails that

$$(4.1) \quad F_1^\varepsilon(x_2|x_1) + F_2^\varepsilon(x_1|x_2) \approx 1$$

if ε is small. ($F_i^\varepsilon(x_j|x_i) = P(X_j^\varepsilon \leq x_j | X_i^\varepsilon = x_i)$ is the probability that i assigns to j making an observation below x_j when he observes x_i .) A corresponding property obtains in the multi-dimensional case.

(ii) In the second step we show that when α is the iterated dominance solution of the unperturbed game g (i.e. $g \in D^\alpha$), α remains iteratively dominant when the players make noisy observations that are close to g . It is also shown that an observation x_i belongs to A_i^ε if and only if at x_i , playing α_i is better than playing β_i whenever j plays α_j exactly on A_j^ε . If we denote the latter strategy by a_j^ε , then, by continuity, $V_i^\varepsilon(a_j^\varepsilon|x_i^\varepsilon) = 0$ for a boundary point x_i^ε of A_i^ε . It will be clear that at such a boundary point x_i^ε , if ε is small, the probability $a_j^\varepsilon(x_i^\varepsilon)^2$ that player i assigns to j playing α_j must be close to the probability $\bar{v}_j(x_i^\varepsilon)$ that makes i indifferent between his two actions in $g(x_i^\varepsilon)$. Finally we prove a Lemma which will help to show that the boundary points that are used in step (iii) must be close to each other.

² The reader should note the distinction between $a_j^\varepsilon(x_j)$, the probability which the strategy a_j^ε assigns to action α_j at the observation x_j on the one hand, and, on the other, $a_j^\varepsilon(x_i)$, $i \neq j$, which denotes the probability which player i will assign to j 's choosing α_j when i observes x_i and j uses strategy a_j^ε .

(iii) In the heart of the proof we argue by contradiction: Assume x' lies on a curve C with $g(C) \subseteq R^\alpha$ and $g(C) \cap D^\alpha \neq \emptyset$, but nevertheless $x' \notin A_1^\varepsilon \cap A_2^\varepsilon$ for ε small. The crucial step of the proof consists in the construction of boundary points x_1^ε and x_2^ε of A_1^ε and A_2^ε , respectively, that are close to each other and for which

$$(4.2) \quad \lim_{\varepsilon \rightarrow 0} \bar{s}_2(x_1^\varepsilon) + \bar{s}_1(x_2^\varepsilon) < 1 \quad \text{and}$$

$$(4.3) \quad \lim_{\varepsilon \rightarrow 0} a_2^\varepsilon(x_1^\varepsilon) + a_1^\varepsilon(x_2^\varepsilon) \geq 1,$$

conditions that cannot be satisfied simultaneously since, by step (ii) above, $a_j^\varepsilon(x_i^\varepsilon) \approx \bar{s}_j(x_i^\varepsilon)$ if ε is small. Now (4.2) will be satisfied since x_1^ε and x_2^ε have a common limit x^* with $g(x^*) \in R^\alpha$. To fulfil (4.3), the boundary points have to be carefully chosen and, in the general case, this requires a rather elaborate construction which is given in the Appendix. In the case where Θ is one-dimensional, however, the argument is straightforward: Assume C connects x' with a point \bar{x} in D^α with $\bar{x} < x'$ and let x_i^ε be the point closest to x' such that $(\bar{x}, x_i^\varepsilon) \subseteq A_i^\varepsilon$. Then $a_j^\varepsilon(x_i^\varepsilon) \geq F_i^\varepsilon(x_j^\varepsilon|x_i^\varepsilon)$ for ε small enough so (4.3) follows from (4.1).

4.2. Posterior Beliefs

We first introduce some notation associated with F^ε . Denote by $F_i^\varepsilon(x, x_j|x_i)$ the distribution function of (X, X_j^ε) conditional on $X_i^\varepsilon = x_i$ and let $f_i^\varepsilon(x, x_j|x_i)$ be the corresponding density. Similarly, let $F_i^\varepsilon(x_j|x_i)$ and $f_i^\varepsilon(x_j|x_i)$ be the conditional distribution and density of X_j^ε . If we write φ^ε for the joint density of $(\varepsilon E_i, \varepsilon E_j)$ (i.e. $\varphi^\varepsilon(z) = \varepsilon^{-2m} \varphi(\varepsilon^{-1}z)$), then, for $i, j = 1, 2, i \neq j$,

$$(4.4) \quad f_i^\varepsilon(x_j|x_i) = \frac{\int h(x) \varphi^\varepsilon(x_1 - x, x_2 - x) dx}{\int \int h(x) \varphi^\varepsilon(x_1 - x, x_2 - x) dx_j dx}.$$

Note that if X is uniformly distributed on Θ and if $B(x_i, \varepsilon)$, the ball with radius ε centered at x_i , is completely contained in Θ , then $h(x)$ is a constant for all x that yield a contribution to any of the integrals in (4.4). Thus $h(x)$ can be factored out from the fraction and, since the denominator of the new expression equals one, we get

$$f_i^\varepsilon(x_j|x_i) = \int \varphi^\varepsilon(x_1 - x, x_2 - x) dx = \psi^\varepsilon(x_1 - x_2),$$

where ψ^ε is the density of $\varepsilon E_1 - \varepsilon E_2$, the difference between the observation errors. Consequently, with a uniform prior, if $B(x_i, \varepsilon) \subseteq \Theta$ for $i = 1, 2$, then

$$f_1^\varepsilon(x_2|x_1) = f_2^\varepsilon(x_1|x_2) = \psi^\varepsilon(x_1 - x_2).$$

Hence

$$\begin{aligned} F_1^\varepsilon(x_2|x_1) &= \int_{y \leq x_2} \psi^\varepsilon(x_1 - y) dy \\ &= \int_{z \geq x_1} \psi^\varepsilon(z - x_2) dz \\ &= 1 - F_2^\varepsilon(x_1|x_2), \end{aligned}$$

which shows that (4.1) holds exactly in this case. The corresponding property in the multi-dimensional case is that, for $p \in \mathbb{R}^m$,

$$(4.5) \quad P(pX_2^\varepsilon \leq px_2 | X_1^\varepsilon = x_1) + P(pX_1^\varepsilon \leq px_1 | X_2^\varepsilon = x_2) = 1.$$

(Geometrically, the terms on the left-hand side denote the probabilities that X_j^ε lies in the negative halfspace $py \leq px_j$ conditional upon observing $X_i^\varepsilon = x_i$.) To see why (4.5) holds note that

$$\begin{aligned} P(pX_2^\varepsilon \leq px_2 | X_1^\varepsilon = x_1) &= \int_{py \leq px_2} \psi^\varepsilon(x_1 - y) dy \\ &= P(p\varepsilon E_1 - p\varepsilon E_2 \geq px_1 - px_2) \end{aligned}$$

and, similarly,

$$\begin{aligned} P(pX_1^\varepsilon \leq px_1 | X_2^\varepsilon = x_2) &= \int_{py \leq px_1} \psi^\varepsilon(y - x_2) dy \\ &= P(p\varepsilon E_1 - p\varepsilon E_2 \leq px_1 - px_2). \end{aligned}$$

Our first lemma shows that, for small ε , (4.5) remains approximately valid even without a uniform prior. The intuition is obvious: If ε is small, the prior is almost constant for all realizations that the players consider possible on the basis of their observations.

LEMMA 4.1: *Let $x_1, x_2 \in \Theta$. Then there exists a constant $k(x_1, x_2)$ such that for all $p \in \mathbb{R}^m$ and all sufficiently small ε*

$$(4.6) \quad |P(pX_2^\varepsilon \leq px_2 | X_1^\varepsilon = x_1) + P(pX_1^\varepsilon \leq px_1 | X_2^\varepsilon = x_2) - 1| \leq k(x_1, x_2)\varepsilon.$$

PROOF: Fix $\delta > 0$ such that $h(x) \geq \delta$ for all $x \in B(x_1, \varepsilon) \cup B(x_2, \varepsilon)$. Since h is continuously differentiable, there exists a constant $l(x_i)$ such that $|h(x) - h(x_i)| < l(x_i)\varepsilon$ for all $x \in B(x_i, \varepsilon)$. For any x_i we have

$$\frac{h(x_i) - l(x_i)\varepsilon}{h(x_i) + l(x_i)\varepsilon} \psi^\varepsilon(x_1 - x_2) \leq f_i^\varepsilon(x_j|x_i) \leq \frac{h(x_i) + l(x_i)\varepsilon}{h(x_i) - l(x_i)\varepsilon} \psi^\varepsilon(x_1 - x_2).$$

(Cf. the argument below (4.4).) Hence, if we choose ε so that $l(x_i)\varepsilon \leq \delta/2$ and define $k(x_i) = 4l(x_i)/\delta$, then

$$(4.7) \quad |f_i^\varepsilon(x_j|x_i) - \psi^\varepsilon(x_1 - x_2)| \leq k(x_i)\varepsilon\psi^\varepsilon(x_1 - x_2).$$

Substituting these bounds for $f_i^\varepsilon(x_j|x_i)$ into the expressions for $P(pX_j^\varepsilon \leq px_j | X_i^\varepsilon = x_i)$ and proceeding as in the case where X is uniform, we obtain (4.6) by defining $k(x_1, x_2) = k(x_1) + k(x_2)$. Q.E.D.

4.3. Properties of the Set A_i^ε

The first lemma shows the suite intuitive result that, if α is the iterative dominance solution of g , then playing α_i is conditionally dominant at g when observations are only slightly noisy.

LEMMA 4.2: *If, for a compact set $S \subseteq \Theta$, $g(S) \subseteq D^\alpha$, then $S \subseteq A_1^\varepsilon \cap A_2^\varepsilon$ for all sufficiently small ε .*

PROOF: For any $x \in S$, either (i) $g(x) \in D_1^\alpha$ so that $g_1^\alpha(x) > 0$, $g_2^\alpha(x) > 0$ and $g_1^\beta(x) < 0$ or (ii) $g(x) \in D_2^\alpha$ so that $g_1^\alpha(x) > 0$, $g_2^\alpha(x) > 0$ and $g_2^\beta(x) < 0$. Since g is continuously differentiable, we may choose ε small enough so that, for all $x \in S$, these inequalities remain satisfied on $B(x, 3\varepsilon)$. It will suffice to consider case (i). Then, if player 1 makes an observation in $B(x, 2\varepsilon)$, he knows that α_1 is a dominant strategy for each possible realization of the game. Consequently, $B(x, 2\varepsilon) \subseteq A_1^{\varepsilon, 1}$. If player 2 observes x , he knows that player 1's observation is in $B(x, 2\varepsilon)$ and, hence, that α_1 is conditionally dominant for 1. He also knows that α_2 is the unique best response to α_1 . Thus, $x \in A_2^{\varepsilon, 2}$ and the proof is complete since the $A_i^{\varepsilon, n}$ are nondecreasing in n . Q.E.D.

The next lemma derives a convenient characterization of A_i^ε when the associated game is well inside G_i^+ : Then, the points in A_i^ε are exactly those for which α_i is strictly better given that j chooses α_j on A_j^ε .

LEMMA 4.3: *Define the strategy a_j^ε of player j in Γ^ε by*

$$(4.8) \quad a_j^\varepsilon(x_j) = \begin{cases} 1 & \text{if } x_j \in A_j^\varepsilon, \\ 0 & \text{otherwise.} \end{cases}$$

Let $x_i \in \Theta_i^\varepsilon$ be such that $g(x) \in G_i^+$ for all $x \in B(x_i, \varepsilon)$. Then $x_i \in A_i^\varepsilon$ if and only if $V_i^\varepsilon(a_j^\varepsilon|x_i) > 0$.

PROOF: Since $A_j^\varepsilon \supseteq A_j^{\varepsilon, n}$ for all n , it is obvious that $a_j^\varepsilon \in S_j^{\varepsilon, n}$. If $x_i \in A_i^\varepsilon$, then $x_i \in A_i^{\varepsilon, n+1}$ for some n . Hence $V_i^\varepsilon(s_j|x_i) > 0$ for all $s_j \in S_j^{\varepsilon, n}$ and, therefore, $V_i^\varepsilon(a_j^\varepsilon|x_i) > 0$.

Conversely, assume $V_i^\varepsilon(a_j^\varepsilon|x_i) > 0$ and define $a_j^{\varepsilon, n}$ as in (4.8) but with A_j^ε replaced by $A_j^{\varepsilon, n}$. For $s_j \in S_j^{\varepsilon, n}$ we have

$$\begin{aligned} V_i^\varepsilon(s_j|x_i) &= \int \int [s_j(x_j)(g_i^\alpha(x) + g_i^\beta(x)) - g_i^\beta(x)] dF_i^\varepsilon(x, x_j|x_i) \\ &\geq \int \int [a_j^{\varepsilon, n}(x_j)(g_i^\alpha(x) + g_i^\beta(x)) - g_i^\beta(x)] dF_i^\varepsilon(x, x_j|x_i) \\ &= V_i^\varepsilon(a_j^{\varepsilon, n}|x_i), \end{aligned}$$

where the inequality follows from the assumption that $g(x) \in G_i^+$ for all $x \in B(x_i, \varepsilon)$. Hence, $x_i \in A_i^{\varepsilon, n+1}$ if and only if $V_i^\varepsilon(a_j^{\varepsilon, n}|x_i) > 0$. Now $a_j^{\varepsilon, n}$ converges pointwise to a_j^ε so, by Lebesgue's bounded convergence theorem

$$V_i^\varepsilon(a_j^{\varepsilon, n}|x_i) \rightarrow V_i^\varepsilon(a_j^\varepsilon|x_i) \quad \text{as } n \rightarrow \infty.$$

Therefore, $V_i^\varepsilon(a_j^{\varepsilon, n}|x_i) > 0$ and, hence, $x_i \in A_i^{\varepsilon, n}$ for n large enough which implies $x_i \in A_i^\varepsilon$. Q.E.D.

For a strategy s_j of player j in Γ^ε , let $s_j^\varepsilon(x_i)$ denote the probability that player i assigns to j choosing action α_j if his observation is x_i :

$$(4.9) \quad s_j^\varepsilon(x_i) = \int s_j(x_j) dF_i^\varepsilon(x_j|x_i).$$

The next lemma shows that in the case where $V_i^\varepsilon(s_j|x_i) = 0$, this probability is close to the probability $\bar{s}_j(x_i)$ (as defined in (2.2)) that makes player i indifferent between his two pure actions in $g(x_i)$. (Of course, if $g_i^\alpha(x_i)$ or $g_i^\beta(x_i)$ is negative, $\bar{s}_j(x_i)$ cannot be interpreted as a probability, but Lemma 4.2 shows that such an x_i cannot be a zero of V_i^ε if ε is small.)

LEMMA 4.4: *Let $g(x_i) \in G_i^+$. There exists a constant $k(x_i)$ such that, for all $s_j \in S_j^\varepsilon$, if $V_i^\varepsilon(s_j|x_i) = 0$, then $|s_j^\varepsilon(x_i) - \bar{s}_j(x_i)| \leq k(x_i)\varepsilon$.*

PROOF: Let k be an upper bound on the absolute values of the partial derivatives of g on Θ and define $k(x_i) = mk/(g_i^\alpha(x_i) + g_i^\beta(x_i))$. Then

$$\begin{aligned} V_i^\varepsilon(s_j|x_i) &= \int \int [s_j(x_j)(g_i^\alpha(x) + g_i^\beta(x)) - g_i^\beta(x)] dF_i^\varepsilon(x, x_j|x_i) \\ &\leq \int [s_j(x_j)(g_i^\alpha(x_i) + g_i^\beta(x_i)) - g_i^\beta(x_i)] dF_i^\varepsilon(x_j|x_i) + mk\varepsilon \\ &= \left[\int [s_j(x_j) - \bar{s}_j(x_i)] dF_i^\varepsilon(x_j|x_i) + k(x_i)\varepsilon \right] [g_i^\alpha(x_i) + g_i^\beta(x_i)] \\ &= [s_j^\varepsilon(x_i) - \bar{s}_j(x_i) + k(x_i)\varepsilon] [g_i^\alpha(x_i) + g_i^\beta(x_i)]. \end{aligned}$$

Consequently, if $V_i^\varepsilon(s_j|x_i) = 0$, then $s_j^\varepsilon(x_i) \geq \bar{s}_j(x_i) - k(x_i)\varepsilon$. The reverse inequality is proved in the same way. Q.E.D.

An important consequence of Lemmas 4.3 and 4.4 is that $a_j^\varepsilon(x_i)$ must be close to $\bar{s}_j(x_i)$ if x_i is a boundary point of A_i^ε . Our final lemma shows that if we start looking for boundary points within a dominance solvable region of R^α , then any boundary point of A_i^ε must be close to a boundary point of A_i^2 and vice versa.

LEMMA 4.5: *If x_i is a boundary point of A_i^ε and $g(x) \in R^\alpha$ for all x in $B(x_i, \varepsilon)$, then $B(x_i, 2\varepsilon) \not\subseteq A_j^\varepsilon$.*

PROOF: If $B(x_i, 2\varepsilon) \subseteq A_j^\varepsilon$, then, by assumption (A4),

$$V_i^\varepsilon(a_j^\varepsilon|x_i) = \int \int g_i^\alpha(x) dF_i^\varepsilon(x, x_j|x_i) > 0$$

since $g_i^\alpha(x) > 0$ for all relevant x . Hence $x_i \in A_i^\varepsilon$ by Lemma 4.3. Q.E.D.

4.4. The Main Step of the Proof

Since in the general, multi-dimensional case the main argument of the proof is obscured by a number of technical details, we defer the concluding step for that case to the Appendix. Here we give the argument for the one-dimensional case where $\Theta \subseteq \mathbb{R}$.

We argue by contradiction: Assume x' is contained in an interval $C = [\underline{x}, \bar{x}]$ such that $g(C) \cap D^\alpha \neq \emptyset$ and $g(C) \subseteq R^\alpha$, but that nevertheless $x' \notin A_1^\varepsilon \cap A_2^\varepsilon$ for arbitrarily small ε . By Lemma 4.2, then, $g_i^\beta(x') \geq 0$ for $i = 1, 2$, and, thus, $g(x') \in G_1^+ \cap G_2^+$. Without loss of generality, we may assume the following:

(i) $g(C) \subseteq G_1^+ \cap G_2^+$ (if this condition is not satisfied, restrict attention to the segment of C around x' that lies inside $G_1^+ \cap G_2^+$; since $\min_i g_i^\alpha(x)$ has a positive lower bound on C , this segment necessarily contains points in D^α as well);

(ii) $g(\underline{x}) \in D^\alpha$; and

(iii) $\underline{x} < x' < \bar{x}$ (as the relevant sets are open, the curve can be slightly extended if desired).

Noting that C is compact and reviewing the constructions of Lemmas 4.1 and 4.4, we see that there exists a constant k such that

$$k(x_1) \leq k \quad \text{and} \quad k(x_1, x_2) \leq k \quad \text{for all} \quad x_1, x_2 \in C.$$

If ε is small enough and such that $x' \notin A_1^\varepsilon \cap A_2^\varepsilon$, then, by Lemma 4.5, both A_1^ε and A_2^ε must have boundary points on C . We let x_i^ε denote the first boundary points starting from \underline{x} :

$$x_i^\varepsilon = \sup \{x \in C : [\underline{x}, x] \subseteq A_i^\varepsilon\}.$$

Lemma 4.3 and the continuity of V_i^ε imply that $V_i^\varepsilon(a_j^\varepsilon|x_i^\varepsilon) = 0$ for $i = 1, 2$. Hence, by Lemma 4.4,

$$(4.10) \quad a_j^\varepsilon(x_i^\varepsilon) \leq \bar{s}_j(x_i^\varepsilon) + k\varepsilon \quad \text{for} \quad i, j = 1, 2, \quad i \neq j.$$

On the other hand, for sufficiently small ε , $x_i^\varepsilon - \underline{x} > 2\varepsilon$ by Lemma 4.2. Hence

$$a_j^\varepsilon(x_i^\varepsilon) \geq P(X_j^\varepsilon \leq x_j^\varepsilon | X_i^\varepsilon = x_i^\varepsilon)$$

and therefore, by Lemma 4.1,

$$(4.11) \quad a_2^\varepsilon(x_1^\varepsilon) + a_1^\varepsilon(x_2^\varepsilon) \geq 1 - k\varepsilon.$$

Combining (4.10) and (4.11) yields

$$\bar{s}_2(x_1^\varepsilon) + \bar{s}_1(x_2^\varepsilon) \geq 1 - 3k\varepsilon.$$

Letting ε go to zero, we see that for the common (Lemma 4.5) limit x^* of x_1^ε and x_2^ε

$$\bar{s}_2(x^*) + \bar{s}_1(x^*) \geq 1,$$

but this contradicts the fact that $g(x^*) \in R^\alpha \cap G_1^+ \cap G_2^+$ (cf. (2.4)) and completes the proof.

5. ROLE OF ASSUMPTIONS

The equilibrium selection result presented above is driven by two fundamental features: the existence—guaranteed by the condition stated in the Theorem—of dominance solvable regions D^α and D^β within the class of games $g(\Theta)$; and the continuity—ensured by assumptions (A1), (A2) and (A3)—of the function g as well as the densities of the prior and the observation errors. In contrast, the assumption (A4) of bounded error supports is made only to simplify our argument but is not really essential. As long as payoffs and the error density are bounded, (A4) can be dispensed with. The following discussion will therefore focus on the condition of the Theorem and the continuity assumptions.

It has already been made clear in the introduction that our result depends crucially on the existence, within the class of games, of regions where each of the two action pairs γ , $\gamma \in \{\alpha, \beta\}$, can be established as strict dominance solutions. These regions are the indispensable starting points for the iterated dominance argument in the global game. The Theorem adds to the nonemptiness of these regions the requirement that from the game $g(x)$ in R^γ there exist a continuous link to another game where γ is the strict dominance solution through a region where the risk dominance relationship is never reversed. The study of simple examples, however, suggests that this second property is likely to be implied by the existence of strict dominance regions. This is indeed the case for the example in the introduction provided that the parameter x is drawn from an interval. The implication also holds for the “natural” parameterization $g(x) = x$ when each payoff entry g^k ($k = 1, \dots, 8$) corresponds to an independent random variable X^k which is defined on a real interval.

The proof of the last assertion may be sketched as follows: Suppose $x' \in \Theta$, $g(x') \in R^\alpha$ and, for some $x'' \in \Theta$, $g(x'') \in D^\alpha$ so that the nonemptiness condition is met. Without loss of generality we may assume that $g_1^\beta(x'') < 0$. To simplify, also assume that $g_2^\beta(x'') < g_2^\beta(x')$ (other cases can be handled in a similar way). Let \bar{x} be the point which agrees with x' for the six payoffs that determine g_1^α , g_2^α , and g_2^β and agrees with x'' for the two payoffs that determine g_1^β . Moreover let C be the straight line between \bar{x} and x' . Then, by construction, $\bar{x} \in \Theta$, $g(\bar{x}) \in D^\alpha$, $C \subseteq \Theta$, and $g(C) \subseteq R^\alpha$ so the condition of the Theorem is satisfied.

Another essential prerequisite for our result is that, given a strategy for the opponent, a player's expected payoff (and thus his expected net gain $V_i^\varepsilon(\cdot | \cdot)$) from using a certain action is continuous in his observation. This ensures that a strict indifference condition ($V_i^\varepsilon(\cdot | x_i) = 0$) has to be satisfied at any point x_i where the player optimally switches from one action to the other, a condition which is exploited in the characterization of the boundary points of the iterated dominance sets A_i^ε . The continuity of the expected payoffs results from the continuity of the prior density (A2) and the joint error density (A3). Note, however, that the latter assumption can be weakened: A careful inspection of the proofs reveals that the Theorem holds as long as the *difference* between the error terms, $E_1 - E_2$, allows a continuous density. A potentially interesting variant which meets this requirement is when one of the players observes exactly the true value while the other player's observation contains an error term.

To see the importance of continuity, it is instructive to consider a variant where, in the global game Γ^ε , each player i observes the true value of X with probability $1 - \varepsilon$, i.e. X_i^ε is given by

$$X_i^\varepsilon = X + Y_i^\varepsilon,$$

where Y_i^ε is a (bounded) random variable satisfying $\text{Prob}\{Y_i^\varepsilon = 0\} = 1 - \varepsilon$. Now consider $x \in \Theta$ such that $g(x)$ has both α and β as strict equilibria. We claim that, for sufficiently small ε , playing any equilibrium of $g(x)$ is compatible with (Bayesian Nash) equilibrium in a global game Γ^ε satisfying (A1) and (A2). To see why, it suffices to note that

$$\lim_{\varepsilon \rightarrow 0} \text{Prob}\{X = x \text{ and } X_j^\varepsilon = x | X_i^\varepsilon = x\} = 1,$$

i.e. as ε goes to zero, each player becomes certain that both he and his opponent have observed the true value. Another variant which yields the same result is when the prior on Θ , instead of having an everywhere positive density, has support on a finite number of points. In this set-up (keeping (A3) and (A4)), for ε smaller than half the minimum distance between any of the points, the true value of X becomes common knowledge to the players.

Somewhat more surprisingly, similar results are obtained even if there is a real link between different games as long as the class of games is discrete. These results are derived in Monderer and Samet (1989) who consider the following set-up: Let (Ω, Σ, μ) be a probability space, for $\omega \in \Omega$, let $g(\omega)$ be a game with a finite set of players and assume that $\{g(\omega) : \omega \in \Omega\}$ is finite. A game Γ of incomplete information is now defined in the natural way. Monderer and Samet address the question under what conditions Γ has an ε -equilibrium which is close to an arbitrary selection of equilibria from the underlying class of games. Formally, if for each $\omega \in \Omega$ $s(\omega)$ is an equilibrium of $g(\omega)$, does there exist an ε -equilibrium of Γ that is close to s ? The answer is shown to be affirmative if, on a large (close to measure one) subset of Ω , each player knows with high probability which game will be played. This result thus suggests that the nonfinite state space is essential for our result.

6. EQUILIBRIUM SELECTION AND THE NASH PROGRAM

In this section we discuss the role of risk dominance and payoff dominance in Harsanyi and Selten's equilibrium selection theory for the case of 2×2 games and compare their argument with the justification for risk dominance which can be derived from our own model. We also summarize the important links between our approach and Nash's work on bargaining and noncooperative games.

6.1. Risk Dominance vs. Payoff Dominance

Harsanyi and Selten offer two complementary justifications of the risk dominance criterion for 2×2 games with two strict equilibria. The one is an axiomatic derivation while the other, more heuristic justification relies on a simplified version of the tracing procedure which the authors use in order to define the risk dominance relation for general games.

The axiomatic derivation is based on three axioms:

- (i) invariance with respect to isomorphisms,
- (ii) best-reply invariance, and
- (iii) payoff monotonicity.

The first axiom rules out any dependence on the labeling of players and actions or on positive linear utility transformations. In particular it implies that, in a symmetric game in which both α and β are strict equilibria, the solution has to be the mixed equilibrium. Axiom (ii) guarantees that the risk dominance relation in a game where α and β are strict equilibria only depends on the four deviation losses g_1^α , g_2^α , g_1^β , and g_2^β . In the same setting the last axiom says essentially that if, starting from a game where β does not risk-dominate α , one or both payoffs associated with the outcome α are increased while all other payoffs remain unchanged, then α risk-dominates β in the new game. Harsanyi and Selten show that the only definition of risk dominance that satisfies these axioms is the one given above, i.e. the equilibrium with the largest Nash product is risk-dominant.

In the heuristic justification, selection of the risk-dominant equilibrium results from postulating an initial state of uncertainty where the players have uniformly distributed second-order beliefs. To be more specific, let j 's initial beliefs about i 's choice of action be represented by the mixture $z\alpha_i + (1-z)\beta_i$, $z \in [0, 1]$. Player i does not know j 's beliefs so, applying the principle of insufficient reason, he considers all values of z to be equally likely. Player i then reasons that, whatever the value of z , j will use a best response against his beliefs, i.e. j will play α_j or β_j according as $z > \bar{s}_i$ or $z < \bar{s}_i$, the critical probability \bar{s}_i being defined as in (2.2). Hence, player i expects j to play $(1 - \bar{s}_i)\alpha_j + \bar{s}_i\beta_j$ and will choose a best response against this mixed action, i.e. he will play α_i if $1 - \bar{s}_i > \bar{s}_j$ and β_i if $1 - \bar{s}_i < \bar{s}_j$. Comparing with (2.3) one sees that, if both players follow this line of reasoning, they will end up coordinating on the risk-dominant equilibrium.

	α_2	β_2
α_1	3,3	3,0
β_1	0,3	4,4

FIGURE 2.—Game $g(3)$.

It is certainly remarkable that two such different approaches lead to the same selection criterion. This fact suggests that the risk dominance relation is a strategically highly significant property in 2×2 games. On the other hand, it has to be admitted that none of the justifications is built on an altogether solid basis. For instance, without going into details, it is not made clear why an equilibrium selection rule should satisfy the axiom of best-reply invariance. It is also difficult to find a compelling reason why the players' initial beliefs have to be uniformly distributed as in the heuristic approach. In this connection, it should be noted that—in contrast to our model where the exact shape of the prior is unimportant—the outcome of the heuristic approach depends crucially on the prior. For instance, in the game shown in Figure 2 (which is the game $g(x)$ in Figure 1 with $x = 3$), the heuristic argument selects the Pareto-dominant equilibrium β if the prior has a sufficient bias in favor of this equilibrium. (Assume, e.g., that player i 's beliefs about j 's beliefs have distribution function $F(z) = z^{1/3}$.)

The shakiness of the justifications appears perhaps most clearly from the fact that Harsanyi and Selten themselves in the end choose not to use risk dominance as a general criterion. Instead they decide that payoff dominance should have precedence in case of conflict between the two criteria. Hence, in the game $g(3)$, the Harsanyi/Selten theory selects the payoff-dominant equilibrium β although α is risk-dominant. Their motivation for giving precedence to payoff dominance is that “(...) risk dominance is important only in those situations where the players would be initially uncertain whether the other players would choose one equilibrium or the other. Yet, if one equilibrium would give every player higher payoffs than the other would (...) every player can be quite certain that the other players will opt for this equilibrium which will make risk dominance considerations irrelevant” (Harsanyi and Selten (1988, p. 358)). In another passage of the book they admit that the players may be justified in relying only on risk dominance in some situations, for instance after a breakdown of preplay communication. Nevertheless, they conclude that “a theory that considers both payoff and risk dominance is more in agreement with the usual image of what constitutes rational behavior” (pp. 89–90).

These arguments for giving precedence to payoff dominance are hardly more compelling than the previous justifications of risk dominance. Rather, Harsanyi and Selten's hesitation on this point seems to be a sign that their theory is located on a level of abstraction where the problem of equilibrium selection does not always have a determinate solution. Without knowing more details about the context of the game, including the possibilities of preplay communication, it is difficult to predict whether the players will conform to risk dominance,

payoff dominance, or some other criterion. In this respect we think that the global game approach presented in this paper constitutes a certain advance. By perturbing the players' payoff information, we obtain an incomplete information game where play in accordance with risk dominance results from a standard noncooperative equilibrium concept. Thus, we embed any given 2×2 game in a context which is sufficiently rich to generate a determinate solution to the equilibrium selection problem.

One advantage of our approach is that it supplies a more precise intuition for risk dominance. The Harsanyi/Selten theory presents risk dominance as a measure of the relative stability of equilibria in the presence of strategic uncertainty, but—granted the ad hoc character of the heuristic derivation—it fails to provide a convincing argument for the relevance of such uncertainty. (Indeed, in the above quotation, the precedence given to the payoff dominance criterion is justified by the *absence* of uncertainty in games where this criterion is applicable.) In our model uncertainty about your opponent's choice of action appears when your observation is close to the point where he switches from one action to the other. Hence, the optimal choice of a switching point requires the players to consider the relative riskiness of the two alternative equilibria. In view of this, it is quite intuitive that the equilibria must be approximately equivalent in terms of risk dominance at the switching points.

6.2. *Nash's Bargaining Theory and the Nash Program*

The preceding analysis of equilibrium selection in 2×2 games has several notable connections with Nash's (1950, 1951, 1953) seminal work on bargaining and noncooperative games. First, there is an obvious link between risk dominance and Nash's bargaining solution since both choose the point corresponding to the largest Nash product (in the bargaining set-up this means the product of utility differences between the agreement and the disagreement point). Actually, in their search for a satisfactory definition of risk dominance, Harsanyi and Selten were guided by the desire "to reproduce the results of Nash's bargaining theory with fixed threats" (Harsanyi and Selten (1988, p. 215)).

Secondly, Nash justifies his bargaining solution by means of an axiomatic derivation where a leading idea consists of imposing conditions on how the solution should be allowed to vary on a whole class of games. Harsanyi and Selten's axiomatic derivation relies on the same type of conditions. Actually all their axioms explicitly force solutions of different games to be linked to each other. Our approach also exploits conditions of this kind, but, rather than simply postulating them, we derive them from more basic conditions of imperfect information.

Moreover, Nash complements his axiomatic derivation with an explicit noncooperative approach where he tries to select the axiomatic solution by adding noise to the bargaining game. This approach was not completely satisfactory for Nash did not succeed in excluding the possibility of multiple solutions, but his basic intuition has been confirmed in recent work by Binmore (1987) and, in a

more detailed model, by Carlsson (1991). These authors add noise to the parties' bids in Nash's noncooperative game and show the existence of a unique solution as the noise vanishes. The formal structure of these bargaining games is very similar to that of global games. In particular, the problem of choosing optimal bids in a noisy bargaining game involves the same kind of risk considerations as the problem of choosing a point where to switch from one action to the other in a global game.

Nash's strategic analysis of bargaining inaugurates a tradition which has become known as the Nash program. The common denominator of this tradition is a dissatisfaction with the ad hoc character of many cooperative or axiomatic solution concepts and the wish to derive predictions within strictly noncooperative models. In our view, the approach of this paper could be viewed as an attempt to extend the Nash program to a more general field of equilibrium selection.

7. COMMON KNOWLEDGE AND PAYOFF UNCERTAINTY

In this section we compare our approach with two other models of payoff uncertainty, namely Rubinstein's electronic mail game and Harsanyi's games with randomly disturbed payoffs. In the final subsection we shortly discuss the role of common knowledge for the justification of the iterated dominance solution concept that we have been using.

7.1. Common Knowledge and Rubinstein's Electronic Mail Game

As we have already pointed out, our result is driven by the remote influence that the dominance solvable regions D^α and D^β exert on the region of games, $G^\alpha \cap G^\beta$, having both α and β as strict equilibria. This influence is due to a certain "lack of common knowledge" in the global game. Once this lack of common knowledge is taken into account, it is no longer paradoxical that, in the game $g(3)$ of Figure 2, the players are forced to select α although, when ε is small, they know that a uniformly better equilibrium β is available: If x is drawn from an interval $[x, \bar{x}]$ with $\bar{x} > 4$, the fact that β is an equilibrium is not *common knowledge*.

To formalize this assertion, let $\Omega = \{\omega: \omega = (x, x_1, x_2)\}$ be the state space and let E be an event. Denote by KE the event that both players know E and write $K^{n+1} = K(K^n E)$. Then, the event E is common knowledge at all states ω belonging to $\bigcap_n K^n E$. Now, let E denote the event that β is an equilibrium, i.e. $E = \{\omega \in \Omega: x \leq 4\}$. Player i knows that E obtains if and only if $x_i^\varepsilon \leq 4 - \varepsilon$, i.e. $KE = \{\omega \in \Omega: x_i^\varepsilon \leq 4 - \varepsilon, i = 1, 2\}$; player i knows that player j knows that E obtains if and only if $x_i^\varepsilon \leq 4 - 3\varepsilon$ ($K^2 E = \{\omega \in \Omega: x_i^\varepsilon \leq 4 - 3\varepsilon, i = 1, 2\}$). Continuing inductively one sees that the event E cannot be common knowledge for any $\varepsilon > 0$.

A similar lack of common knowledge occurs in Rubinstein's (1989) electronic mail game, where it also gives rise to a kind of action from a distance. To




FIGURE 3

describe Rubinstein’s model, refer to Figure 3 where we let $g = g(3)$ and $g' = g(5)$, $g(x)$ being defined as in Figure 1. Then g' is dominance solvable with solution α while g has both α and β as strict equilibria. Note that α is risk-dominant but β is Pareto-dominant in g . Rubinstein considers the following situation: First, one of the games, g or g' , is selected, each with probability $1/2$. Player 1 is always informed about the game selected. Player 2 is not informed if the game is g' . If the game is g , player 1 automatically sends player 2 a message (saying that the game is g) which player 2, upon receiving it, automatically acknowledges, an acknowledgement which, when received by 1 is again automatically acknowledged, etc. The communication technology, however, is imperfect: Each message gets lost with a small probability ε and, as soon as a message is lost, the communication process stops. As a result it never becomes common knowledge which game has been selected.

Rubinstein shows that even if many messages have been exchanged and, hence, the players are sure that the game is g , each knows that the other knows this, etc., rational players will not choose the payoff-dominant equilibrium β . The argument is simple and uses induction on the number of messages received. Obviously, conditional on the fact that he has not received any message, player 2 will choose α since he will consider it very likely that the game is g' . Conditional on the fact that he has received k but not $k + 1$ messages, player i considers it about equally likely that it was his acknowledgement that got lost as that j ’s acknowledgement of the k th message got lost. By the induction hypothesis, i knows that j will play α_j in the first case and thus concludes that the probability of j choosing β_j is at most around $1/2$. Therefore, the expected payoff associated with β_i is at most around 2 so it is conditionally dominant for i to choose α_i .

The reader notices that the forces at work in Rubinstein’s model closely parallel those in our global game model: The result is driven by the fact that α is dominant in g' and risk-dominant in g . Rubinstein considers it paradoxical that the Pareto-dominant equilibrium is excluded since it implies “that the game-theoretic ‘prediction’ for the ‘almost common knowledge’ situation is very different from the situation with common knowledge” (Rubinstein (1989, p. 385)). We are not convinced. We do not know what *the* game-theoretic “prediction” is for the common knowledge situation. In our view, it is not obvious that this “prediction” should be identified with the Pareto-dominant

equilibrium. The prediction for a game with multiple equilibria—in particular if there is a conflict between Pareto dominance and risk dominance—cannot be determined independently of the context in which the game is played. Rubinstein's electronic mail game as well as global games are examples of contexts where the informational conditions force the players to choose the risk-dominant equilibrium.

7.2. Harsanyi's Randomly Disturbed Games

Harsanyi (1973) analyzes games where each player is slightly uncertain about the payoffs of his opponents. His major finding is that, in this perturbation, there is no need for the players to actively randomize: A mixed strategy equilibrium of a normal form game can be approximated by pure strategy equilibria of slightly perturbed games. For the present context, however, the most significant result is that, in generic normal form games, all Nash equilibria are robust against Harsanyi's perturbation. Hence, it is interesting to compare Harsanyi's model and our own approach for 2×2 games in order to see why their results are so different.

Let G be the set of 2×2 games in which each player i has the action set $S_i = \{\alpha_i, \beta_i\}$ and let $S = S_1 \times S_2$. Any $g \in G$ can be written as $g = (g_1, g_2)$ where $g_i \in \mathbb{R}^4$ describes i 's payoffs in g . For $s \in S$ and $i = 1, 2$, let $E_i(s)$ be a random variable which takes values in \mathbb{R}^4 and has a continuous density; assume E_1 and E_2 are independent. For $g \in G$ and $\varepsilon > 0$, Harsanyi considers the incomplete information game $\Gamma^\varepsilon(g)$ described by the following rules:

1. A realization (e_1, e_2) of (E_1, E_2) is drawn.
2. Player i is informed about e_i (and about nothing else).
3. The players choose actions $s = (s_1, s_2)$ simultaneously.
4. Each player i receives payoff $g_i(s) + \varepsilon e_i(s)$.

Comparing Harsanyi's approximating sequence $\{\Gamma^\varepsilon(g)\}_{\varepsilon \downarrow 0}$ with the sequence of global games $\{\Gamma^\varepsilon\}_{\varepsilon \downarrow 0}$ as defined in Section 3, one notes three important differences:

(i) Harsanyi lets the prior uncertainty vanish while we keep the prior constant and let the residual uncertainty, which remains after the players have made their observations, vanish for each observation.

(ii) In Harsanyi's set-up different players' payoffs are independently distributed. Our model allows both for independence and correlation in this respect.

(iii) In Harsanyi's model a player learns his own payoffs exactly but does not receive any information about the opponent's payoffs. Since payoffs are independent, the player's observation does not tell him anything about the opponent's payoffs. In our model, players make (imperfect) observations of the entire game, i.e. they learn something about both players' payoffs. Since observations are correlated with the actual game, observations of different players are obviously correlated as well.

The third point in particular is of major importance. A necessary prerequisite for the action from a distance which occurs in our model is that a player's beliefs about what the opponent is going to do are correlated with the player's observation: If player i observes a game close to D^α , then he will think that his opponent is likely to play α_j (if only because it is likely that j will have observed a game actually belonging to D^α) and, hence, i will play α_i . Such a phenomenon cannot occur in Harsanyi's model where player i 's beliefs about j 's behavior are independent of his observation. Actually, in Harsanyi's model, the players' first-order beliefs are common knowledge. Hence, in contrast with what happens in our model, there is no increase in the range of uncertainty as you ascend to higher levels in the hierarchy of beliefs and, as ε goes to zero, it becomes common knowledge that the actual game is very close to g (provided that the error terms E_1 and E_2 are bounded.)

Hence, there are considerable differences between the two models and it should not be surprising that the results are completely different as well. For a 2×2 game g with three equilibria Harsanyi shows that all three equilibria can be approximated by the beliefs associated with equilibria of the game $\Gamma^\varepsilon(g)$; i.e. for every equilibrium s of g , there exists an equilibrium s^ε of $\Gamma^\varepsilon(g)$ for which the associated beliefs (defined as in (4.9) but independent of x_i) converge to s as ε tends to zero. On the other hand, in our model, only the risk-dominant equilibrium can be approximated; i.e. if s is the risk-dominant equilibrium of g and s^ε is an equilibrium of Γ^ε , then $s^\varepsilon(g)$ converges to s . Note the difference between the two notions of convergence: We investigate pointwise convergence of equilibrium strategies while Harsanyi studies convergence of equilibrium beliefs. In Appendix B, we construct a hybrid model which contains our own model and Harsanyi's as extreme special cases. The mathematical analysis clearly brings out the differences between the various models.

In view of the important differences between the results of the two models one would, of course, like to know which one provides the most relevant description of payoff uncertainty for economic and other applications. The answer to this question requires more knowledge about how agents acquire information than we possess today, but the following remarks may nevertheless be of some interest. In Harsanyi's model, the players' first-order posterior beliefs are assumed to be common knowledge, i.e. although each player is uncertain about his opponent's exact payoffs, the exact distribution from which they are drawn are common knowledge. This kind of assumption has now become standard in game-theoretic applications, but it is open to a rather obvious objection: Granted that it is unrealistic to assume that the exact payoffs are common knowledge, why would this assumption be more plausible when applied to first-order beliefs? Our model escapes this difficulty since there is residual uncertainty at any level of posterior beliefs; in other words there is no n such that the players' n th-order beliefs are common knowledge. On the other hand, as long as we want to interpret our result within the traditional game-theoretic framework, we still need the common knowledge assumption at the

level of the rules of the global game. In the following subsection we briefly discuss alternative interpretations.

7.3. *Iterated Dominance and Common Knowledge*

Our Theorem is based on a relatively weak equilibrium concept, viz. iterated elimination of dominated strategies. An attractive feature of iterated dominance—in contrast with stronger concepts such as Nash equilibrium or Bayesian Nash equilibrium—is that we know a general condition which ensures its implementation: In the context of our model, if the rules of the global game—as given by ε and the functions g , h , and φ —and the players' rationality are common knowledge, then the players *will* choose serially undominated strategies and, thus, act in conformity with our Theorem.

This justification of our solution concept, however, may not be totally convincing since a major motivation for our approach has been the wish to relax the common knowledge assumption. Hence, it would be desirable to find alternative justifications which—instead of simply moving this assumption from one level of analysis to another—dispense with it altogether. The kind of stories that naturally come to one's mind are those where the strategy choices, instead of being determined by strictly rational considerations, result from some learning or evolutionary process. Intuitively, granted that such a process would favor the use of strategies that are relatively successful in terms of average payoffs, a stable state seems to require that only serially undominated strategies are used. Of course, this intuition remains to be formalized and proved, but we feel that recent progress with learning models are very promising in this respect. In particular, Milgrom and Roberts (1989) show that, for an important class of games and a great variety of learning processes, the sequence of strategy choices will eventually be confined to the set of strategies which survive iterated elimination of strictly dominated strategies.

8. CONCLUSION

Two central aims in recent game-theoretic research have been to arrive at unique solutions³ on the one hand, and to incorporate more realistic informational assumptions on the other. There is often thought to exist a conflict and, thus, a necessary trade-off between these two goals. The concept of global games that has been presented here indicates that this need not be the case. On

³The importance of uniqueness is stressed by Robert Aumann who writes in the foreword to Harsanyi and Selten (1988): "Nash equilibrium makes sense only if each player knows which strategies the others are playing; if the equilibrium recommended by the theory is not unique, the players will not have this knowledge. Thus it is essential that for each game, the theory selects one unique equilibrium from the set of all Nash equilibria."

the contrary, the use of an informational set-up implying a considerable weakening of the common knowledge assumption has been shown to generate a model with interesting equilibrium selection properties.

The paper's main message is that something can be gained by moving from the conventional local analysis of individual games to a global analysis of classes of games: an equilibrium of a given game need not be consistent with an equilibrium rule for the entire class of games. The global game approach provides a natural way to force players to link games together and analyze them simultaneously. Unfortunately, the approach turned out not to be easily tractable mathematically and we have been unable to pursue its implications, in a general set-up, beyond the restricted class 2×2 games. Extensions to other classes of games are therefore urgently called for. In this connection two distinct questions become relevant. The first concerns the classes of games for which the global game is dominance solvable in the limit, while the second concerns the nature of the solution. To conclude, we discuss some results in Carlsson (1989) and Carlsson and van Damme (1991) which concern other classes of games but are based on somewhat more restrictive assumptions than the ones used above.

Carlsson (1989) analyzes $m \times m$ unanimity games defined on a one-dimensional parameter space and shows, roughly, that an equilibrium which risk-dominates all other equilibria will be selected by the global game approach. Unfortunately, equilibrium selection rules based on Nash products of deviation losses become intransitive—and, thus, incapable of selecting a solution—in general two-person games (see Harsanyi and Selten (1988, p. 112 and pp. 216–217)). Hence, even for this restricted class of games, there is some reason to be pessimistic about the chances of always generating uniqueness by the present approach.

In Carlsson and van Damme (1993) we consider a class of n -person binary choice games. Each player has the choice between a safe strategy that yields a fixed payoff x and a risky strategy the payoff of which depends on the total number of players that choose the same strategy. If the number of players choosing the risky strategy is small, it yields a low payoff, but if many choose it, it yields a payoff that is higher than x . Hence, there is again a conflict between risk dominance and payoff dominance. We analyze the global game in which the value of x is observed with some noise and show that this game is approximately dominance solvable. It turns out that the derived selection rule does not coincide with selection based on a naïve comparison of Nash products and also differs from the selection rule proposed by Harsanyi and Selten, which shows that global games need not justify the risk dominance principle outside the class of 2×2 games. Quite interestingly, Kim (1992) has shown that the outcome selected by the global game approach is also obtained in an entirely different dynamic context: In his model the game is played by randomly and anonymously matched players who are drawn from large populations, have perfect foresight about the evolution of play, but face costs in adjusting their actions. Kim shows that, when the friction vanishes (i.e. the players become very patient or the

adjustment cost becomes negligible), the play settles down at the equilibrium selected by the global game perturbation.

Dept. of Economics, University of Gothenburg, Viktoriagatan 30, 41125 Gothenburg, Sweden

and

Center for Economic Research, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, the Netherlands

Manuscript received June, 1990; final revision received March, 1993.

APPENDIX A

THE FINAL STEP OF THE PROOF IN THE MULTI-DIMENSIONAL CASE

We argue by contradiction: Assume x' lies on a curve C in Θ with $g(C) \cap D^\alpha \neq \emptyset$ and $g(C) \subseteq R^\alpha$, but that arbitrarily small ε can be found such that $x' \notin A_1^\varepsilon \cap A_2^\varepsilon$. By the same arguments as in Section 4.4, we may assume that C is compact, $g(C) \subseteq G_1^+ \cap G_2^+$ and $g(\underline{x}) \in D^\alpha$ where \underline{x} denotes one of the end points of the curve C . Since Θ , D^α , R^α , G_1^+ and G_2^+ are open, we can find $\eta > 0$ such that

$$x \in \Theta \quad \text{and} \quad g(x) \in R^\alpha \cap G_1^+ \cap G_2^+ \quad \text{for all } x \text{ with } d(x, C) \leq 4\eta,$$

$$g(x) \in D^\alpha \quad \text{for all } x \in B(\underline{x}, \eta).$$

Let K be the compact set of all points in Θ that are at a distance of at most 4η of C , i.e., $K = \{x \in \Theta : d(x, C) \leq 4\eta\}$. In the remainder of the proof only points in K will play a role. Choose k so that, for the constants of Lemmas 4.1 and 4.4, $k(x_1) \leq k$ and $k(x_1, x_2) \leq k$ for all $x_1, x_2 \in K$.

Now assume $\varepsilon < \eta/2$ is such that $x' \notin A_1^\varepsilon \cap A_2^\varepsilon$. Also assume ε is small enough so that $B(\underline{x}, \eta) \subseteq A_1^\varepsilon \cap A_2^\varepsilon$ by Lemma 4.2. To construct a point on the boundary of $A_1^\varepsilon \cap A_2^\varepsilon$, let us move along the curve C from \underline{x} in the direction of x' until we first reach a point, say y , for which the closed ball $B(y, \eta)$ is no longer contained in $A_1^\varepsilon \cap A_2^\varepsilon$. Then y is strictly in between \underline{x} and x' and there exists a point, say x_1^ε , on the boundary of $B(y, \eta)$ such that $x_1^\varepsilon \notin A_1^\varepsilon \cap A_2^\varepsilon$. Without loss of generality, we may assume $x_1^\varepsilon \notin A_1^\varepsilon$ so that x_1^ε is a boundary point of A_1^ε . Thus, $V_1^\varepsilon(a_2^\varepsilon | x_1^\varepsilon) = 0$ by Lemma 4.3 and the continuity of $V_1^\varepsilon(a_2^\varepsilon | \cdot)$. By Lemma 4.5 and since $\varepsilon < \eta$, $B(x_1^\varepsilon, 2\varepsilon)$ must contain boundary points of A_2^ε .

To construct a particular boundary point of A_2^ε in $B(x_1^\varepsilon, 2\varepsilon)$, let $p = x_1^\varepsilon - y$ and walk along the ray from y passing through x_1^ε until you reach the first point, say z , such that the halfspace $\{x : px \leq pz\}$ intersected with $B(x_1^\varepsilon, 2\varepsilon)$ is not completely contained in A_2^ε . (See Figure 4 for an illustration.) Let x_2^ε be a point on the hyperplane $\{x : px = pz\}$ within $B(x_1^\varepsilon, 2\varepsilon)$ such that $x_2^\varepsilon \notin A_2^\varepsilon$. Then, obviously, x_2^ε is a boundary point of A_2^ε and, thus, $V_2^\varepsilon(a_1^\varepsilon | x_2^\varepsilon) = 0$. Now note that

$$a_2^\varepsilon(x_1^\varepsilon) \geq P(pX_2^\varepsilon \leq px_2^\varepsilon | X_1^\varepsilon = x_1^\varepsilon)$$

and

$$a_1^\varepsilon(x_2^\varepsilon) \geq P(X_1^\varepsilon \in B(y, \eta) | X_2^\varepsilon = x_2^\varepsilon).$$

In order to apply Lemma 4.1 and to derive a contradiction, we will show that the probability in the right-hand side of the latter inequality is approximately equal to $P(pX_1^\varepsilon \leq px_1^\varepsilon | X_2^\varepsilon = x_2^\varepsilon)$. Note that by construction

$$B(y, \eta) \subseteq \{x : px \leq px_1^\varepsilon\}$$

and let D be the set difference of these sets. We will show that

$$P(X_1^\varepsilon \in D | X_2^\varepsilon = x_2^\varepsilon) \rightarrow 0 \quad \text{as} \quad \varepsilon \rightarrow 0.$$

Let D' be the set of all points x in D with $d(x, x_1^\varepsilon) < 4\varepsilon$ (see Figure 5). Since $d(x_1, x_2^\varepsilon) < 2\varepsilon$ for all




FIGURE 4




FIGURE 5.—The shaded area is D' .

x_1 with $f_2^\varepsilon(x_1|x_2^\varepsilon) > 0$ and since $d(x_1^\varepsilon, x_2^\varepsilon) < 2\varepsilon$, it suffices to show that

$$(A.1) \quad P(X_1^\varepsilon \in D' | X_2^\varepsilon = x_2^\varepsilon) \rightarrow 0 \quad \text{as } \varepsilon \rightarrow 0.$$

Recalling Lemma 4.1 we have

$$P(X_1^\varepsilon \in D' | X_2^\varepsilon = x_2^\varepsilon) \leq (1 + k\varepsilon) \int_{D'} \psi^\varepsilon(x_1 - x_2^\varepsilon) dx_1.$$

By assumption (A3) the density ψ is bounded, say by l , so ψ^ε is bounded by $\varepsilon^{-m}l$. Thus

$$\int_{D'} \psi^\varepsilon(x_1 - x_2^\varepsilon) dx_1 \leq \varepsilon^{-m}l\lambda(D')$$

where λ denotes Lebesgue measure. To prove (A.1), therefore, it suffices to show that $\lambda(D')$ is of order ε^{m+1} . Making a change of variables so that $y = (0, \dots, 0)$ and $x_i^\varepsilon = (0, \dots, 0, \eta)$ we can write D' as

$$D' = \left\{ \xi \in \mathbb{R}^m : \xi_m \leq \eta, \sum_i \xi_i^2 \geq \eta^2, \sum_{i=1}^{m-1} \xi_i^2 + (\xi_m - \eta)^2 \leq 16\varepsilon^2 \right\}.$$

Obviously, $|\xi_i| \leq 4\varepsilon$ for $i = 1, \dots, m - 1$ so by combining the last two inequalities in the definition of D' we obtain that $\xi_m \geq \eta - 8\varepsilon^2/\eta$. Consequently

$$D' \subseteq \{ \xi \in \mathbb{R}^m : |\xi_i| \leq 4\varepsilon \text{ for } i = 1, \dots, m - 1 \text{ and } \eta - 8\varepsilon^2/\eta \leq \xi_m \leq \eta \}$$

so that

$$\lambda(D') \leq (4\varepsilon)^{m-1} 8\varepsilon^2/\eta,$$

and, indeed, $\lambda(D')$ is of the order ε^{m+1} .

The steps of the proof performed so far allow us to conclude that

$$(A.2) \quad \liminf_{\varepsilon \rightarrow 0} [a_2^\varepsilon(x_1^\varepsilon) + a_1^\varepsilon(x_2^\varepsilon)] \geq \liminf_{\varepsilon \rightarrow 0} [P(pX_2^\varepsilon \leq px_2^\varepsilon | X_1^\varepsilon = x_1^\varepsilon) + P(pX_1^\varepsilon \leq px_1^\varepsilon | X_2^\varepsilon = x_2^\varepsilon)] = 1,$$

where the last inequality follows from Lemma 4.1. On the other hand, by Lemma 4.4,

$$a_j^\varepsilon(x_i^\varepsilon) - \bar{s}_j(x_i^\varepsilon) \rightarrow 0 \quad \text{as } \varepsilon \rightarrow 0$$

so that, by (A.2), for the common limit x^* of x_1^ε and x_2^ε we must have

$$\bar{s}_2(x^*) + \bar{s}_1(x^*) \geq 1$$

since \bar{s}_j is continuous on $G_1^+ \cap G_2^+$. This conclusion, however, is inconsistent with the fact that $K \subseteq R^\alpha$. Thus, we have derived a contradiction and the proof is complete.

APPENDIX B

COMPARISON WITH HARSANYI'S MODEL

Here, we will analyze a hybrid model of payoff uncertainty which contains both Harsanyi's and our own model as special cases. To keep the computations simple, we restrict ourselves to a one-dimensional and symmetric set-up.

For $x \in \mathbb{R}$, let $g(x)$ be as in Figure 1. Let E_0, E_1 , and E_2 be independent random variables, each normally distributed with mean 0 and variance 1. Let $x^* \in (0, 4)$ and for $\delta, \varepsilon > 0$, define the random variable $X_i = X_i(x^*, \delta, \varepsilon)$ by means of

$$(B.1) \quad X_i = x^* + \delta E_0 + \varepsilon E_i, \quad i = 1, 2.$$

Moreover, define the incomplete information game $G(x^*, \delta, \varepsilon)$ by the following rules:

1. A realization (x_1, x_2) of (X_1, X_2) is drawn.
2. Player i is informed about x_i (and about nothing else).
3. The players choose actions $s = (s_1, s_2)$ simultaneously.
4. Player i receives payoff $g_i(x_i, s) = s_i x_i + 4(1 - s_i)(1 - s_j)$.

This model is a hybrid of ours and Harsanyi's: each player knows his own payoffs, but knowing the own payoffs gives information about the payoffs of the opponent. The case where $\delta = 0$ corresponds to Harsanyi's model, as players' payoffs are independent in this case. The case with fixed $\delta > 0$ and ϵ tending to zero corresponds more or less to our model. A difference is that now a player always knows his own payoffs exactly, which is not true in the model described in Section 3. However, it is easily checked that substituting the above condition 4 for the corresponding condition in the definition of a global game in Section 3 does not alter our main result. We will investigate the Bayesian equilibria of $G(x^*, \delta, \epsilon)$ and show that, in the first case, one indeed obtains results as in Harsanyi (1973), whereas, in the second case, one replicates the equilibrium selection result of this paper.

The game $G(x^*, \delta, \epsilon)$ is symmetric (if player i receives the information " $X_i = x$ " he is in exactly the same situation as player j receiving the information " $X_j = x$ ") so it is natural to look for symmetric equilibria. Actually, one can show that restricting attention to such equilibria, as we will do here, implies no loss of generality. Noting that player i will play α_i if $x_i > x$ and β_i if $x_i < x$, we will look for simple equilibria of the form

$$s_i(x_i) = \begin{cases} 0 & \text{if } x_i < x, \\ 1 & \text{if } x_i > x, \end{cases}$$

where $s_i(x_i)$ is the probability that i chooses α_i if his observation is x_i . The condition that player i is indifferent if $x_i = x$ may be written as

$$(B.2) \quad 4P(X_j < x | X_i = x) = x.$$

Now, conditional on X_i taking the value x , X_j is normally distributed with mean μ and standard deviation σ , where

$$(B.3) \quad \mu = \frac{\delta^2 x + \epsilon^2 x^*}{\delta^2 + \epsilon^2} \quad \text{and}$$

$$(B.4) \quad \sigma = \left[\frac{2\epsilon^2 \delta^2 + \epsilon^4}{\delta^2 + \epsilon^2} \right]^{1/2}.$$

Hence, (B.2) is equivalent to

$$\Phi \left[\frac{x - \mu}{\sigma} \right] = \frac{x}{4}$$

with Φ being the standard normal distribution function. The last equation may be written

$$(B.5) \quad \Phi(\eta(x - x^*)) = x/4$$

where

$$(B.6) \quad \eta = [\epsilon^2 / (2\delta^2 + \epsilon^2)(\delta^2 + \epsilon^2)]^{1/2}.$$

Note that the extreme case corresponding to Harsanyi's model ($\delta = 0, \epsilon \rightarrow 0$) can be approximated by choosing a sequence $(\delta, \epsilon) \rightarrow 0$ where $\delta \ll \epsilon$ and, thus, $\delta^2 / (\delta^2 + \epsilon^2) \rightarrow 0$ and $\eta \rightarrow \infty$. The global game model corresponds to the case where $\epsilon \rightarrow 0$ and $\epsilon \ll \delta$ so that $\delta^2 / (\delta^2 + \epsilon^2) \rightarrow 1$ and $\eta \rightarrow 0$. Obviously, equation (B.5) has at least one solution. Furthermore, by drawing a couple of graphs the reader may easily convince himself that there exists $\eta(x^*) > 0$ such that equation (B.5) has three solutions if $\eta > \eta(x^*)$ and a unique solution if $\eta < \eta(x^*)$. For $\eta > \eta(x^*)$, let $x^1(\eta) < x^2(\eta) < x^3(\eta)$ be the solutions. Then

$$(B.7) \quad x^1(\eta) \rightarrow 0 \quad \text{as } \eta \rightarrow \infty,$$

$$(B.8) \quad x^3(\eta) \rightarrow 4 \quad \text{as } \eta \rightarrow \infty \quad \text{and}$$

$$(B.9) \quad x^2(\eta) \rightarrow x^* \quad \text{with } \Phi(\eta(x^2(\eta) - x^*)) \rightarrow x^*/4 \quad \text{as } \eta \rightarrow \infty.$$

If the equilibrium of $G(x^*, \delta, \epsilon)$ corresponding to the switching point $x^1(\eta)$ is played, then, in the limit as $\eta \rightarrow \infty$, players will coordinate on the equilibrium α : Since $X_j \sim N(x^*, (\delta^2 + \epsilon^2)^{1/2})$ and $x^1(\eta) \ll x^*$, each player i chooses α_i with a probability that tends to 1 as $\eta \rightarrow \infty$. Similarly, the

switching point $x^3(\eta)$ induces beliefs corresponding to the equilibrium β as $\eta \rightarrow \infty$. If players adopt the switching point $x^2(\eta)$, then each player i will choose β_i with a probability

$$\Phi\left((x^2(\eta) - x^*) (\delta^2 + \varepsilon^2)^{-1/2}\right) = \Phi\left[\eta (x^2(\eta) - x^*) ((2\delta^2 + \varepsilon^2)/\varepsilon^2)^{1/2}\right],$$

which tends to $x^*/4$ as $\eta \rightarrow \infty$ in view of (B.9) and since $((2\delta^2 + \varepsilon^2)/\varepsilon^2)^{1/2} \rightarrow 1$. Hence, the beliefs associated with this equilibrium correspond to the mixed equilibrium of the game $g(x^*)$. Consequently, for $\eta \rightarrow \infty$, we replicate Harsanyi's results.

In the other extreme case where $\eta \rightarrow 0$ it is easily seen that $x(\eta) \rightarrow 2$ for the unique solution of (B.5). Hence, in the limit player i chooses α_i if $x_i > 2$ and β_i if $x_i < 2$. This choice is in agreement with equilibrium selection according to risk dominance and, thus, with the main result of the present paper. Note that in the special case where $x^* = 2$, the unique solution of (B.5) is given by $x(\eta) = x^*$. Hence, in this case, the beliefs generated by the Bayesian equilibrium do not converge to a Nash equilibrium of the limit game $g(x^*)$, but rather to the correlated equilibrium in which both α and β are played with probability $1/2$.

REFERENCES

- BINMORE, K. G. (1987): "Nash Bargaining and Incomplete Information," in *The Economics of Bargaining*, ed. by K. G. Binmore and P. Dasgupta. Oxford: Basil Blackwell.
- CARLSSON, H. (1989): "Global Games and the Risk Dominance Criterion," Mimeo, University of Lund.
- (1991): "A Bargaining Model where Parties Make Errors," *Econometrica*, 59, 1487–1496.
- CARLSSON, H., AND E. VAN DAMME (1989): "Global Payoff Uncertainty and Risk Dominance," Center Discussion Paper No. 8933.
- (1993): "Equilibrium Selection in Stag Hunt Games," in *Frontiers of Game Theory*, ed. by K. Binmore, A. Kirman, and P. Tani. Cambridge, MA: MIT Press.
- HARSANYI, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1–23.
- HARSANYI, J. C., AND R. SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- KIM, Y. (1992): "Adjustment Dynamics with Patient Players in Coordination Games," Mimeo, UCLA.
- KOHLBERG, E., AND J.-F. MERTENS (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–1039.
- MILGROM, P., AND J. ROBERTS (1989): "Adaptive and Sophisticated Learning in Repeated Normal Form Games," *Games and Economic Behavior*, 3, 82–100.
- MONDERER, D., AND D. SAMET (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170–190.
- NASH, J. F. (1950): "The Bargaining Problem," *Econometrica*, 18, 155–162.
- (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286–295.
- (1953): "Two-Person Cooperative Games," *Econometrica*, 21, 128–140.
- RUBINSTEIN, A. (1989): "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge,'" *American Economic Review*, 79, 385–391.
- SELTEN, R. (1975): "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 22–55.