

CARESS Working Paper #95-18
The Robustness of Equilibria to Incomplete Information[⌘]

Atsushi Kajii and Stephen Morris
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia PA 19104

March 1995
this version: November 1995

Abstract

A number of papers have shown that a strict Nash equilibrium action profile of a game may never be played if there is a small amount of incomplete information (see, for example, Carlsson and van Damme (1993a)). We present a general approach to analyzing the robustness of equilibria to a small amount of incomplete information. A Nash equilibrium of a complete information game is said to be robust to incomplete information if every incomplete information game with payoffs almost always given by the complete information game has an equilibrium which generates behavior close to the Nash equilibrium. We show that an open set of games has no robust equilibrium and examine why we get such different results from the refinements literature. We show that if a game has a unique correlated equilibrium, it is robust. Finally, a natural many-player many-action generalization of risk dominance is shown to be a sufficient condition for robustness.

[⌘] We are grateful to Eddie Dekel, Drew Fudenberg and George Mailath for valuable comments.

1. Introduction

Before Harsanyi's seminal contribution (Harsanyi (1967)), game theory was subject to the apparently compelling criticism that all conclusions relied completely on the assumption that payoffs were common knowledge. Harsanyi's formulation of games of incomplete information allowed analysis of situations where payoffs are not common knowledge. In this paper, we return to the classic question of how sensitive the conclusions of (complete information) game theory are to the common knowledge of payoffs assumption.

Suppose we fix a Nash equilibrium of a complete information game. Say that it is robust to incomplete information if behavior close to it is an equilibrium of every nearby incomplete information game. By "nearby" incomplete information game, we mean that the sets of players and actions are the same as in the complete information game and, with high probability, each player knows that his payoffs are the same. By "behavior close to it", we mean that the distribution over actions generated by the Nash equilibrium is close to the distribution over actions generated by an equilibrium of the nearby incomplete information game.

Although this definition has the flavor of the "refinements" literature, we show that it has very different properties. In particular, we show that there exists an open set of games which have a unique (strict) Nash equilibrium that is not robust. The argument is based on an "infection argument" similar to that of Rubinstein (1989): fix a complete information game and suppose that payoffs of an incomplete information game are always the same as the complete information game, except that at some information set of small probability, one player has a dominant strategy to play an action which is not part of the unique Nash profile. If there is an information set where a second player attaches high probability to the first player's information set, we can guarantee that the second player has a unique best response (not in the unique Nash profile) at that information set. We can iterate this argument to ensure that the unique Nash profile of actions is never played in any equilibrium of the incomplete information game. Since this can be done no matter how small the probability of the information set of the first player, the Nash equilibrium is shown not to be robust. We will discuss below why we get such different results from the refinements literature.

We present positive results which show two different kinds of sufficient conditions for robustness. The first comes from the observation that if an incomplete information game is near a complete information game, then any equilibrium of the incomplete information game generates a distribution over actions which is an approximate correlated equilibrium of the complete information game (the payoff uncertainty allows correlation of actions). Thus if a complete information game has a unique correlated equilibrium, then that equilibrium - which must be also a Nash equilibrium - is robust (by the upper hemicontinuity of the correlated equilibrium correspondence).

This is the first sufficient condition for robustness.

Our second, more interesting, sufficient condition works by showing a necessary connection between the ex ante probability of an event and the probability that players have certain higher order beliefs about that event. In order to develop this result, we first review some earlier work.

Following Morris, Rob and Shin (1995), say that an action profile $a = (a_1; \dots; a_I)$ in an I player game is a $(p_1; \dots; p_I)$ -dominant equilibrium if each action a_i is a best response to any conjecture putting probability at least p_i on other players choosing a_{-i} . Write $p = (p_1; \dots; p_I)$ for such a profile of probabilities. Following Monderer and Samet (1989), say that an event is p -evident if each individual i attaches probability at least p_i to the event whenever it is true. These definitions can be related together to prove a result about the equilibria of incomplete information games. Suppose that a complete information game has a p -dominant equilibrium a ; suppose also that an incomplete information game contains an event E which [1] is p -evident; and [2] has the property that payoffs are given by the complete information game at all states in E . Then the incomplete information game has an equilibrium where a is played at all states in E .

This result suggests a strategy for proving robustness. Suppose a complete information game has a p -dominant equilibrium a . If we could show that every nearby incomplete information game contains a high probability event E which satisfies the two properties cited above, we would be done. In fact, Monderer and Samet (1989) have provided an algorithm which - for any given information structure - finds the largest probability event E satisfying those two properties. Say that an event is p -believed if each individual i believes it with probability at least p_i . Say that an event is common p -belief if it is p -believed, it is p -believed that it is p -believed, etc... Then the largest event E satisfying the two properties above will be the set of states where it is common p -belief that payoffs are given by the complete information game.

But what is the connection between the probability of an event and the probability of the set of states where that event is common p -belief? We show that if $\prod_{i=1}^I p_i < 1$, then as the probability of an event tends to 1, the probability that the event is common p -belief tends to 1 uniformly across information systems. Conversely, if $\prod_{i=1}^I p_i > 1$, then it is possible to construct an information structure which has an event with probability arbitrarily close to 1, which is never common p -belief. This result shows a surprising necessary connection between the ex ante probability of an event and the ex ante probability of individuals having certain higher order beliefs. As an immediate corollary, we have that if an action profile a is a p -dominant equilibrium with $\prod_{i=1}^I p_i < 1$, then a is robust.

These sufficient conditions are rather strong. However, they can be used to give a complete characterization of robustness in one special but much studied class of games. In generic two player, two action games, there is always exactly one robust equilibrium: if the game has a unique pure strategy equilibrium, it is robust; if the game has two pure strategy equilibria, then the risk

dominant equilibrium (Harsanyi and Selten (1988)) is robust; if the game has no pure strategy equilibrium, then the unique mixed strategy equilibrium is robust.

What should we conclude about the robustness of game theoretic predictions to the common knowledge of payoffs assumption? We have shown that, in some contexts, an arbitrarily small probability that the analyst has mis-specified the payoffs of a game might imply very different behavior and outcomes from the game without the misspecification. In this sense, game theoretic predictions are not robust to the common knowledge of payoffs assumption. On the other hand, there are circumstances when there is a robust prediction.

These are the main results of the paper. We will now relate them back to the extensive related literature. First, there is the "trembles" literature. In that literature, perturbations have typically been either directly on players' action choices (the "trembles" approach) or on the payoffs of a complete information game. We, on the other hand, perturb payoffs indirectly, via the information structure. However, these qualitative differences do not account for why we get such different results - in particular, the result that even unique strict Nash equilibria need not be robust. The differences arise because of the richness of the perturbed games which we consider. Thus we show (in section 7) that if we restricted attention to nearby incomplete information games with either bounded state spaces or independent signals about payoffs, then all strict equilibria would be robust.

Our work follows Fudenberg, Kreps and Levine (1988) and Dekel and Fudenberg (1990) in studying the robustness of game theoretic predictions (they considered, respectively, strict equilibrium and iterated deletion of weakly dominated strategies) when there is a small amount of incomplete information. Our technique of "embedding" a complete information game in "nearby" incomplete information game closely follows theirs. But they tested whether an outcome can be justified by some sequence of perturbed games, while we require robustness to all such sequences. Thus our work relates to their work as stability type refinements relate to perfection type refinements.¹

Our work builds most strongly on a literature relating higher order beliefs to the equilibria of incomplete information games, especially Monderer and Samet (1989) and Morris, Rob and Shin (1995). While we use extensively techniques and results from those papers, we go one step further. Instead of making assumptions about players' higher order beliefs about payoffs, we make assumptions about the ex ante probabilities of payoffs and deduce the required properties of players' higher order beliefs, which can then be used to characterize equilibria.

The robustness question which we study complements the main question previously studied

¹Another difference is that infinite state space and independence assumptions in those papers ensure that allowing only a small amount of ex ante uncertainty about payoffs implies that, with high probability, payoffs are common p-belief in the sense of Monderer and Samet (1989), for some p close to 1 (thus Bärgers (1994) establishes essentially the same conclusion as Dekel and Fudenberg (1990) by directly assuming that payoffs are common p-belief, for p close to 1). We will allow infinite state spaces and correlated signals where such equivalences do not work.

in the higher order beliefs literature. Given that some Nash equilibria are not robust to the inclusion of a small amount of incomplete information, there are two natural questions to ask. First, which Nash equilibria of a complete information game can always be played in equilibria of nearby incomplete information games. This is the question addressed in this paper. Second, how must the notion of a "nearby" incomplete information games be strengthened in order to ensure that, for every Nash equilibrium of a complete information game, similar behavior is generated by an equilibrium of every nearby incomplete information game. This lower hemicontinuity question has been studied in the literature.² Say that an incomplete information game is nearby a complete information game if, with high probability, it is common p -belief - with p close to 1 - that payoffs are those of the complete information game. This notion is sufficient to ensure that all Nash equilibria are robust in the sense described above. This result is implied by Monderer and Samet (1989). Monderer and Samet (1990) and Kajii and Morris (1994b) can be interpreted as showing (in different settings) that this notion of closeness is necessary and sufficient for the lower hemicontinuity.

Finally, Carlsson and van Damme (1993a) have considered a closely related robustness question. They suppose that each player of a two player, two action game observes a noisy signal of the payoffs. They show that as the noise goes to zero, the unique (Bayesian Nash) equilibrium has the risk dominant (Nash) equilibrium of the complete information game being played. Similar techniques also work in some many player settings - see Carlsson and van Damme (1993b) and Kim (1993). There are many reasons why our results are not directly comparable; but our work should be seen as replicating some of their results in a different setting, as well as providing new results. The key difference is that while Carlsson and van Damme consider a certain critical class of payoff perturbations with continuous signals, we characterize robustness to all perturbations with countable state spaces. The relation is discussed in more detail in the appendix (section 9.4).

The paper is organized as follows. In section 2, we present our approach to embedding complete information games in nearby incomplete information games and define the notion of robustness. In section 3, we show that a unique correlated equilibrium must be robust but that a (strict) Nash equilibrium need not be. In section 4, we review results on belief operators and common p -belief and present new results on ex ante probabilities and higher order beliefs. In section 5, these results are applied to give further positive robustness results. In section 6, we provide a complete characterization of robustness for two player, two action games. In section 7, we discuss alternative notions of robustness and the relation to refinements. Section 8 concludes.

²For non-strict Nash equilibria, it is necessary to weaken the solution concept in the incomplete information game to interim ϵ -equilibrium; that is, each player's payoff conditional on each information set must be within ϵ of the best response.

2. Framework

2.1. Complete Information Games

Throughout this analysis, we fix a complete information game G consisting of a finite collection of players $I = \{1, \dots, I\}$ and, for each player i , a finite action set A_i and payoff function $g_i : A \rightarrow \mathbb{R}$, where $A = A_1 \times \dots \times A_I$. Thus $G = (I; \{A_i\}_{i \in I}; \{g_i\}_{i \in I})$. We shall denote $\prod_{j \in I} A_j$ by A_{-i} and a generic element of A_{-i} by a_{-i} . Similar conventions will be used whenever it is clear from the context.

For any finite set S , denote by $\Phi(S)$ the set of all probability measures on S .

Definition 2.1. An action distribution, $\mu \in \Phi(A)$, is a correlated equilibrium of G if, for all $i \in I$ and $a_i^0 \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} g_i(a_i; a_{-i}) \mu(a_i; a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} g_i(a_i^0; a_{-i}) \mu(a_i^0; a_{-i})$$

An action distribution μ is a Nash equilibrium if it is a correlated equilibrium and, for all $a \in A$,

$$\mu(a) = \prod_{i \in I} \mu_i(a_i)$$

where $\mu_i \in \Phi(A_i)$ is the marginal distribution of μ on A_i .

This indirect way of defining (mixed strategy) Nash equilibrium is equivalent to the standard one.

2.2. Embedding Complete Information Games in Incomplete Information Games

We will require a way of comparing complete and incomplete information games. An incomplete information game is said to "embed" the complete information game $G = (I; \{A_i\}_{i \in I}; \{g_i\}_{i \in I})$ if it has the same sets of players and actions.

For our purposes, an incomplete information game U consists of (1) the collection of players, $I = \{1, \dots, I\}$; (2) their action sets, A_1, \dots, A_I ; (3) a countable state space, Ω ; (4) a probability measure on the state space, P ; (5), for each player i , a partition of the state space, Q_i ; and (6), for each player i , a state dependent payoff function, $u_i : A \times \Omega \rightarrow \mathbb{R}$. Thus $U = (I; \{A_i\}_{i \in I}; \Omega; P; \{Q_i\}_{i \in I}; \{u_i\}_{i \in I})$. We write $P(\omega)$ for the probability of the singleton event $\{\omega\}$ and $Q_i(\omega)$ for the (unique) element of Q_i containing ω . Throughout the paper we will restrict attention to incomplete information games where every information set of every player is possible, that is $P[Q_i(\omega)] > 0$ for all $i \in I$ and $\omega \in \Omega$. Under this assumption the conditional probability of state ω given information set $Q_i(\omega)$, written $P[\omega | Q_i(\omega)]$, is well-defined by the

rule $P(\cdot | jQ_i(\cdot)) = \frac{P(\cdot)}{P(Q_i(\cdot))}$. If U satisfies all the above properties, we say that U embeds G ; we write $E(G)$ for the set of incomplete information games which embed G :

A (mixed) strategy for player i is a Q_i -measurable function $\sigma_i : \Omega \rightarrow \Delta(A_i)$. We denote by $\sigma_i(a_j | \cdot)$ the probability that action a_i is chosen given \cdot under σ_i . A strategy profile is a function $\sigma = (\sigma_i)_{i \in I}$ where σ_i is a strategy for player i . We write S for the collection of such strategy profiles. We denote by $\sigma(a_j | \cdot)$ the probability that action profile a is chosen given \cdot under σ ; we write σ_i for $(\sigma_j)_{j \in I}$; when no confusion arises, we extend the domain of each u_i to mixed strategies and thus write $u_i(\sigma(\cdot); \cdot)$ for $\sum_{a \in A} u_i(a; \cdot) \sigma(a_j | \cdot)$. Now the payoff of strategy profile σ to player i is given by the expected utility $\sum_{\omega \in \Omega} u_i(\sigma(\omega); \omega) P(\omega)$ which can also be written as $\sum_{\omega \in \Omega} u_i(\sigma(\omega); \omega) P(\omega)$.

Definition 2.2. A strategy profile σ is a Bayesian Nash equilibrium of U if, for all $a_i \in A_i$ and $\omega \in \Omega$,

$$\sum_{\omega \in Q_i(\omega)} u_i(\sigma_i(\omega); \omega) P(\omega | jQ_i(\omega)) \geq \sum_{\omega \in Q_i(\omega)} u_i(a_i; \sigma_{-i}(\omega); \omega) P(\omega | jQ_i(\omega))$$

A strategy profile σ specifies the probability of a given action profile being played at a given state. We will be interested in a reduced form representation of the strategy profile, where we only report the ex ante probability of certain actions being played.

Definition 2.3. An action distribution, $\lambda \in \Delta(A)$, is an equilibrium action distribution of U if there exists a Bayesian Nash equilibrium σ of U such that $\lambda(a) = \sum_{\omega \in \Omega} \sigma(a_j | \omega) P(\omega)$.

2.3. Robustness

We want to formalize the idea that an incomplete information game U is close to a complete information game G if the payoff structure under U is equal to that under G with high probability. Thus, for each incomplete information game $U \in E(G)$, write Ω_U for the set of states where payoffs are given by G , and every player knows his payoffs:

$$\Omega_U = \{\omega \in \Omega : u_i(a; \omega) = g_i(a) \text{ for all } a \in A, \omega \in Q_i(\omega), \text{ and } i \in I\}$$

Definition 2.4. The incomplete information game U is an ϵ -elaboration of G if $U \in E(G)$ and $P(\Omega_U) = 1 - \epsilon$. Let $E(G; \epsilon)$ be the set of all ϵ -elaborations of G .

To clarify the idea of ϵ -elaborations, consider 0-elaborations. The degenerate 0-elaboration is the game where Ω is a singleton set, and the original complete information game is played with

probability 1. The set of equilibrium action distributions of the degenerate incomplete information game is just the set of Nash equilibria of the complete information game. But 0-elaborations can entail more complicated information structures that allow players to correlate their actions. Indeed, an action distribution is an equilibrium action distribution of some 0-elaboration of G if and only if it is correlated equilibrium of G (see Aumann (1987)). Thus while two 0-elaborations may appear close to an outside observer, it should be clear that they are not close in any game theoretic sense.

We will measure the distance between action distributions by the max norm:

$$k^1_j \text{ }^\circ k = \text{Max}_{a \in A} |j^1(a) - j^\circ(a)|:$$

Deñition 2.5. An action distribution μ^1 is robust to incomplete information in G if, for every $\epsilon > 0$, there exists $\delta > 0$ such that, for all $\mu \in \mathcal{U}$, every $U \in E(G; \mu)$ has an equilibrium action distribution μ° with $k^1_j \text{ }^\circ k \leq \epsilon$.

In words, an action distribution μ^1 is robust if every slightly perturbed game has an equilibrium action distribution that is close to μ^1 . Notice that if μ^1 is robust to incomplete information, it is a Nash equilibrium since the degenerate 0-elaboration belongs to $E(G; 0)$.

Remark 1. It is important that we allow payoffs under an μ -elaboration U to be very different from G outside the set μ . Quite different results would follow if, in addition, we required the payoffs outside μ to be uniformly within δ of G.

Remark 2. For simplicity, our deñition focuses on elaborations where, with high probability, each player knows that his payoffs are given by G. Essentially the same results would follow if we replaced "knows" by "believes with high probability," as long as we imposed a bound on payoffs in the elaboration (see section 9.4.2 in the appendix). Players' uncertainty about their own payoff function is orthogonal to the issues considered in this paper.

Remark 3. An alternative formulation would require every $U \in E(G; \mu)$ to have an equilibrium μ° with ex ante payoffs within ϵ of payoffs under μ^1 . If, in addition, we bounded payoffs outside μ , this would give the same results.

Remark 4. The deñition requires a property uniformly over μ -elaborations, as we vary all characteristics of the incomplete information games. It will be clear from the arguments which follow that we could allow δ to depend on the state space and information partitions as long as we still required uniformity with respect to probability distributions on that state space. There is further discussion of this issue in the appendix (section 9.3).

Remark 5. In the related work of Monderer and Samet (1989), allowing for interim ϵ -equilibrium in the incomplete information game makes a large difference to the result. We will note when we give our non-existence example that this would not make a difference to our exercise.

Remark 6. We allow players to correlate their actions via information in the elaborations. Yet we require (in our definition of robustness) that an action distribution be (nearly) played in an equilibrium of every nearby elaboration, including degenerate ones where no correlation is possible. A more reasonable definition might allow players access to payoff-irrelevant randomizing devices with private signals, uncorrelated with the states of the elaborations (see Cotter (1991)). Thus say that $\sigma \in \Sigma(A)$ is a correlated equilibrium action distribution of U if there exists a correlated equilibrium $\sigma \in \Sigma(S)$ of U with

$$\sigma(a) = \sum_{s \in S} \sigma(s) P(a|s)$$

and say that $\sigma \in \Sigma(A)$ is robust with correlation to incomplete information if it satisfies definition 2.5 with "equilibrium action distribution" replaced by "correlated equilibrium action distribution". By definition, any robust σ is also robust with correlation; thus our positive results would continue to hold with this definition. We will verify that our negative result also holds with this definition.

3. The Robustness of Unique Nash Equilibria

In this section, we consider complete information games which have a unique Nash equilibrium and examine when it is robust. Intuitively, this should be the easiest setting in which demonstrate robustness. However, we first provide an open class of games with a unique (strict) Nash equilibrium which is not robust.

3.1. Non-existence

Example 3.1. The Cyclic Matching Pennies Game. Consider the following game G . There are 3 players and each player has three possible actions: Heads (H), Tails (T) and Safe (S). The following tables of payoffs show 1's action on the row, 2's action on the column. Now if 3 chooses Heads, payoffs are:-

	H	T	S
H	1; 1; 1	1; 1; 1	1; $\frac{1}{4}$; 0
T	1; 1; 1	1; 1; 1	1; $\frac{1}{4}$; 0
S	$\frac{1}{4}$; 0; 1	$\frac{1}{4}$; 0; 1	$\frac{1}{4}$; $\frac{1}{4}$; 0

if 3 chooses Tails, payoffs are:-

	H	T	S
H	1; $\frac{1}{4}$; 1	1; 1; $\frac{1}{4}$	1; $\frac{1}{4}$; 0
T	$\frac{1}{4}$; 1; 1	$\frac{1}{4}$; 1; $\frac{1}{4}$	$\frac{1}{4}$; 1; 0
S	$\frac{1}{4}$; 0; 1	$\frac{1}{4}$; 0; 1	$\frac{1}{4}$; $\frac{1}{4}$; 0

if 3 chooses Safe payoffs are:-

	H	T	S
H	0; $\frac{1}{4}$; $\frac{1}{4}$	0; 1; $\frac{1}{4}$	0; $\frac{1}{4}$; $\frac{1}{4}$
T	0; 1; $\frac{1}{4}$	0; $\frac{1}{4}$; $\frac{1}{4}$	0; $\frac{1}{4}$; $\frac{1}{4}$
S	$\frac{1}{4}$; 0; $\frac{1}{4}$	$\frac{1}{4}$; 0; $\frac{1}{4}$	$\frac{1}{4}$; $\frac{1}{4}$; $\frac{1}{4}$

The game has the following interpretation. Each player's payoff depends only on his own action and the action of his "adversary". Adversaries are determined by the cycle 1 → 3 → 2 → 1, so that 3 is 1's adversary etc.. Thus, for example, 1's payoff is completely independent of 2's action. Each player has a safe action under which he gets $\frac{1}{4}$ (independent of his adversary's action). If he does not play his safe action, then he is playing a cyclic matching pennies game, where he tries to choose the face of the coin different from his adversary's. Thus player 2 gets 1 if he doesn't match 1's choice, $\frac{1}{4}$ otherwise. Player 3 gets 1 if he doesn't match 2's choice, $\frac{1}{4}$ otherwise. Player 1 gets 1 if he doesn't match 3's choice, $\frac{1}{4}$ otherwise. Each player gets 0 if his adversary chooses a safe action and he does not.

Note for future reference that if any player puts probability strictly greater than $\frac{5}{8}$ on his adversary choosing H or T, he has a strict best response to do the opposite (and not play S). Thus if 1 thinks that player 3 will play H with probability $q > \frac{5}{8}$ his payoff to T is at least $q + (1 - q) = 2q - 1 > \frac{1}{4}$, while the payoff to H is at most $q + (1 - q) = 1 - 2q < \frac{1}{4}$, while the payoff to S is $\frac{1}{4}$.

This game has a unique Nash equilibrium where all players choose S. To see why, first suppose that 1 plays H in equilibrium with positive probability, but not T. Then 2 never plays H. So 3 never plays T. So 1 never plays H, a contradiction. Suppose now that 1 plays both H and T in equilibrium. Then 3 must play H and T with equal probability. So the payoff to 1 of playing H and T is 0 which is strictly less than $\frac{1}{4}$, the payoff to playing S. So we have another contradiction.

Now consider the following "elaboration" of G. Let \mathbb{N} be the set of non-negative integers; $P(\mathbb{N}) = \{ (1 - \epsilon)^i \epsilon^j \mid i, j \in \mathbb{N} \}$. Partitions are given by $Q_1 = (f_0; 1; 2g; f_3; 4; 5g; \dots)$; $Q_2 = (f_0g; f_1; 2; 3g; f_4; 5; 6g; \dots)$ and $Q_3 = (f_0; 1g; f_2; 3; 4g; f_5; 6; 7g; \dots)$. Let $u_i(a; \mathbb{N}) = g_i(a)$ unless either $i = 3$ and $\mathbb{N} \geq f_0; 1g$ or $i = 2$ and $\mathbb{N} = 0$. Let $u_3((a_1; a_2; H); \mathbb{N}) = 1$, $u_3((a_1; a_2; T); \mathbb{N}) = 0$ and $u_3((a_1; a_2; S); \mathbb{N}) = 0$ for all a_1, a_2 and $\mathbb{N} \geq f_0; 1g$; and let $u_2((a_1; T; a_3); 0) = 0$,

$u_2((a_1; H; a_3); 0) = 0$ and $u_3((a_1; S; a_3); 0) = 0$ for all a_1, a_3 . Thus player 3 has a dominant strategy to play H at information set $f_0; 1g$, player 2 has a dominant strategy to play T at information set f_0g , and payoffs are given by G everywhere else. Note that $-u = f_2; 3; 4; \dots; g$, so $P[-u] = 1; 1; (1; \dots)^{\frac{1}{2}}; 1; 1; (1; \dots)^{\frac{1}{2}}; (1; \dots)^{\frac{1}{2}} = 1; 1; \dots$.

Now at information set $f_0; 1; 2g$, 1 assigns probability at least $\frac{2}{3}$ to 3 playing H, so 1 must play T. But by a similar argument, 2 must then play H at $f_1; 2; 3g$, 3 must play T at $f_2; 3; 4g$, 1 must play H at $f_3; 4; 5g$ and so on. Safe is played nowhere. Thus for any $\epsilon > 0$, we can construct an ϵ -elaboration where the unique Nash equilibrium is never played in the unique Bayesian Nash equilibrium. The following table describes the actions in that unique Bayesian Nash equilibrium.

!	0	1	2	3	4	5	6	...
1's action	T	T	T	H	H	H	T	...
2's action	T	H	H	H	T	T	T	...
3's action	H	H	T	T	T	H	H	...

This negative example does, however, suggest a strategy for identifying robust equilibria in other games. The unique Bayesian Nash equilibrium $\frac{3}{4}$ and the unique equilibrium action distribution $\frac{1}{6}$ it implies are summarized in the following:

Action Profile: a	$f! : \frac{3}{4}$ (aj!) = 1g	$\frac{1}{6}$ (a)
(T; T; H)	$f_0; 6; 12; \dots; g$	$\frac{1}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(T; H; H)	$f_1; 7; 13; \dots; g$	$\frac{(1; \dots)^{\frac{1}{2}}}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(T; H; T)	$f_2; 8; 14; \dots; g$	$\frac{(1; \dots)}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(H; H; T)	$f_3; 9; 15; \dots; g$	$\frac{(1; \dots)^{\frac{3}{2}}}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(H; T; T)	$f_4; 10; 16; \dots; g$	$\frac{(1; \dots)^2}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(H; T; H)	$f_5; 11; 17; \dots; g$	$\frac{(1; \dots)^{\frac{5}{2}}}{1+(1; \dots)^{\frac{1}{2}}+(1; \dots)+(1; \dots)^{\frac{3}{2}}+(1; \dots)^2+(1; \dots)^{\frac{5}{2}}}$
(H; H; H)	;	0
(T; T; T)	;	0

As $\epsilon \rightarrow 0$ (and thus $P[-u] \rightarrow 1$), the equilibrium action distribution converges to a correlated equilibrium of G with each of $f_H; H; Tg, f_H; T; Hg, f_H; T; Tg, f_T; H; Hg, f_T; H; Tg, f_T; T; Hg$ played with probability $\frac{1}{6}$ (thus the bad outcomes $f_H; H; Hg$ and $f_T; T; Tg$ are avoided). Each player's expected payoff is $\frac{1}{3}$ in this correlated equilibrium. We will see that convergence to a correlated equilibrium is going to be a general property, which we will be able to exploit in

providing our first positive robustness result in the next section. But first we will verify that the selection of this equilibrium is robust to a number of features of the construction.

- ² In the above argument, each optimal action gives a payoff of at least $\frac{1}{3}$, and thus gives a payoff at least $\frac{1}{12}$ higher than the next best action. Thus the table describes not only the unique equilibrium but also the unique interim $\frac{1}{12}$ -equilibrium. Thus allowing interim ϵ -equilibria in our definition of robustness would not help in this example.
- ² The above argument also shows that the strategies identified are the unique (incomplete information game) strategies which survive iterated deletion of strictly dominated strategies. This ensures that the unique correlated equilibrium of the incomplete information game entails each player choosing those strategies with probability one. Thus if any action distribution is going to be robust with correlation (see definition in remark 6 on page 9), it must be the correlated equilibrium identified above. But notice that we could easily alter the above sequence of ϵ -elaborations to ensure a different limit. Suppose that (for each ϵ) we increased the probability of states 0, 6, 12, etc... by a small amount while decreasing the probability of all other states. If the change was not too large, the same strategy profile would remain the unique one surviving iterated deletion. But the limit equilibrium action distribution would put strictly higher probability on action profile (T; T; H) and strictly lower probability on the others. Thus the cyclic matching pennies game also has no correlated equilibrium which is robust with correlation to incomplete information.
- ² If we perturbed payoffs slightly, the argument (relying as it does on iterated deletion of strictly dominated strategies) would go through unchanged. Thus the set of complete information games where there is no robust equilibrium is open in the set of payoff matrices.

3.2. A Sufficient Condition for Robustness

The previous discussion suggests that the limit of equilibrium action distributions of a sequence of ϵ -elaborations of a complete information game must be a correlated equilibrium. This gives the following:

Proposition 3.2. (Unique Correlated Equilibrium). If G has a unique correlated equilibrium σ^* , then σ^* is robust.

A related result appears in Shin and Williamson (1994). The condition that there is a unique correlated equilibrium is very strong but far from vacuous. If a two player, two action game has no pure strategy Nash equilibrium, then the unique mixed strategy equilibrium is the unique correlated equilibrium (see section 6). Dominance solvable games (where a unique action profile survives iterated deletion of strictly dominated strategies) have unique correlated equilibria; so do

two player zero sum games with a unique Nash equilibrium (Aumann (1987) and Forges (1990)). Neyman (1991) and Cripps (1994) give sufficient conditions for all correlated equilibria to be convex combinations of Nash equilibria, and thus for uniqueness of Nash equilibrium to imply uniqueness of correlated equilibrium.

The proof of proposition 3.2 requires some preliminary definitions and lemmas.

Definition 3.3. Action distribution σ is an ϵ -correlated equilibrium of G if for all $i \in I$ and $f : A_i \rightarrow A_i$,

$$\sum_{a \in A} \sigma_i(a) \geq \sum_{a \in A} \sigma_i(f(a_i); a_{-i}) - \epsilon \sum_{a \in A} \sigma_i(a)$$

Note that 0-correlated equilibrium is equivalent to the definition of correlated equilibrium (definition 2.1) given on page 6.

Lemma 3.4. For any complete information game G and $\epsilon > 0$, there exists $\delta > 0$ such that every equilibrium action distribution of every δ -elaboration of G with $\delta \leq \delta$ is an ϵ -correlated equilibrium.

Proof. Let U be an δ -elaboration of G and let σ be any Bayesian Nash equilibrium of U . Let $M = \max_{a \in A} g_i(a) - \min_{a \in A} g_i(a)$. Let σ be the induced action distribution; that is, $\sigma_i(a) = \sum_{j \in J} \sigma_i(a, j) P(j)$. Fix $f : A_i \rightarrow A_i$ and let $\sigma_i^0(a, j) = \sum_{i' \in I} \sigma_i^0(a_i', j)$. By the assumption

that σ is an equilibrium, we have, for all $i \in I$,

$$\sum_{j \in J} \sum_{a \in A} \sigma_i(a, j) u_i(a, j) - \sum_{j \in J} \sum_{a \in A} \sigma_i^0(a, j) u_i(a, j) \leq \epsilon \sum_{j \in J} P(j) Q_i(j) \leq 0 \quad (3.1)$$

Let $u_i(a, j) = f_j(a; j) = g_i(a)$ for all $a \in A$ and $j \in J$. Since $u_i(a, j) = P(j) Q_i(j)$, $\sum_{j \in J} P(j) Q_i(j) = 1$. Thus by equation (3.1) and the definition of u_i , we have, for all $i \in I$,

$$\sum_{j \in J} \sum_{a \in A} \sigma_i(a, j) g_i(a, j) - \sum_{j \in J} \sum_{a \in A} \sigma_i^0(a, j) g_i(a, j) \leq \epsilon \sum_{j \in J} P(j) Q_i(j) \leq 0 \quad (3.2)$$

and so

$$\sum_{j \in J} \sum_{a \in A} \sigma_i(a, j) g_i(a, j) - \sum_{j \in J} \sum_{a \in A} \sigma_i^0(a, j) g_i(a, j) \leq \epsilon \sum_{j \in J} P(j) Q_i(j) \leq 0$$

But

$$\sum_{j \in J} \sum_{a \in A} \sigma_i(a, j) g_i(a, j) - \sum_{j \in J} \sum_{a \in A} \sigma_i^0(a, j) g_i(a, j) \leq \epsilon \sum_{j \in J} P(j) Q_i(j) \leq \epsilon M$$

information game U^α has a unique Bayesian Nash equilibrium where player 1 plays H (with probability 1) at information sets $f_0; 1g, f_4; 5g, \text{ etc...}$; while he plays T at information sets $f_2; 3g, f_6; 7g, \text{ etc...}$; similarly, player 2 plays H at information sets $f_0g, f_3; 4g, \text{ etc...}$; and he plays T at information sets $f_1; 2g, f_5; 6g, \text{ etc...}$. Thus the unique equilibrium action distribution, 1^α , of U^α has $1^\alpha(H; H) = \frac{\alpha}{1_i(1_i - \alpha)^4}$, $1^\alpha(H; T) = \frac{\alpha(1_i - \alpha)}{1_i(1_i - \alpha)^4}$, $1^\alpha(T; T) = \frac{\alpha(1_i - \alpha)^2}{1_i(1_i - \alpha)^4}$ and $1^\alpha(T; H) = \frac{\alpha(1_i - \alpha)^3}{1_i(1_i - \alpha)^4}$. We can verify that $1^\alpha(a) \neq 1^\alpha(a) = \frac{1}{4}$ as $\alpha \neq 0$, for all $a \in A$. This is also an implication of corollary 3.5. But notice that for every α and at every state, the unique equilibrium behavior is a pure strategy and thus very different from the unique Nash equilibrium strategy.

4. p-Belief

As noted in the introduction, our second positive robustness result uses new results about the ex ante probability of higher order belief events. Thus, in this section, we introduce belief operators and common p-belief (in section 4.1) and report the critical path result which shows that for sufficiently small p, there is a connection between ex ante probability and statements about higher order beliefs (in sections 4.2 and 4.3). In section 5, we will use this result in generating a positive robustness result.

4.1. Common p-Belief

Fix the information system part of an incomplete information game, i.e. $\{I; -; P; fQ_i g_{i \in I}\}$. We maintain the assumption that $P[Q_i(!)] > 0$, for all $i \in I$ and $! \in \Omega$. For any number $p_i \in (0; 1]$, define

$$B_i^{p_i}(E) = \{! : P[E|Q_i(!)] \geq p_i\}$$

That is, $B_i^{p_i}(E)$ is the set of states where player i believes E with probability at least p_i . For any row vector $p = (p_1; \dots; p_I) \in (0; 1]^I$, $B^p(E) = \bigcap_{i \in I} B_i^{p_i}(E)$; $B^p(E)$ is the set of states where E is p-believed, i.e. each player i believes E with probability at least p_i . An event is said to be common p-belief if it is p-believed, it is p-believed that it is p-believed, etc...; thus E is common p-belief at state $!$ if $! \in C^p(E) = \bigcap_{n=1} B^n(E)$. Monderer and Samet (1989) introduced such belief operators and characterized common p-belief for symmetric p, i.e. $C^{(p; \dots; p)}(E)$. Their results remain true essentially as stated in the case of asymmetric p. This section reports the trivial extension.

Write F_i for the σ -field generated by Q_i and say that event E is simple if $E = \bigcup_{i \in I} E_i$, each $E_i \in F_i$. We will use the following straightforward properties of belief operators which we state without proof.

Fact 1. If $E_i \in F_i$, then $B_i^{p_i}(E_i) = E_i$.

Fact 2. For all events $E \subseteq F$, $B_i^{p_i}(E) \subseteq B_i^{p_i}(F)$.

Fact 3. For all events E , $B_i^{p_i}(E) \subseteq F_i$.

Fact 4. For all events E , $B_i^p(E)$ is simple.

Fact 5. If E is simple, then $B_i^p(E) \subseteq E$.

Fact 6. If E^n is decreasing, then $B_i^{p_i}(\bigcap_{n=1}^{\infty} E^n) = \bigcap_{n=1}^{\infty} B_i^{p_i}(E^n)$.

Definition 4.1. E is p -evident if $E \subseteq B_i^p(E)$.

The following result follows Monderer and Samet (1989).

Theorem 4.2. (The Common p -Belief Theorem). E is common p -belief at i (i.e. $i \in C^p(E)$) if and only if there is a p -evident event F with $i \in F \subseteq B_i^p(E)$.

Proof. By properties 4 and 5, $B_i^{p_i^n}(E)$ is decreasing in n (for $n \geq 1$). Now for each $i \in I$ and $n \geq 1$,

$$C^p(E) \subseteq [B_i^p]^{n+1}(E) \subseteq B_i^{p_i}([B_i^p]^n(E)).$$

Thus by fact 6,

$$C^p(E) \subseteq \bigcap_{n=1}^{\infty} B_i^{p_i}([B_i^p]^n(E)) = B_i^{p_i}(\bigcap_{n=1}^{\infty} [B_i^p]^n(E)) = B_i^{p_i}(C^p(E)).$$

Thus $C^p(E)$ is p -evident; since $C^p(E) \subseteq B_i^p(E)$, the "only if" part of the proposition is satisfied setting $F = C^p(E)$.

Conversely, suppose F is p -evident, $i \in F$ and $F \subseteq B_i^p(E)$. By fact 2, we have for all $n \geq 0$,

$$[B_i^p]^n(F) \subseteq [B_i^p]^{n+1}(E):$$

By F p -evident and fact 2, we have for all $n \geq 1$, $F \subseteq B_i^{p_i^n}(F)$. Thus we have

$$i \in F \subseteq \bigcap_{n=0}^{\infty} [B_i^p]^n(F) \subseteq \bigcap_{n=1}^{\infty} [B_i^p]^n(E) = C^p(E): \quad \square$$

4.2. The Critical Path Result

What is the connection between the (ex ante) probability of an event E and the ex ante probability of the event $C^P(E)$? The following proposition shows that if $\prod_{i \in I} p_i < 1$, then $P[C^P(E)]$ is close to 1 whenever $P[E]$ is close to 1, regardless of the state space.

Proposition 4.3. (The Critical Path Result). If $\prod_{i \in I} p_i < 1$, then in any information system $\mathcal{I}; -; P; fQ_i g_{i \in I}$, all simple events E satisfy:

$$P[C^P(E)] \geq \prod_{i \in I} (1 - P[E])^{1 - \min_{i \in I} p_i}$$

This result is tight in the following two senses³:

1. If $\prod_{i \in I} p_i < 1$ and $\prod_{i \in I} (1 - \min_{i \in I} p_i) < 1$, then there exists an information system $\mathcal{I}; -; P; fQ_i g_{i \in I}$ and a simple event E , with $\prod_{i \in I} P[E] = \prod_{i \in I} (1 - \min_{i \in I} p_i)$ and $P[C^P(E)]$ arbitrarily close to $\prod_{i \in I} (1 - \min_{i \in I} p_i)$.
2. If $\prod_{i \in I} p_i \geq 1$ and $q < 1$, then there exists an information system $\mathcal{I}; -; P; fQ_i g_{i \in I}$ and a simple event E , with $P[E] = q$ and $C^P(E) = \emptyset$.

Proposition 4.3 is of some interest beyond its use in this paper, so before giving a proof, let us make a few observations about the proposition. First, as Monderer and Samet (1989) observed, the characterization of common p -belief described in the previous sub-section was independent of the common prior assumption. That is, if we endowed each individual with a different prior measure P_i , and defined belief operators by $B_i^{p_i}(E) = \{ \omega : P_i[E | Q_i(\omega)] \geq p_i \}$, theorem 4.2 and its proof would remain unchanged. Indeed, the game theoretic results which Monderer and Samet prove using common p -belief would also be unchanged. By contrast, proposition 4.3 relies on the common prior assumption. Suppose $I = \{1, 2, \dots, N\}$;

$$P_1(\omega) = \begin{cases} \frac{1}{N}, & \text{if } \omega \text{ is odd} \\ \frac{1}{N}, & \text{if } \omega \text{ is even} \end{cases}; \quad P_2(\omega) = \begin{cases} \frac{1}{N}, & \text{if } \omega \text{ is odd} \\ \frac{1}{N}, & \text{if } \omega \text{ is even} \end{cases}$$

³This can be shown by construction; example 9.9 on page 44 in the appendix shows the second form of tightness.

$Q_1 = (f_1; 2g; f_3; 4g; \dots; f_{2N}; 1; 2Ng)$ and $Q_2 = (f_1g; f_2; 3g; \dots; f_{2N}g)$. Consider the event $E = f_3; 4; \dots; 2Ng$; $P_1[E] = P_2[E] = \frac{N-1}{N}$, so we can make both ex ante probabilities of event E arbitrarily close to 1 by choosing N sufficiently large. But $C^P(E) = \emptyset$, for all p with $p_1 > \frac{1}{2}$ and $p_2 > \frac{1}{2}$.

Second, note that while proposition 4.3 concerns the ex ante probability of events, it can be used to prove results about the existence of certain events as follows; write $\#Q_i$ for the (perhaps infinite) number of elements of i 's information partition.

Corollary 4.4. Fix any information system $\mathcal{I}; \dots; P; f; Q_i; g; i \in I$. If $\sum_{i \in I} p_i < 1$ and $\#Q_j > \frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)}$ for some $j \in I$, then there exists an event E such that $C^P(E) \neq \emptyset$.

Proof. Since $\#Q_j > \frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)}$ and each element of Q_j has positive probability, some element $F \in Q_j$ has $0 < P[F] < \frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)}$. Now $C^P(-nF) \cap B_j^{P_j}(-nF) = -nF$, so

$C^P(-nF) \neq \emptyset$. But by proposition 4.3, $P[C^P(-nF)] \leq \sum_{i \in I} (1 - P[-nF])^{\frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)}} > 1 - \sum_{i \in I} \frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)} = 0$. But $P[C^P(-nF)] > 0$ $C^P(-nF) \neq \emptyset$.

Morris (1993) showed the same conclusion with an independent argument under the assumptions that all Q_i finite are finite, $I = 2$ and $p_1 = p_2 = \frac{1}{2}$. This case was critical in the results of Morris, Rob and Shin (1995).

Finally, let us note that if we did not restrict attention to simple events, a version of proposition 4.3 would still hold with a weaker bound.

Corollary 4.5. In any information system $\mathcal{I}; \dots; P; f; Q_i; g; i \in I$, if $\sum_{i \in I} p_i < 1$, then for all events E ,

$$P[C^P(E)] \leq \sum_{i \in I} (1 - P[E])^{\frac{1}{\sum_{i \in I} p_i} \frac{1}{\min(p_i)}}.$$

Since $B_i^P(E)$ is a simple event and $C^P(B_i^P(E)) = C^P(E)$, the corollary is a straightforward consequence of proposition 4.3 and lemma 4.7 (which will be proved in the next section).

4.3. Proof of the Critical Path Result

We prove the critical path result in this sub-section. Because the argument is somewhat intricate and - we believe - of more general interest, we present a proof which - while long - clarifies the basic logic. The argument is structured as follows. In section 4.3.1, we provide some basic inequalities which relate together the ex ante probability of an event and the ex ante probability of first order beliefs about that event. In section 4.3.2, we provide an inequality relating the ex ante probability of an event (i.e. $P[E]$) and the ex ante probability that that event is K th order p -belief (i.e. $P[B_i^p]^K(E)$). Finally, in section 4.3.3, we examine what happens to that inequality as K becomes large. If $\sum_i p_i > 1$, the inequality provides no useful bound. But if $\sum_i p_i < 1$, the limit of the K -inequality provides a proof of the critical path result (proposition 4.3).

4.3.1. Basic Ex Ante Probability Properties of Belief Operators

The proof will use the following results which relate the ex ante probabilities of event E and the events $B_i^{p_i}(E)$. The basic insight is that if p_i is small and event E has high probability, then there will be a high probability that i p_i -believes event E . For purposes of this sub-section, allow $P[E|F]$ to take any value in $[0; 1]$ if $P[F] = 0$. The following result appears in Fudenberg and Tirole (1991).

Lemma 4.6. For all events E , (1) $P[E] \geq p_i P[B_i^{p_i}(E)]$ and (2) $P[B_i^{p_i}(E)] \geq \frac{P[E] p_i}{1 - p_i}$.

Proof. By the definition of conditional probability, we have

$$P[E] = P[E | B_i^{p_i}(E)] P[B_i^{p_i}(E)] + P[E | \neg B_i^{p_i}(E)] P[\neg B_i^{p_i}(E)]; \quad (4.1)$$

Since $P[E | Q_i(\!)] \geq p_i$ if $\! \in B_i^{p_i}(E)$, we have $P[E | B_i^{p_i}(E)] \geq p_i$ if $B_i^{p_i}(E) \neq \emptyset$; and thus (by (4.1))

$$P[E] \geq p_i P[B_i^{p_i}(E)];$$

which gives (1). Since $P[E | Q_i(\!)] < p_i$ if $\! \in \neg B_i^{p_i}(E)$, we have $P[E | \neg B_i^{p_i}(E)] < p_i$ if $\neg B_i^{p_i}(E) \neq \emptyset$; since $P[E | B_i^{p_i}(E)] \leq 1$ by definition, we have (by (4.1)):

$$P[E] \leq P[B_i^{p_i}(E)] + p_i P[\neg B_i^{p_i}(E)];$$

Noting that $P[\neg B_i^{p_i}(E)] = 1 - P[B_i^{p_i}(E)]$, we have

$$P[E] \leq (1 - p_i) P[B_i^{p_i}(E)] + p_i;$$

which gives (2). \square

Lemma 4.7. For all events E , $1 - \prod_{i \in I} P[B_i^p(E)] \leq \prod_{i \in I} \frac{1 - P[E]}{1 - p_i}$:

Proof. By lemma 4.6, $P[B_i^p(E)] \geq \frac{P[E] + p_i}{1 + p_i}$; thus

$$P[-B_i^p(E)] \leq 1 - \frac{P[E] + p_i}{1 + p_i} = \frac{1 - P[E]}{1 + p_i}$$

Taking the union gives:

$$P\left[\bigcup_{i \in I} -B_i^p(E)\right] \leq (1 - P[E]) \prod_{i \in I} \frac{1}{1 + p_i}$$

So

$$1 - \prod_{i \in I} P[B_i^p(E)] \leq (1 - P[E]) \prod_{i \in I} \frac{1}{1 + p_i}$$

Lemma 4.8. (1) For all events E , $P[E \cap -B_i^p(E)] \leq p_i P[-B_i^p(E)]$.

(2) If $F \subseteq F_i$ and $F \subseteq -B_i^p(E)$ then $P[E \setminus F] \leq \frac{p_i}{1 + p_i} P[F \cap E]$:

Proof. Note that $P[E \setminus F | \mathcal{Q}_i(\omega)] = P[E | \mathcal{Q}_i(\omega)]$ if $\omega \notin F$ and F is F_i -measurable. Let $F \subseteq F_i$ and $F \subseteq -B_i^p(E)$. Then since $P[E | \mathcal{Q}_i(\omega)] < p_i$ if $\omega \in F$ and $P[E | F] < p_i$ if $P[F] > 0$,

$$P[E \setminus F] = P[E | F] P[F] \leq p_i P[F]$$

So (1) follows by setting $F = -B_i^p(E)$; (2) follows since $P[F] = P[E \setminus F] + P[F \cap E] \leq 2 P[F \cap E]$.

4.3.2. The Kth order belief inequality

Since $P[C^p(E)] = \lim_{K \rightarrow \infty} \prod_{i \in I} P[B_i^{p^K}(E)]$, we would like to provide an upper bound for $1 - \prod_{i \in I} P[B_i^{p^K}(E)]$ as a function of $1 - \prod_{i \in I} P[E]$. Iterated application of lemma 4.7 gives us that

$$1 - \prod_{i \in I} P[B_i^{p^K}(E)] \leq \prod_{i \in I} \frac{1 - P[E]^K}{1 - p_i^K}$$

This implies that for any given $p \in (0, 1)$ and K , we can guarantee (uniformly across information systems) that $\prod_{i \in I} P[B_i^{p^K}(E)]$ is large by setting $P[E]$ sufficiently close to 1. But for any $p \in (0, 1)$, $\prod_{i \in I} \frac{1 - P[E]^K}{1 - p_i^K} \rightarrow 1$ as $K \rightarrow \infty$, so this inequality will be of no use in bounding

$P[C^P(E)]$. We need a tighter bound. The bound is constructed from the following $I \times I$ matrix R :

$$R = \begin{pmatrix} 0 & \frac{p_2}{1_i p_2} & \frac{p_3}{1_i p_3} & \dots & \frac{p_l}{1_i p_l} \\ \frac{p_1}{1_i p_1} & 0 & \frac{p_3}{1_i p_3} & \dots & \frac{p_l}{1_i p_l} \\ \frac{p_1}{1_i p_1} & \frac{p_2}{1_i p_2} & 0 & \dots & \frac{p_l}{1_i p_l} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{p_1}{1_i p_1} & \frac{p_2}{1_i p_2} & \frac{p_3}{1_i p_3} & \dots & \frac{p_l}{1_i p_l} \end{pmatrix}; \quad (4.2)$$

note that R depends on p_i . Writing $[x^0]_i$ for the i th element of column vector x^0 , let $\gg^\pi(p; K) = \max_{i \geq 1} |I + R + \dots + R^{K-1}|_i$.

Lemma 4.9. In any information system $\mathcal{I}; -; P; \mathcal{F}Q_i \mathcal{G}_{i \geq 1}$, for any simple event E ,

$$1_i P^h [B_i^p]^K(E) \leq (1_i P[E]) \gg^\pi(p; K):$$

Intuitively, we want to maximize (over all information systems and all simple events) the value of $1_i P^h [B_i^p]^K(E)$, subject to the constraint that $P[E] \leq 1_i$, and show that the maximand of this problem is $\gg^\pi(p; K)$. For the remainder of this sub-section, we will be considering this maximization problem, which is (implicitly) parameterized by p , K and π .

But how can we formalize the idea of maximizing over information systems? The trick is to observe that we are not interested in the whole structure of the information system, but only in the ex ante probability of certain events. Thus we may, without loss of generality, focus on the probability of events which concern us. But the definitions of those events impose certain linear inequalities on their ex ante probabilities (such as those of the previous sub-section). Our maximization problem "over all information systems" can thus be reduced to a linear programming problem. This is most easily explained by considering special cases. First, let $I = 2$ and $K = 1$. Fix any $\mathcal{I}; -; P; \mathcal{F}Q_i \mathcal{G}_{i \geq 1}$ and any simple event $E = E_1 \setminus E_2$, where each E_i is measurable with respect to Q_i . Now $B_i^{p_1}(E) \mu E_1$, so the state space - can be represented by the following box:

$\frac{1}{4}(0; 0)$	$\frac{1}{4}(0; 1)$	$\frac{1}{4}(0; 2)$
$\frac{1}{4}(1; 0)$	$\frac{1}{4}(1; 1)$	$\frac{1}{4}(1; 2)$
$\frac{1}{4}(2; 0)$	$\frac{1}{4}(2; 1)$	

where $\frac{1}{4}(i; j)$ is the probability of the corresponding event $(i; j)$; thus, for example, $\frac{1}{4}(2; 1) = P(B_1^{p_1}(E) \setminus E_2 \cap B_2^{p_2}(E))$. Notice that although $\frac{1}{4}(:; :)$ depends on the choice of information

system and event E , the relevant relationship between the events $E_1; E_2; B_1^{p_1}(E)$ and $B_2^{p_2}(E)$ is completely captured by this box representation. Thus our problem reduces to maximizing:

$$\sum_i P[B_i^p(E)] = \left(\begin{array}{l} \frac{1}{4}(0;0) + \frac{1}{4}(0;1) + \frac{1}{4}(0;2) + \frac{1}{4}(1;0) \\ + \frac{1}{4}(1;1) + \frac{1}{4}(1;2) + \frac{1}{4}(2;0) + \frac{1}{4}(2;1) \end{array} \right) \quad (4.3)$$

subject to

$$\sum_i P[E] = \left(\begin{array}{l} \frac{1}{4}(0;0) + \frac{1}{4}(0;1) + \frac{1}{4}(0;2) \\ + \frac{1}{4}(1;0) + \frac{1}{4}(2;0) \end{array} \right) \cdot \dots \quad (4.4)$$

But there are additional restrictions on the $\frac{1}{4}$ implicit in the above construction. In particular, $P[E_2 | Q_1] < p_1$ if $\exists E_1 \cap B_1^{p_1}(E_2)$. Thus $P[E_1 \cap B_1^{p_1}(E_2) \setminus E_2] \cdot p_1 P[E_1 \cap B_1^{p_1}(E_2)]$ and so

$$\frac{1}{4}(1;1) + \frac{1}{4}(1;2) \cdot p_1 (\frac{1}{4}(1;0) + \frac{1}{4}(1;1) + \frac{1}{4}(1;2))$$

Re-arranging gives:

$$\frac{1}{4}(1;1) + \frac{1}{4}(1;2) \cdot \frac{p_1}{1 - p_1} \frac{1}{4}(1;0) \quad (4.5)$$

By a symmetric argument:

$$\frac{1}{4}(1;1) + \frac{1}{4}(2;1) \cdot \frac{p_2}{1 - p_2} \frac{1}{4}(0;1) \quad (4.6)$$

Suppose (without loss of generality) that $p_1 \geq p_2$. Then the maximum possible value of (4.3) subject to (4.4), (4.5) and (4.6) is $1 + \frac{p_1}{1 - p_1}$. One choice of $\frac{1}{4}$ attaining that maximum is:

0	0	0
"	0	" $\frac{p_1}{1 - p_1}$
0	0	

Intuitively, we want to distribute the " outside event E in such a way that it can be used to "knock out" the maximum probability when the belief operators are applied once. Given that $p_1 \geq p_2$, " $\frac{p_1}{1 - p_1}$ " is the most probability that can be knocked out at the next round.

Now let $K = 2$; the relevant division of the state space - becomes:

$\frac{1}{4}(0;0)$	$\frac{1}{4}(0;1)$	$\frac{1}{4}(0;2)$	$\frac{1}{4}(0;3)$
$\frac{1}{4}(1;0)$	$\frac{1}{4}(1;1)$	$\frac{1}{4}(1;2)$	$\frac{1}{4}(1;3)$
$\frac{1}{4}(2;0)$	$\frac{1}{4}(2;1)$	$\frac{1}{4}(2;2)$	$\frac{1}{4}(2;3)$
$\frac{1}{4}(3;0)$	$\frac{1}{4}(3;1)$	$\frac{1}{4}(3;2)$	

We seek to maximize

$$1_i P^h [B_i^{p_i}]^2(E) = \begin{matrix} \text{NW} & \text{NE} \\ \text{SW} & \text{SE} \end{matrix} \begin{matrix} \text{NW} & \text{NE} \\ \text{SW} & \text{SE} \end{matrix} \begin{matrix} \frac{1}{4}(0;0) + \frac{1}{4}(0;1) + \frac{1}{4}(0;2) + \frac{1}{4}(0;3) \\ \frac{1}{4}(1;0) + \frac{1}{4}(1;1) + \frac{1}{4}(1;2) + \frac{1}{4}(1;3) \\ \frac{1}{4}(2;0) + \frac{1}{4}(2;1) + \frac{1}{4}(2;2) + \frac{1}{4}(2;3) \\ \frac{1}{4}(3;0) + \frac{1}{4}(3;1) + \frac{1}{4}(3;2) \end{matrix} \quad (4.7)$$

subject to

$$1_i P[E] = \begin{pmatrix} \frac{1}{4}(0;0) + \frac{1}{4}(0;1) + \frac{1}{4}(0;2) + \frac{1}{4}(0;3) \\ \frac{1}{4}(1;0) + \frac{1}{4}(2;0) + \frac{1}{4}(3;0) \end{pmatrix} \quad (4.8)$$

$$\frac{1}{4}(1;1) + \frac{1}{4}(1;2) + \frac{1}{4}(1;3) \cdot \frac{p_1}{1_i p_1} \frac{1}{4}(1;0) \quad (4.9)$$

$$\frac{1}{4}(2;2) + \frac{1}{4}(2;3) \cdot \frac{p_1}{1_i p_1} (\frac{1}{4}(2;0) + \frac{1}{4}(2;1)) \quad (4.10)$$

$$\frac{1}{4}(1;1) + \frac{1}{4}(2;1) + \frac{1}{4}(3;1) \cdot \frac{p_2}{1_i p_2} \frac{1}{4}(0;1) \quad (4.11)$$

$$\frac{1}{4}(2;2) + \frac{1}{4}(3;2) \cdot \frac{p_2}{1_i p_2} (\frac{1}{4}(0;2) + \frac{1}{4}(1;2)) \quad (4.12)$$

This programming problem has solution $1 + \frac{p_1}{1_i p_1} + \frac{p_1}{1_i p_1} \frac{p_2}{1_i p_2}$. One choice of $\frac{1}{4}$ attaining that maximum is:

0	0	0	0
"	0	" $\frac{p_1}{1_i p_1}$	0
0	0	0	0
0	0	" $\frac{p_1}{1_i p_1}$	" $\frac{p_2}{1_i p_2}$

Intuitively, there is a "critical path" which is the optimal way to use the " to knock out probability in all future rounds; we use this intuition in the following general formulation.

The General Linear Programming Formulation. We now formalize the above approach for arbitrary I and K . Fix $I; ; ; P; fQ_i g_{i \in I}$ and, for each $i, E_i^1 \in F_i$. Thus $E^1 = \bigcup_{i \in I} E_i^1$ is a simple event. Define inductively $fE_i^n; ; ; E_i^n; E_i^{n+1} g_{n=1}^{K+1}$ as follows: $E^n = \bigcup_{i \in I} E_i^n$ and $E_i^{n+1} = B_i^{p_i}(E^n)$. Thus $E_i^2 = B_i^{p_i}(E_i^1)$ and $E_i^n = B_i^{p_i} \circ B_i^{p_i} \circ \dots \circ B_i^{p_i}(E_i^1)$ for all $n \geq 2$. By convention, let $E_i^0 = ;$, for all i . Let $D_i^n = E_i^n \cap E_i^{n+1}$ for all $n = 0; K$ and $D_i^{K+1} = E_i^{K+1}$. Observe that by fact 6, the sets $fD_i^n g_{n=0}^{K+1}$ partition $;$ for each i . In particular, $fD_i^n g_{n=0}^{K+1}$ is a coarser partition than Q_i . Let

$n = (n_1; \dots; n_l)$ be a typical element of $f_0; 1; \dots; K + 1g^l$. We shall denote by $\min(n)$ the smallest number in $f_{n_1; \dots; n_l}g$. Define $L(n) = \bigcup_{i \geq 1} D_i^{n_i}$ and $\frac{1}{4}(n) = P[L(n)]$ for all $n \in f_0; \dots; K + 1g^l$. This notation allows us to characterize every relevant region of the state space by an "address" $n \in f_0; \dots; K + 1g^l$; note that the $L(n)$ are disjoint. With two individuals, we can represent the state space by the following box:

$L(0; 0)$	$L(0; 1)$	\dots	$L(0; n_2)$	\dots	$L(0; K + 1)$	$\tilde{A} D_1^0$
$L(1; 0)$	$L(1; 1)$		$L(1; n_2)$		$L(1; K + 1)$	$\tilde{A} D_1^1$
\dots		\ddots				
$L(n_1; 0)$	$L(n_1; 1)$		$L(n_1; n_2)$		$L(n_2; K + 1)$	$\tilde{A} D_1^{n_1}$
\dots				\ddots		
$L(K + 1; 0)$	$L(K + 1; 1)$		$L(K + 1; n_2)$		$L(K + 1; K + 1)$	$\tilde{A} D_1^{K+1}$
"	"		"		"	
D_2^0	D_2^1		$D_2^{n_2}$		D_2^{K+1}	

Thus for all $n = 1; K + 1$ and $i \geq 1$,

$$D_i^n = \left[\begin{matrix} f_{n_2 f_0; \dots; K+1g^l} \\ n_i = n_g \end{matrix} L(n); \right] \quad (4.13)$$

$$E^n = \left[\begin{matrix} f_{n_2 f_0; \dots; K+1g^l} \\ \min(n) \leq n_g \end{matrix} L(n); \right] \quad (4.14)$$

Now for all $n = 1; K$ and $i \geq 1$, $D_i^n \mu - n E_i^{n+1} = -n B_i^{n_i}(E^n)$; so we have by lemma 4.8,

$$P[D_i^n \setminus E^n] \leq \frac{p_i}{1 - p_i} P[D_i^n \cap E^n]; \quad (4.15)$$

Combining (4.13), (4.14) and (4.15), gives, for all $n = 1; K$ and $i \geq 1$,

$$\frac{1}{4}(n) \leq \frac{p_i}{1 - p_i} \left[\begin{matrix} \text{O} \\ \text{B} \\ \text{A} \end{matrix} \right] \frac{1}{4}(n) \quad (4.16)$$

$f_{n_2 f_0; \dots; K+1g^l} : n_i = n$ and $\min(n) = n_g$

 $f_{n_2 f_0; \dots; K+1g^l} : n_i = n$ and $\min(n) < n_g$

These inequalities follow from the construction of the events. Now our maintained hypothesis that $P[E] \leq 1 - \epsilon$ implies:

$$\frac{1}{4}(n) \leq \epsilon; \quad (4.17)$$

$f_{n_2 f_1; \dots; K+1g^l} : \min(n) = 0_g$

Observe that $\mathbf{E}^{\mathbf{B}_i^{\mathbf{P}^{\mathbf{K}}} \mathbf{1}} = \mathbf{E}^{\mathbf{K}+1}$, so $\mathbf{1}_i \mathbf{P}^{\mathbf{h} \mathbf{E}^{\mathbf{B}_i^{\mathbf{P}^{\mathbf{K}}} \mathbf{1}}} = \mathbf{P}^{\mathbf{1}} \mathbf{1}(\mathbf{n})$.

Thus we are interested in the following linear programming problem (P)⁴:

$$\begin{aligned} \max \quad & \mathbf{1}(\mathbf{n}) \\ \text{subject to} \quad & \mathbf{1}(\mathbf{n}) \geq 0, \quad (4.16) \text{ and } (4.17). \end{aligned} \tag{4.18}$$

Thus the statement of lemma 4.9 has been re-interpreted as a linear programming problem. It remains only to show that the maximand is $\mathbf{1}(\mathbf{p}; \mathbf{K})$.

We can guess the form of the solution to P from the $\mathbf{K} = 2$ and $\mathbf{K} = 1; 2$ cases outlined above. The solutions had the property that only certain critical locations had positive probability and this property generalizes. Write $\mathbf{c}(i; \mathbf{n})$ for the location where all components are equal to $\mathbf{n} + 1$ except the i th which is \mathbf{n} , i.e. $\mathbf{c}(i; \mathbf{n}) = (\mathbf{n} + 1; \dots; \mathbf{n} + 1; \mathbf{n}; \mathbf{n} + 1; \dots; \mathbf{n} + 1)$, and say that \mathbf{n} is a critical location if it can be written in this form. The critical path constraint requires:

$$\mathbf{1}(\mathbf{n}) = 0, \text{ if } \mathbf{n} \text{ is not a critical location.} \tag{4.19}$$

The critical path intuition suggests that the following problem (P⁰) has the same value as problem (P).

$$\begin{aligned} \max \quad & \mathbf{1}(\mathbf{n}) \\ \text{subject to} \quad & \mathbf{1}(\mathbf{n}) \geq 0, \quad (4.16), (4.17) \text{ and } (4.19). \end{aligned} \tag{4.20}$$

We will prove this by looking at the dual problems of (P) and (P⁰).

The Dual Problem. First, the dual problem (D) of (P) is as follows. Call the constraint corresponding to individual i and \mathbf{n} in (4.16) $(i; \mathbf{n})$, and denote by $\mathbf{p}_i(i; \mathbf{n})$ the shadow price of the constraint. Similarly, denote by \mathbf{q}_i the shadow price of the constraint in (4.17). Consider any $\mathbf{n} = (\mathbf{n}_1; \dots; \mathbf{n}_1)$ with $0 < \min(\mathbf{n}) \leq \mathbf{K}$. Then $\mathbf{1}(\mathbf{n})$ appears in the left hand side of (4.16) if and only if $\mathbf{n}_i = \min(\mathbf{n})$, i.e., $\mathbf{1}(\mathbf{n})$ appears in all inequalities $(i; \mathbf{n})$ with $\mathbf{n}_i = \min(\mathbf{n})$; $\mathbf{1}(\mathbf{n})$ appears in the right hand side of every inequality $(i; \mathbf{n})$ with $\mathbf{n}_i = \mathbf{n} > \min(\mathbf{n})$. If $\mathbf{k} = 0$, then $\mathbf{1}(\mathbf{n})$ also appears in inequality in (4.17). Thus the dual problem has the following constraints:

for each $\mathbf{n} = (\mathbf{n}_1; \dots; \mathbf{n}_1)$ with $\min(\mathbf{n}) = 0$;

$$\mathbf{q}_i + \sum_{\mathbf{f}_i: \mathbf{n}_i > 0} \frac{\mathbf{p}_i}{\mathbf{1}_i \mathbf{p}_i} (i; \mathbf{n}_i) \leq 1; \tag{4.21}$$

⁴Note that equation (4.16) must hold as a strict inequality unless both sides of the equation are identically equal to zero. Thus the maximum of the programming problem P will be the supremum of the achievable values of $\mathbf{1}_i \mathbf{P}^{\mathbf{h} \mathbf{E}^{\mathbf{B}_i^{\mathbf{P}^{\mathbf{K}}} \mathbf{1}}}$, but may not be attainable.

for n with $0 < \min(n) < K$;

$$\sum_{f_i: n_i = \min(n)} x_{(i; n_i)} + \sum_{f_i: n_i > \min(n)} \frac{p_i}{1 + p_i} x_{(i; n_i)} \leq 1; \quad (4.22)$$

and for n with $\min(n) = K$;

$$\sum_{f_i: n_i = K} x_{(i; K)} \leq 1; \quad (4.23)$$

Thus the dual problem (D) of (P) is:

$$\begin{aligned} & \min \{ \zeta \} \\ & \text{subject to } x_{(i; n)} \geq 0, \forall x_{(i; n)} \geq 0, (4.21), (4.22) \text{ and } (4.23). \end{aligned} \quad (4.24)$$

Next, let us construct the dual (D^0) of the constrained primal problem (P^0). Look at the constraints corresponding critical addresses, which will give us the dual form of the critical path: for $c(i; 0)$, from (4.21), we have:

$$x_{(i; 0)} + \sum_{j \in I} \frac{p_j}{1 + p_j} x_{(j; 1)} \leq 1 \text{ for each } i; \quad (4.25)$$

for $c(i; n)$, $n < K$, from (4.22), we have:

$$x_{(i; n)} + \sum_{j \in I} \frac{p_j}{1 + p_j} x_{(j; n+1)} \leq 1 \text{ for each } i \text{ and } n; \quad (4.26)$$

and for $c(i; n)$, $n = K$, from (4.23), we have:

$$x_{(i; K)} \leq 1; \quad (4.27)$$

Thus the dual problem (D^0) is:

$$\begin{aligned} & \min \{ \zeta \} \\ & \text{subject to } x_{(i; n)} \geq 0; \forall x_{(i; n)} \geq 0, \text{ and } (4.25), (4.26) \text{ and } (4.27). \end{aligned} \quad (4.28)$$

By construction, the minimum of (D^0) is no larger than (D) since (D^0) has less constraints.

Now let $x^0(i; k) = 1 + R + \dots + R^{K-i} \cdot 1^0$, and $\zeta^0 = \max_{i \in I} x^0(i; 0)$. This implies that $x^0(i; K) = 1$ for all $i \in I$, $x^0(i; k) = 1 + \sum_{j \in I} \frac{p_j}{1 + p_j} x^0(j; k+1)$ and for all $i \in I$ and $0 \leq k < K$; and $\zeta^0 = 1 + \max_{i \in I} \sum_{j \in I} \frac{p_j}{1 + p_j} x^0(j; 1)$.

Lemma 4.10. (λ^*, μ^*) is a solution to (D^0) .

Proof. Note that given any $\lambda, \mu = 1 + \max_{i \geq 1} \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; 1)$ is clearly optimal since $\lambda \leq 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; 1)$ for all i . Since $\lambda(i; 1) \leq 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; 2)$ from (4.26) and $\max_{i \geq 1} \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; 1)$ is a weakly increasing function of vector $(\lambda(i; 1))_{i \geq 1}$, the minimum is attained at $\lambda(i; 1) = 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; 2)$ for any given $\lambda(j; 2)$. The iterative argument shows $\lambda(i; k) = 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda(j; k + 1)$ and thus the statement of the lemma. \square

Lemma 4.11. (λ^*, μ^*) is a solution to (D) .

Proof. Since problem (D) includes all constraints of (D^0) , by lemma 4.10 it suffices to show that (λ^*, μ^*) satisfy (4.21), (4.22) and (4.23). (4.21) is clear by construction. Note that for any k , $\lambda^*(i; k) \leq \lambda^*(i; k + 1)$ holds for every i : it is straightforward to check that this is true for $K \leq i \leq 1$, and if it is true for k , then $\lambda^*(i; k + 1) = 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda^*(j; k) \leq 1 + \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda^*(j; k + 1) = \lambda^*(i; k)$, thus it is true for $k \leq i \leq 1$.

Now for any n with $\min(n) = 0$, $\mu^*(i) \leq \sum_{f: n_f > 0} \frac{p_f}{1 - p_f} \lambda^*(i; n_f) \leq \sum_{f: n_f > 0} \mu^*(i) \leq \mu^*(i) \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda^*(j; 1) \leq 1$, where the first inequality holds due to the monotonic property of λ^* shown above. Thus (4.21) is satisfied. Similarly, for n with $0 < \min(n) = n < K$, $\mu^*(i) \leq \sum_{f: n_f > 0} \frac{p_f}{1 - p_f} \lambda^*(i; n_f) \leq \sum_{f: n_f = n} \mu^*(i; n) \leq \sum_{f: n_f > 0} \frac{p_f}{1 - p_f} \lambda^*(i; n + 1) \leq \sum_{f: n_f = n} \mu^*(i; n) \leq \sum_{j \in i} \frac{p_j}{1 - p_j} \lambda^*(j; n + 1) = 0$, where i^* can be any i with $n_i = n$. Therefore (4.23) holds. \square

By the duality theorem of linear programming, the value of problem (D) is the same as the value of problem P , completing the proof of lemma 4.9, since $\mu^*(p; K) = \mu^*$ (making the dependence on p and K explicit).

4.3.3. Evaluating μ^*

It remains only to consider what happens to $\mu^*(p; K)$ as $K \rightarrow 1$. First we establish some properties of the matrix R . We will use the following decomposition of R . Write I for the identity matrix. Write D for the diagonal matrix with i th diagonal element p_i . Write X for the matrix of 1's.

$$I = \begin{pmatrix} 0 & 1 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \end{pmatrix}, D = \begin{pmatrix} 0 & p_1 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & p_2 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & p_i \end{pmatrix}, X = \begin{pmatrix} 0 & 1 & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & 1 & 1 & \dots & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & 1 & 1 & \dots & 1 \end{pmatrix}$$

Observe that $R = (X_i I) D (I_i D)^{i-1}$.

Lemma 4.12. $I + R + R^2 + \dots$ is bounded if $\sum_{i=2}^P p_i < 1$.

Proof. $\sum_{j=2}^P p_j < 1 \implies \sum_{j=2}^P p_j \cdot \sum_{i=1}^j p_i < 1 \implies p_i \sum_{j=i}^P p_j < 1$ for all i . $\implies \max_{i=2}^P \frac{\sum_{j=i}^P p_j}{1 - \sum_{j=i}^P p_j} < 1$.

Write 1 for the column vector on 1's. So $R(1^0_i p^0) = \sum_{i=2}^P p_i \cdot 1^0_i p^0 \cdot \sum_{j=i}^P p_j (1^0_j p^0)$ and thus $I + R + R^2 + \dots = \sum_{i=2}^P p_i (1^0_i p^0) \cdot (1 + \sum_{j=i}^P p_j + \sum_{j=i}^P p_j^2 + \dots) (1^0_j p^0) = \frac{1}{1 - \sum_{j=i}^P p_j} (1^0_i p^0)$.
Write $\pm = 1_i \sum_{i=2}^P p_i$.

Lemma 4.13. $(I_i R)^{i-1} = I_i (X_i I) D (I_i D)^{i-1} 1^i = (I_i D)^3 I + \frac{1}{\pm} X D$.

Proof. First observe that $X D = \begin{pmatrix} 0 & p_1 & p_2 & p_3 & \dots & p_i \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ p_1 & p_2 & p_3 & \dots & p_i \end{pmatrix}$ and so $X D X D = (1_i \pm) X D$ (4.29)

Now:

$$\begin{aligned} I_i (X_i I) D (I_i D)^{i-1} 1^i &= (I_i D)^3 I + \frac{1}{\pm} X D \\ &= \left(I_i (X_i I) D (I_i D)^{i-1} 1^i \right) \\ &= \left(I_i (X_i I) D (I_i D)^{i-1} 1^i + \frac{1}{\pm} X D (I_i D)^{i-1} 1^i \right) \\ &= \frac{1}{\pm} X D + I_i \frac{1}{\pm} (1_i \pm) X D (I_i D)^{i-1} 1^i, \text{ by (4.29)} \\ &= \frac{1}{\pm} X D + I_i \frac{1}{\pm} (1_i \pm) X D (I_i D)^{i-1} 1^i \end{aligned}$$

These two lemmas imply that:

Corollary 4.14. If $\prod_{i=2}^I p_i < 1$, then $\sigma^p(p; K) \rightarrow \sigma_{i=2}^{1, \min(p_i)} A$ as $K \rightarrow 1$.

Proof. Recall that $\sigma^p(p; K) = \max_{i=2}^I \{I + R + \dots + R^{K-1}\}^i$. Lemma 4.12 implies that $\sigma^p(p; K) \rightarrow \max_{i=2}^I [I + R]^{i-1} 1^i$ as $K \rightarrow 1$. Lemma 4.13 shows that the $(i; i)$ th co-ordinate of $[I + R]^{i-1} 1^i$ is $\sum_{k=2}^I \frac{1-p_k}{1-p_k} A_{i-1} p_j$ and the $(i; j)$ th element is $\sum_{k=2}^I \frac{1-p_k}{1-p_k} A_{i-1} p_j$ if $i \neq j$. Thus

$$[I + R]^{i-1} 1^i \text{ is a column vector with } i\text{th coordinate } \sum_{k=2}^I \frac{1-p_k}{1-p_k} A_{i-1} p_j. \text{ Thus } \lim_{K \rightarrow 1} \sigma^p(p; K) = \max_{i=2}^I \sum_{k=2}^I \frac{1-p_k}{1-p_k} A_{i-1} p_j = \sum_{i=2}^I \frac{1-p_i}{1-p_i} A_{i-1} p_i.$$

Now proposition 4.3 follows from corollary 4.14 and lemma 4.9.

5. Robustness and p-Dominance

5.1. The Robustness of p-Dominant Equilibria

It is an implication of the work of Monderer and Samet (1989) that any strict Nash equilibrium action profile will be played with high probability in some equilibrium action distribution if there is common p-belief of payoffs, for p close to 1, with high ex ante probability. But what if we have common p-belief of payoffs for some p which is not close to 1? We will extend earlier results of Morris, Rob and Shin (1995) to answer this question.

Definition 5.1. Action profile a^p is a p-dominant equilibrium of G if for all $i \in I$, $a_i \in A_i$ and all $s_i \in \Phi(A_i)$ with $s_i(a_i^p) \geq p_i$,

$$\sum_{a_i \in A_i} s_i(a_i) g_i(a_i^p; a_i) \geq \sum_{a_i \in A_i} s_i(a_i) g_i(a_i; a_i^p):$$

(This definition extends the definition in Morris, Rob and Shin (1995) to the many player case with asymmetric p; we have also replaced strict inequalities by weak inequalities). Note that a dominant strategies equilibrium is a (0; ...; 0)-dominant equilibrium, and that any pure strategy Nash equilibrium is a 1-dominant equilibrium. It is easy to see if a^p is a p-dominant equilibrium and $p^0 \leq p$ (with the usual vector ordering), then a^p is a p^0 -dominant equilibrium. So for any pure strategy equilibrium a^p , we are interested in the smallest p for which a^p is a p-dominant equilibrium.

The following result uses the core idea of Monderer and Samet's (1989) main result.

Lemma 5.2. Suppose action profile a^* is a p -dominant equilibrium of G . Consider any $U \in E(G)$ and let F be a p -evident event such that $F \subseteq \mu_{-i}$. Then U has a Bayesian Nash equilibrium where $\sigma_i(a_j^*) = 1$ for all $i \in I$ and $j \in F$.

Proof. Let $F_i = B_i^{p_i}(F)$, so $F \subseteq \bigcap_{i \in I} F_i$ by assumption. Consider the modified incomplete game U^0 where each player's strategy must satisfy $\sigma_i(a_j^*) = 1$, for all $j \in F_i$ and $i \in I$. There exists an equilibrium σ of the modified game. We shall show that σ is in fact an equilibrium of U . By construction, for every i , at any $j \in F_i$, σ_i is a best response to σ_{-i} . Let $j \in F$. Then by definition, $P[F_j | Q_i(j)] \geq p_i$, thus by the construction of σ , the conditional probability of a_j^* being played is at least p_i , so a_j^* is a best response against σ_{-i} . Thus σ is also an equilibrium of the original game. \square

Thus a p -dominant equilibrium is robust if we can find a large p -evident F contained in μ_{-i} for any U which is near G . When can we find such an event? Consider the cyclic matching pennies game (example 3.1). The unique (strict) Nash equilibrium in that example is in fact $(\frac{4}{5}; \frac{4}{5}; \frac{4}{5})$ -dominant. If there was a high probability $(\frac{4}{5}; \frac{4}{5}; \frac{4}{5})$ -evident event contained in μ_{-i} in every nearby incomplete information game, we could prove robustness. In this case, there was not. However, this strategy may work for some values of p .

Observe that $C^p(\mu_{-i})$ is the largest set which is p -evident and contained in μ_{-i} (by theorem 4.2). Thus we need to show that there is $\epsilon > 0$ such that $P[C^p(\mu_{-i})] \geq 1 - \epsilon$ for any μ -elaboration U where μ is small enough. But this is what proposition 4.3 provides. In particular:

Proposition 5.3. Suppose a^* is p -dominant with $\prod_{i \in I} p_i < 1$. Then a^* is robust to incomplete information.

Proof. Write μ^* for the distribution putting probability 1 on a^* . Fix any $\epsilon > 0$. By proposition 4.3, we can choose $\mu > 0$, such that $P[E] > 1 - \epsilon$ implies $P[C^p(E)] > 1 - \epsilon$. Thus by construction of μ , for any $U \in E(G; \mu)$, we have $P[C^p(\mu_{-i})] > 1 - \epsilon$. By lemma 5.2, there exists a Bayesian Nash equilibrium of U with $\sigma_i(a_j^*) = 1$, for all $j \in C^p(\mu_{-i})$. Thus there exists an equilibrium action distribution of U with $\sigma_j(a^*) \geq P[C^p(\mu_{-i})] > 1 - \epsilon$. Therefore, $|\sigma_j(a) - \mu^*(a)_j| < \epsilon$ for all $a \in A$, so μ^* is robust. \square

Example 5.4. Co-ordination Game. Let $A_i = \{a^1; \dots; a^N\}$ for each $i \in I$. Let

$$g_i(a) = \begin{cases} x_i^n, & \text{if } a_j = a^n, \text{ for all } j \in I \\ 0, & \text{otherwise} \end{cases}$$

where $x_i^n > 0$ for all $i \in I$ and $n = 1; \dots; N$.

This game has N symmetric pure strategy strict Nash equilibria. Equilibrium $a = (a^n; \dots; a^1)$ is p -dominant if $p_i x_i^n \geq (1 - p_i) x_i^m$ for all $m \in n$ and $i \in I$, i.e. if $p_i \geq \frac{x_i^m}{x_i^m + x_i^n}$ for all $i \in I$. Thus, by proposition 5.3, $a = (a^n; \dots; a^1)$ is robust if $\prod_{i \in I} \max_{m \in n} \frac{x_i^m}{x_i^m + x_i^n} < 1$. Thus there is a sense in which the condition $\prod_{i \in I} p_i < 1$ becomes increasingly severe as the number of players increases.

5.2. Non-robustness of p -Dominated Action Profiles

Actions which survive iterated deletion of strictly dominated actions are never played in any Nash equilibrium. In this section, we provide an analogous, stronger, condition on a set of actions sufficient to ensure that they are never played in any robust equilibrium.

An action subset $A^\#$ is a subset of A of the form $A^\# = A_1^\# \times A_2^\# \times \dots \times A_I^\#$ where $A_i^\# \subseteq A_i$ for each $i \in I$.

Definition 5.5. Action subset $A^\#$ is p -dominated in G if, for all $i \in I$ and all $a_i \in A_i \setminus A_i^\#$ with $\sum_{a_i \in A_i^\#} p_i < p_i$,

$$A_i^\# \setminus a_i : \arg \max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}^\#} (a_i, a_{-i}) g_i(a_i; a_{-i}) > \sum_{a_{-i} \in A_{-i}^\#} (a_i, a_{-i}) g_i(a_i; a_{-i}) ;$$

That is, $A^\#$ is p -dominated if for every i , taking an action a_i in $A_i^\#$ is not a best response unless he assigns probability at least p_i on the others playing some actions in $A_{-i}^\#$.

Proposition 5.6. If σ is robust in G and $A^\#$ is p -dominated in G with $\prod_{i \in I} p_i > 1$, then $\sigma(a) = 0$ if $a_i \in A_i^\#$ for some $i \in I$.

The proposition is an immediate consequence of the following lemma.

Lemma 5.7. If $A^\#$ is p -dominated in G with $\prod_{i \in I} p_i > 1$, then, for all $\epsilon > 0$, there exists $U \subseteq E(G; \epsilon)$ such that in any equilibrium action distribution σ of U , $\sigma(a) = 0$ if $a_i \in A_i^\#$ for some $i \in I$.

Thus actions in the p -dominated action subset are never played in any equilibrium of the ϵ -elaboration.

Proof. The proof is by construction of U . Note that if $A^\#$ is p -dominated in G with $\prod_{i \in I} p_i > 1$, then there exists q with $\prod_{i \in I} q_i = 1$ such that $A^\#$ is q -dominated in G . Now let $\epsilon = \frac{1}{I} \in \mathbb{Z}_+$. Let

$P(i; k) = \prod_{j \in I} q_j$. Let each Q_i consist of (1) the singleton event $f(i; 0)g$; and (2) all events of the form $(j; k)_{j \in I}; (i; k + 1)$, for each integer $k \geq 0$. Let

$$u_i(a; i) = \begin{cases} g_i(a), & \text{if } i \in (i; 0) \\ 1, & \text{if } i = (i; 0) \text{ and } a_i \geq A_i^a \\ 0, & \text{if } i = (i; 0) \text{ and } a_i < A_i^a \end{cases};$$

and let σ_i be any equilibrium of U . Write σ_i^j for the probability i attaches to action set A_i^j given strategy profile σ_{-i} , i.e.

$$\sigma_i^j = \sum_{i \in Q_i(i)} P[i \in Q_i(i)] \sum_{a_i \in A_i^j} \sigma_i^j(a_i);$$

By construction, $\sigma_i^j(a_i; (i; 0)) = 0$ if $a_i < A_i^j$ (any such a_i is strictly dominated). If $i = (i; 1)$, then $P[f(j; 0) : j \in I | Q_i(i)] > \prod_{j \in I} q_j$, so $\sigma_i^j < q_j$. Since A_i^j is q_j -dominated, this shows that $\sigma_i^j(a_i; i) = 0$ if $a_i < A_i^j$ and $i = (i; 1)$. The argument iterates to establish the result. \square

The following example shows how this result can be used even when our earlier results have no bite.

Example 5.8.

	H	T	X
H	12; 10	10; 12	0; 0
T	10; 12	12; 10	0; 0
X	0; 0	0; 0	1; 1

Since the game has multiple Nash equilibria, proposition 3.2 has no implications here. The only pure strategy equilibrium is $(X; X)$, but since this is a p -dominant equilibrium only if $p_1 \geq \frac{12}{13}$ and $p_2 \geq \frac{12}{13}$, proposition 5.3 also has no implications here. But consider the singleton action subset $f(X; X)g$; $f(X; X)g$ is p -dominated if $p_1 < \frac{11}{12}$ and $p_2 < \frac{11}{12}$. Thus by proposition 5.6, X is never played by either player in any robust equilibrium.

5.3. Strong p -dominant equilibria and unique robust equilibria

The following stronger concept of p -dominance will help clarify the notions introduced so far.

Definition 5.9. Action profile a^a is a strong p -dominant equilibrium of G if for all $i \in I$, $a_i \in A_i^a$ and all $\sigma_{-i} \in \Sigma_{-i}(A_{-i}^a)$ with $\sum_{a_i \in A_i^a} \sigma_i(a_i; a_{-i}) > p_i$,

$$\sum_{a_i \in A_i^a} \sigma_i(a_i; a_{-i}) g_i(a_i; a_{-i}) > \sum_{a_i \in A_i^a} \sigma_i(a_i; a_{-i}) g_i(a_i; a_{-i});$$

That is, playing a_i^* is a unique best response if i believes that the probability of some player j choosing a_j^* is greater than p_i . By continuity, a strong p -dominant equilibrium must be p -dominant, since $\exists a_{-i} : a_j = a_j^*$ for some $j \in I$ with $\pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$. The converse does not hold, since typically we will have $\exists a_{-i} : a_j = a_j^*$ for some $j \in I$ with $\pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$. The two player case is special, because then $\exists a_{-i} : a_j = a_j^*$ for some $j \in I$ with $\pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$. In the two player case, for a generic choice of payoffs, a p -dominant equilibrium must be a strong p -dominant equilibrium. We will use this almost equivalence for the two player case in the next section. However, consider a game with non-generic payoffs where all players are indifferent between all actions. Every action profile will be a p -dominant equilibrium, and no action profile will be a strong p -dominant equilibrium, for any number of players.

Recall that $1_i - p_i$ is the row vector with i th component $1 - p_i$. Now we have:

Lemma 5.10. Action profile a^* is a strong p -dominant equilibrium if and only if $A_1 n f a_1^* g \in A_2 n f a_2^* g \in \dots \in A_n n f a_n^* g$ is $(1_i - p_i)$ -dominated.

Proof. Let $A^* = A_1^* \in \dots \in A_n^* = A_1 n f a_1^* g \in A_2 n f a_2^* g \in \dots \in A_n n f a_n^* g$. Suppose $\exists i \in I$. Now

$$\exists a_{-i} : \pi_i(a_i, a_{-i}) < \pi_i(1_i - p_i, a_{-i}), \quad \exists a_{-i} : a_j = a_j^* \text{ for some } j \in I \text{ with } \pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$$

So A^* is $(1_i - p_i)$ -dominated if and only if,

$$A_i^* \setminus : \arg \max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} (a_i, a_{-i}) g_i(a_i, a_{-i}) = a_i^*$$

for all $i \in I$ and all $\exists a_{-i} : a_j = a_j^*$ for some $j \in I$ with $\pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$, which holds if and only if $\arg \max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} (a_i, a_{-i}) g_i(a_i, a_{-i}) = a_i^*$ for all $i \in I$ and all $\exists a_{-i} : a_j = a_j^*$ for some $j \in I$ with $\pi_i(a_i, a_{-i}) > \pi_i(a_i^*, a_{-i})$, i.e. a^* is a strong p -dominant equilibrium. 2

This immediately implies:

Corollary 5.11. (Uniqueness). If a^* is a strong p -dominant equilibrium with $\prod_{i \in I} p_i < 1$, then a^* is the unique robust equilibrium of the game. 2

Proof. If a^* is strong p -dominant, then a^* is p -dominant, so a^* is robust by proposition 5.3. But by lemma 5.10 and proposition 5.6, no action profile other than a^* is played in any robust equilibrium.

In appendix 9.1, we show how our new results on p -belief can be used to address related questions.

6. Two Player, Two Action Games

Two player games are special because p -dominant equilibria are - generically - strong p -dominant equilibria. We will exploit this in the analysis that follows. We will give a complete characterization of robust equilibria for generic two player, two action games. Recall that, except for some non-generic cases, a two player, two action game has (1) a unique (strict) pure strategy equilibrium; (2) two (strict) pure strategy equilibria; or (3) no pure strategy equilibrium and a unique mixed strategy equilibrium. Let us consider each case in turn.

- ² If there is a unique pure strategy equilibrium, then at least one player has a dominant strategy to play his action in that equilibrium. Say that it is player 1. Then the unique equilibrium is $(0; p)$ -dominant for some $p < 1$. Thus the unique Nash equilibrium is robust by proposition 5.3.
- ² If there are two pure strategy equilibria, then (generically) exactly one of them is risk dominant in the sense of Harsanyi and Selten (1988). In a two player, two action game, an equilibrium is risk dominant exactly if it is strong $(p_1; p_2)$ -dominant equilibrium for some $(p_1; p_2)$ with $p_1 + p_2 < 1$. Now by corollary 5.11, the risk dominant equilibrium is the unique robust equilibrium.
- ² If a generic two player two action game has no pure strategy Nash equilibrium, then the (unique) mixed strategy equilibrium is the unique correlated equilibrium. So that unique equilibrium must be robust by proposition 3.2.

Proofs and more detailed discussion of these results is provided in the appendix (section 9.2).

7. Alternative Formulations of Robustness

In this section we present a framework for understanding the relation between our notion of robustness and the existing "refinements" literature. In section 2.2, we constructed the set $E(G)$ of incomplete information games "embedding" G and the subset $E(G; \mu)$ with $P[-U] = 1_j \mu$. We were extremely liberal in what we allowed to occur in constructing the perturbed games, and thus our concept of robustness is very strong: we allowed unbounded state spaces, correlated signals and more. One might wish to know whether an equilibrium is robust to a more restrictive class of elaborations, and this is what we study in this section.

Consider a mapping with $F(G) \mu E(G)$ and write $F(G; \mu) = F(G) \setminus E(G; \mu)$.

Definition 7.1. Action distribution σ^1 is robust under F if for every $\epsilon > 0$, there exists $\delta > 0$, such that every $U \in F(G; \mu)$ with $\mu \cdot \sigma^1$ has an equilibrium action distribution σ^0 with $k^1_j \sigma^0 \cdot \epsilon$.

Thus our notion of "robust to incomplete information" corresponds to "robustness under E." As we put shrink F , "robust under F " becomes a weaker requirement. In the following, we consider various examples of mapping F and their implications.

7.1. Bounded Elaborations

Let $E_F(G)$ be the set of incomplete information games embedding G where the state space is finite. There would be no change in our results if we replaced E by E_F in our definition. Our positive results would go through with identical proofs. The required ϵ -elaborations in the cyclic pennies matching game (example 3.1) could be generated in a finite state space: for any given $\epsilon > 0$, choose $N > \frac{2}{\epsilon}$, let $\Omega = \{0, 1, \dots, N\}$, let P be uniform on the state space, let payoffs and partitions be essentially as before. Now for each $\epsilon > 0$, we have a finite ϵ -elaboration where the unique strict Nash equilibrium action profile is never played. Note, however, that as $\epsilon \rightarrow 0$, a larger and larger state space would be required in the construction of the example.

Consider, then, a set of bounded elaborations. Let $E_M(G)$ be the set of incomplete information games embedding G where the state space has at most M elements. Now we have the following result:

Lemma 7.2. For fixed M , if a^* is a strict equilibrium of G , then a^* is robust under E_M .

Proof. In any information system, $\mathbb{E} B_{\epsilon}^{p^n}(E)$ is decreasing in n for $n \geq 1$ (by belief operator properties 4 and 5). So in any information system with at most M states, $C^p(E) = \mathbb{E} B_{\epsilon}^{p^M}(E)$ for all events E . By lemma 4.9,

$$1 - \mathbb{P} [B_{\epsilon}^p]^M(E) \leq (1 - \mathbb{P}[E])^{\epsilon} (p; M);$$

and so

$$\mathbb{P} [C^p(E)] = \mathbb{P} [B_{\epsilon}^p]^M(E) \geq 1 - (1 - \mathbb{P}[E])^{\epsilon} (p; M); \tag{7.1}$$

Now since a^* is a strict equilibrium of game G , there exists $p < 1$ such that a^* is p -dominant. Write π^* for the probability distribution in $\Phi(A)$ putting probability 1 on a^* . Fix any $\epsilon > 0$. By (7.1), we can choose $\delta > 0$ such that $\mathbb{P}[E] > 1 - \delta$ implies $\mathbb{P}[C^p(E)] > 1 - \epsilon$. Then in particular, for any ϵ -elaboration U of G , we have $\mathbb{P}[C^p(-U)] > 1 - \epsilon$. By lemma 5.2, there exists a Bayesian Nash equilibrium of U with $\sum_i \pi_i^*(a_i^*) = 1$, for all $\pi \in C^p(-U)$. Thus there exists an equilibrium action distribution of U with $\pi^*(a^*) \geq \mathbb{P}[-U] > 1 - \epsilon$. Therefore, $|\pi^*(a) - \pi^*(a^*)| < \epsilon$ for all $a \in A$, so π^* is robust. \square

Equation 7.1 is related to theorem 14.5 in Fudenberg and Tirole (1991, page 567): they show that, for a fixed finite state space, if $\mathbb{P}^k[E] \geq 1 - \delta$, then there exists $p^k \geq 1 - \epsilon$ such that $\mathbb{P}^k[C^{(p^k; \dots; p^k)}(E)] \geq 1 - \epsilon$. They use this result to make essentially the same point.

Strictness is by no means a necessary condition for even a pure strategy equilibrium to be robust under E_M . In particular, by the argument above, any equilibrium which is p -dominant for some $p < 1$ will be robust under E_M . An action profile a^* is a p -dominant equilibrium for some $p < 1$ if and only if for each $i \in I$ and $a_i \in A_i$, either $g_i(a_i^*; a_{-i}^*) > g_i(a_i; a_{-i}^*)$ or $g_i(a_i^*; a_{-i}^*) \geq g_i(a_i; a_{-i}^*)$ for all $a_{-i} \in A_{-i}$. In fact, pure strategy equilibria which are not p -dominant for any $p < 1$ may be robust under E_M . But a necessary condition for robustness under E_M is strict perfection in the sense of Okada (1981) (an equilibrium is strictly perfect if it is stable in the sense of Kohlberg and Mertens (1986) as a singleton set). Since strictly perfect equilibria need not exist, equilibria robust under E_M need not exist.

Lemma 7.2 shows that our claim that many strict Nash equilibria are not robust relies on unbounded state space elaborations. However, we would argue that if the definition of robustness is altered to consider only state spaces with number of elements bounded by M , it would be reasonable to change the definition of "sufficiently small". The logic of our arguments suggests that if we replaced the requirement "for all $\epsilon > 0$ " with the requirement "for all ϵ of the order of $\frac{1}{M}$ ", essentially the same results would continue to hold.

7.2. Independent Elaborations

Write $E_I(G)$ for the set of games embedding G where $P[E_i \cap Q_i(\! \cdot \!)] = P[E_i]$ for all $i \in I$, $i \in I$ and all events $E_i = \bigcup_{j \in I} E_j$, each $E_j \in F_j$. Thus $E_I(G)$ is the set of games embedding G where players have independent types.

Now we have:

Lemma 7.3. If a^* is a strict Nash equilibrium of G , then a^* is robust under E_I .

Proof. Again, a^* is a p -dominant equilibrium for some $p < 1$. Now choose $\epsilon > 0$ but sufficiently small such that $1 - \prod_{j \in I} \frac{1}{1 - p_j} \epsilon > p_i$ for all $i \in I$. Suppose event E has probability at least $1 - \prod_{j \in I} \frac{1}{1 - p_j} \epsilon$.

Now by the independence assumption and lemma 4.7, we have that for all $i \in I$ and $\! \cdot \! \in \Omega$:

$$P \left[\bigcup_{j \in I} B_j^{p_j}(\! \cdot \!) \cap Q_i(\! \cdot \!) \right] = P \left[\bigcup_{j \in I} B_j^{p_j}(\! \cdot \!) \right] \cdot P[B_i^p(\! \cdot \!)] \geq 1 - \prod_{j \in I} \frac{1}{1 - p_j} \epsilon > p_i,$$

so that $P[B_i^p(\! \cdot \!) \cap Q_i(\! \cdot \!)] > p_i$ if $\! \cdot \! \in B_i^{p_i}(\! \cdot \!)$. Since $P[B_i^p(\! \cdot \!) \cap Q_i(\! \cdot \!)] = 0$ if $\! \cdot \! \notin B_i^{p_i}(\! \cdot \!)$, we have that $B_i^{p_i}(\! \cdot \!) \cap Q_i(\! \cdot \!) = B_i^{p_i}(\! \cdot \!)$, so $B_i^{p_i, n}(\! \cdot \!) = B_i^{p_i}(\! \cdot \!)$ for all $n \geq 1$, $C^p(\! \cdot \!) = B_i^p(\! \cdot \!)$ and so

$P[C^p(\! \cdot \!)] > 1 - \prod_{j \in I} \frac{1}{1 - p_j} \epsilon$. Now a similar argument to lemma 7.2 completes the proof. \square

Again, this result is far from tight. Strict perfection is a necessary condition for robustness under E_I ; we conjecture that it is sufficient.

Lemma 7.3 shows that our results rely on the correlation of signals. Is this reasonable? Let us just emphasize why independence matters so much. It is a consequence of independence of signals that $B_i^{p_i} \cap B_j^{p_j}(E) \neq \emptyset$ for all events E and $i \neq j$ and thus that $B_i^{p_i} \cap B_j^{p_j} \cap B_k^{p_k}(E) = B_j^{p_j} \cap B_k^{p_k}(E)$ for all events E and $i, j, k \in I$ with $j \neq k$. In other words, independence rules out any interesting questions about higher order beliefs.

8. Conclusion

The technical contribution of the paper is threefold. First, we have introduced a general notion of the robustness of an equilibrium to a small amount of incomplete information. While closely related to a number of approaches, certain subtle details of our approach ensure that we get very different results. Second, we have demonstrated the existence of games where no robust equilibrium exists. There is an open set of such games. Third, we have provided two different kinds of positive results. The correlated equilibrium result is straightforward. The common p -belief results utilize new work showing the existence of a surprising amount of structure in arbitrary belief hierarchies with common priors. We believe that this "hidden content" of the common prior is an important area of future work.

We certainly do not wish to interpret this work as a new refinement. We have shown that some equilibria of some games have the property that they can be played in equilibria of nearby incomplete information games. As economists, we can be relatively confident in those predictions. But in other games, there is no robust equilibrium. This tells us that the equilibrium outcome must be sensitive to changes in the information structure which are small when measured in terms of the ex ante probability of payoff relevant events. The correct way to analyze this sensitivity appears to involve paying attention to "higher order beliefs". The importance of "higher order beliefs" emerges mathematically in our analysis.

Rubinstein (1989) and others have argued that "boundedly rational" players might ignore subtleties of higher order beliefs and - in certain circumstances - behave as if there was common knowledge of payoffs even when the underlying information structure does not justify this presumption. This may be true in some circumstances. But we would like to argue against a presumption that bounded rationality should imply a systematic bias in favor of playing "as if" there is common knowledge. A reasonable boundedly rational rule for playing games might be the following. If there really is common knowledge of payoffs - say, because payoffs are publicly announced - play according to a Pareto-efficient or otherwise focal Nash equilibrium. If there is not common knowledge of payoffs, and there is some action which is an equilibrium action whatever the fine details of the situation, play that. This leaves open the question of how a boundedly rational player should play in a game with no robust equilibrium. Consider the cyclic matching pennies game (example 3.1). What is reasonable play in that game when it is not common knowledge, with probability one, that the payoffs are correct? For fully rational players, iterated

deletion of dominated strategies alone ensures extreme sensitivity to the information structure. Boundedly rational players might not be able to perform the required iterated deletion. But there is every reason to believe that actual reasonable play - like fully rational play - will involve players trying to outguess others' choices in ways that are highly sensitive to the information structure.

References

- [1] Aumann, R.: Agreeing to Disagree. *Annals of Statistics* 4, 1236-1239 (1976).
- [2] Aumann, R.: Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica* 55, 1-18 (1987).
- [3] Bargers, T.: Weak Dominance and Approximate Common Knowledge. *Journal of Economic Theory* 64, 265-276 (1994).
- [4] Carlsson, H. and van Damme, E.: Global Games and Equilibrium Selection. *Econometrica* 61, 989-1018 (1993a).
- [5] Carlsson, H. and van Damme, E.: Equilibrium Selection in Stag Hunt Games. in *Frontiers of Game Theory*, edited by K. Binmore, A. Kirman and P. Tani. Cambridge: M.I.T. Press (1993b).
- [6] Cotter, K.: Correlated Equilibrium in Games with Type-Dependent Strategies. *Journal of Economic Theory* 54, 48-68 (1991).
- [7] Cripps, M.: Extreme Correlated and Nash Equilibria in Two-Person Games. University of Warwick (1994).
- [8] Debreu, G.: A Social Equilibrium Existence Theorem. *Proceedings of the National Academy of Sciences* 38, 886-893 (1990).
- [9] Dekel, E. and Fudenberg, D.: Rational Behavior with Payo[®] Uncertainty. *Journal of Economic Theory* 52, 243-267 (1990).
- [10] Forges, F.: Correlated Equilibrium in Two-Person Zero-Sum Games *Econometrica* 58, 515 (1990).
- [11] Fudenberg, D., Kreps, D. and Levine, D.: On the Robustness of Equilibrium Re⁻nements *Journal of Economic Theory* 44, 354-380 (1980).
- [12] Fudenberg, D. and Tirole, J.: *Game Theory*. Cambridge: M.I.T. Press (1991).

- [13] Harsanyi, J. and Selten, R.: A General Theory of Equilibrium Selection in Games. Cambridge: M.I.T. Press (1988).
- [14] Kajii, A. and Morris, S.: Common p -Belief: the General Case. CARESS Working Paper 94-15 (1994a).
- [15] Kajii, A. and Morris, S.: Payo[®] Continuity in Incomplete Information Games. CARESS Working Paper 94-17 (1994b).
- [16] Kim, Y.: Equilibrium Selection in n -person Coordination Games. forthcoming in Games and Economic Behavior (1993).
- [17] Kohlberg, E. and Mertens, J.-F.: On the Strategic Stability of Equilibria *Econometrica* 54, 1003-1038 (1986).
- [18] Monderer, D. and Samet, D.: Approximating Common Knowledge with Common Beliefs. *Games and Economic Behavior* 1, 170-190 (1989).
- [19] Monderer, D. and Samet, D.: Proximity of Information in Games with Incomplete Information. forthcoming in *Mathematics of Operations Research*, (1990).
- [20] Morris, S.: A Curious Property of Belief Operators. University of Cambridge Economic Theory Discussion Paper #195 (1993).
- [21] Morris, S., Rob, R. and Shin, H.: p -Dominance and Belief Potential. *Econometrica* 63, 145-157 (1995).
- [22] Okada, A.: On the Stability of Perfect Equilibrium Points. *International Journal of Game Theory* 10, 67-73 (1981).
- [23] Neyman, A.: Correlated Equilibrium and Potential Games. Seminar notes (1991).
- [24] Rubinstein, A.: The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge. *American Economic Review* 79, 385-391 (1989).
- [25] Selten, R.: Re-Examination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* 4, 25-55 (1965).
- [26] Shin, H. and Williamson, T.: How Much Common Belief is Necessary for a Convention? forthcoming in *Games and Economic Behavior* (1994).

9. Appendix

9.1. A Generalization of Morris, Rob and Shin (1995)

Here we report a generalization of a result of Morris, Rob and Shin (1995). That result was for two player symmetric games with a finite state space. We give a result for many player asymmetric games with an infinite state space.

Fix an incomplete information game $U = (I; f; A_i; g_{i2I}; -; P; f; Q_i; g_{i2I}; f; u_i; g_{i2I})$. For any given $! , G^a(!) = (I; f; A_i; g_{i2I}; f; u_i; g_{i2I})$ is a complete information game.

Proposition 9.1. Suppose that [1] Players know their own payoffs, so $! \in Q_i(!) \Rightarrow u_i(a; !^0) = u_i(a; !)$ for all $a \in A$; [2] a^a is a strong p -dominant equilibrium of $G^a(!)$, for all $! \in -$; [3] a_j^a is a strictly dominant strategy for some player j at some state $!$; [4] $C^{1i P}(E) \subseteq f; -g$ for all simple events E . Then $\%_i(a_j^a) = 1$ for all $! \in -$ and $i \in I$ is the unique Bayesian Nash equilibrium.

Proof. Suppose $\%_i$ is a Bayesian Nash equilibrium. Write $F_i^a = \{! : \%_i(a_j^a) \in [1-g] \text{ and } F^a = \bigcup_{i \in I} F_i^a$. Let us consider two cases. [i] Suppose first that $F^a = -$; if $! \in F_i^a$, then for every $! \in Q_i(!)$, there exists $j \in I$ such that $! \in F_j^a$. Thus player i assigns probability 1 to states where some other player j is playing a_j^a with probability 1, so by [1] and [2], he has a strict best response to play a_i^a , so $! \in F_i^a$, a contradiction. Thus $F_i^a = -$ for all $i \in I$ and the proposition is proved. [ii] Suppose that $F^a \neq -$. By [3], we know that $F_j^a \neq -$ for some player j and, by construction, we have $F_j^a \subseteq F_j$. Thus $C^{1i P}(F^a) \subseteq F^a$ and in particular $C^{1i P}(F^a) \neq -$. Thus by [4], $C^{1i P}(F^a) = -$. Since $F^a \neq -$, we cannot have $F^a \subseteq C^{1i P}(F^a)$. Thus there exists a player $k \in I$ and a state $!^a \in F_k^a$ such that $P[F^a | Q_k(!^a)] < 1 - p_k$. So $P(! : \%_j(a_j^a) = 1 \text{ for some } j \in I - Q_k(!^a)) > p_k$. Thus by [1] and [2], $\%_k(a_k^a) = 1 \Rightarrow !^a \in F_k^a$, a contradiction. Thus $F^a = -$, which is impossible by [ii]. \square

It is straightforward to show that in fact always playing a^a is the unique behavior surviving iterated deletion of dominated strategies. This result will obviously only be of interest if there exist information systems with property [4] of the proposition. In Morris, Rob and Shin (1995), the belief potential of an information system was the smallest number p such that $C^{(1-p; 1-p)}(E) \subseteq f; -g$ for all simple events E . The belief potential was shown (in a finite setting) to be less than $\frac{1}{2}$. The analogous belief potential in our setting would be multidimensional: for any given information system, we would be interested in the set $B = \{p \in [0; 1]^I : C^{1i P}(E) \subseteq f; -g \text{ for all simple events } E\}$. If some individual's partition has an infinite number of elements, it is a consequence of corollary 4.4 (on page 18) that if $C^{1i P}(E) \subseteq f; -g$ for all simple events E , we must have $\prod_{i \in I} p_i < 1$. Thus proposition 9.1 is vacuous if a^a is not strong p -dominant for some p with $\prod_{i \in I} p_i < 1$.

9.2. Two Player, Two Action Games

Here we prove more formally the assertions of section 6 and provide some further discussion. We will give a complete characterization of robust equilibria for generic two player, two action games. Thus we consider games of the form

	L	R
L	a,b	c,d
R	e,f	g,h

Let us consider each of three generic cases in turn.

9.2.1. Unique Pure Strategy Equilibrium

Suppose (without loss of generality) that (L; L) is the unique (strict) Nash equilibrium. Then at least one player has a dominant strategy to play L (if not, (R; R) would be a Nash equilibrium). Thus (L; L) is either (p; 0)-dominant or (0; p)-dominant for some $p < 1$. Thus (L; L) is robust by proposition 5.3.

Lemma 9.2. If a generic two player, two action game has a unique pure strategy Nash equilibrium, then that equilibrium is robust.

9.2.2. Two Pure Strategy Equilibria

Without loss of generality, assume (L; L) and (R; R) are strict Nash equilibria, so $a > e$, $g > c$, $b > d$ and $h > f$. Harsanyi and Selten (1988) said that (L; L) ((R; R)) is risk dominant if $(a - e)(b - d) > (<) (g - c)(h - f)$. For a generic game with two pure strategy equilibria, exactly one will be risk dominant.

Lemma 9.3. If a generic two player, two action game has two pure strategy Nash equilibria, then the risk dominant equilibrium is robust while the risk dominated equilibrium and the mixed strategy equilibrium are not robust.

Proof. Suppose (without loss of generality) that (L; L) is risk dominant (and so $(g - c)(h - f) < (a - e)(b - d)$). Now if we set $p_1 = \frac{(g - c)}{(g - c) + (a - e)}$ and $p_2 = \frac{(h - f)}{(h - f) + (b - d)}$, we have $p_1 + p_2 < 1$, $p_1 a + (1 - p_1)c > p_1 e + (1 - p_1)g$ and $p_2 b + (1 - p_2)f > p_2 d + (1 - p_2)h$. So (L; L) is a strong p-dominant equilibrium for some p with $p_1 + p_2 < 1$ and thus is the unique robust equilibrium, by corollary 5.11.

9.2.3. No Pure Strategy Equilibrium

Lemma 9.4. If a generic two player two action game has no pure strategy Nash equilibrium, then the (unique) mixed strategy equilibrium is the unique correlated equilibrium.

Proof. Suppose (without loss of generality) that $a > e$, $g > c$, $d > b$ and $f > h$. This game has a unique Nash equilibrium where 1 plays L with probability $\frac{1}{2} = \frac{(g_i - c)}{(g_i - c) + (a_i - e)}$ and 2 plays L with probability $\frac{1}{2} = \frac{(f_i - h)}{(f_i - h) + (d_i - b)}$. Suppose there were a correlated equilibrium where (L; L) was played with probability π_1 , (L; R) with probability π_2 , (R; R) with probability π_3 and (R; L) with probability π_4 , i.e.

	L	R
L	π_1	π_2
R	π_4	π_3

Suppose $\pi_1 = 0$. Then if $\pi_2 > 0$, 1's optimal best response is to play R when he is supposed to play L. Thus $\pi_2 = 0$. Analogous arguments show that $\pi_2 = 0 \Rightarrow \pi_3 = 0$, $\pi_3 = 0 \Rightarrow \pi_4 = 0$ and $\pi_4 = 0 \Rightarrow \pi_1 = 0$. Thus we must have $\pi_i > 0$ for all $i = 1; 2; 3; 4$. But now optimality requires that $\frac{\pi_1}{\pi_2} \geq \frac{1/2}{1 - 1/2}$, $\frac{\pi_4}{\pi_3} \geq \frac{1/2}{1/2}$, $\frac{\pi_3}{\pi_1} \geq \frac{1 - 1/2}{1/2}$ and $\frac{\pi_2}{\pi_4} \geq \frac{1/2}{1 - 1/2}$. If any of these inequalities were strict inequalities, then multiplying inequalities together would give $1 > 1$. Thus they are all equalities and $\pi_1 = \pi_4 \pi_2$, $\pi_2 = \pi_1 (1 - 1/2)$, $\pi_3 = (1 - \pi_1) (1 - 1/2)$ and $\pi_4 = (1 - \pi_1) 1/2$. This is the original Nash equilibrium.

Lemma 9.5. If a generic two player two action game has no pure strategy Nash equilibrium, then the (unique) mixed strategy equilibrium is robust.

Proof. Follows immediately from lemma 9.4 and proposition 3.2.

9.2.4. The Non-Generic Case

For non-generic payoffs, a robust equilibrium need not exist. Consider the following example.

Example 9.6. Simple Co-ordination Game.

	L	R
L	1,1	0,0
R	0,0	1,1

No equilibrium is robust to incomplete information. Consider the following ϵ -elaboration: $I = (f_1; 2g; -; -; f_0; 1; 2; \dots; g; P(\cdot)) = (1 - \epsilon)$; $Q_1 = (f_0; 1g; f_2; 3g; \dots)$ and $Q_2 = (f_0g; f_1; 2g; f_3; 4g; \dots)$;

$u_i(a; \mu) = g_i(a)$ for all $i \in I$ and $a \in A$ except that at state 0, player 2 has a dominant strategy to play L. The unique equilibrium of this μ -elaboration has L played everywhere by both players. Thus (R; R) and the mixed strategy equilibrium are not robust to incomplete information. However, a symmetric argument implies that (L; L) is not robust to incomplete information.

9.3. Changing Probability Distributions

We have been allowing the whole state space to vary as we perturb the game. Thus in particular, the set of elaborations does not have a good topological structure, and so the reader may wonder if this is the cause of our non-robustness results. We study this issue in this subsection.

Endow the set of all probability measures with the weak topology, which is metrizable; let d be a metric. Write $E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}$ for the set of incomplete information games embedding G , with state space structure $\Sigma; f; Q_i; g_{i21}$ and indexed by a "limit probability distribution" P^μ with $P^\mu[\Sigma] = 1$. Now introduce a more restrictive notion of elaborations; let

$$E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}^a = \bigcup_{U \in E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}} \{d(P; P^\mu) \leq \epsilon\};$$

that is, it is required that not only P assigns a high probability on Σ but also P is close to P^μ weakly. So if $U^k \in E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}^a$, $k = 1, 2, \dots$ and $\mu^k \rightarrow \mu$, then the corresponding P^k converges to P^μ weakly. Note that weak convergence is equivalent to pointwise convergence under our assumptions.

Definition 9.7. Action distribution μ is robust against elaboration sequences if, for every $\Sigma; f; Q_i; g_{i21}; P^\mu; \mu$ and $\epsilon > 0$, there exists $\delta > 0$, such that every $U \in E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}^a$ with $\mu \cdot \delta$ has an equilibrium action distribution ν with $\|\nu - \mu\| \leq \epsilon$.

It is clear from the definitions that if μ is robust to incomplete information, then it is robust against elaboration sequences. How important is the strengthening? Notice that there are two parts to this strengthening. First, we fixed the state space structure $\Sigma; f; Q_i; g_{i21}$ and allowed the μ to depend on it. Second, we specified a limit distribution P^μ on Σ and used a stricter notion of closeness. The first strengthening does not change any of our arguments: we nowhere use the fact that μ is chosen independently of $\Sigma; f; Q_i; g_{i21}$. The second is more important: consider, for example, the cyclic matching pennies game (example 3.1). There we generated a class of incomplete information games indexed by $\mu > 0$. As $\mu \rightarrow 0$, the sequence of probability measures do not converge, so there is no $\Sigma; f; Q_i; g_{i21}; P^\mu; \mu$ such that each μ -elaboration belongs to $E_X \{G; \Sigma; f; Q_i; g_{i21}; P^\mu; \mu\}^a$. As it happens, the example can be altered to show that the strict Nash equilibrium is not robust against elaboration sequences. But a somewhat more complex sequence of μ -elaborations must be constructed.

How big is the difference between the two notions? To understand the difference, it is useful to state proposition 4.3 in terms of sequences. Given a sequence P^k of probability distributions, write $B_i^{p_i:k}$, $B_x^{p_i:k}$ and $C^{p_i:k}$ for the p -belief and common p -belief operators generated by P^k .

Proposition 9.8. Fix a state space structure $\langle \Omega, \mathcal{G} \rangle$ and consider a sequence of probability distributions P^k and event E such that $P^k[E] \rightarrow 1$. (1) If $\liminf_{i \rightarrow \infty} p_i < 1$, then $P^k C^{p_i:k}(E) \rightarrow 1$. (2) Conversely, if $\liminf_{i \rightarrow \infty} p_i > 1$, it is possible to construct a state space structure $\langle \Omega, \mathcal{G} \rangle$, a sequence of probability distributions P^k and event E such that $P^k[E] \rightarrow 1$ but $C^{p_i:k}(E) = 0$ for all k .

(1) follows immediately from proposition 4.3; the following example shows (2) for the (hardest) case where $\liminf_{i \rightarrow \infty} p_i = 1$.

Example 9.9. Fix $\epsilon > 0$ and p such that $\sum_{i \in \mathbb{Z}_+} p_i = 1$. Let $\mathcal{G} = \{ \{i\} \in \mathbb{Z}_+ \}$. Let $P^k(i; m) = (1 - \epsilon)^k \epsilon^m p_i$. Let each Q_i consist of (1) the singleton event $\{i\}$; and (2) all events of the form $\{j; m\}_{j \in \mathbb{Z}_+; (i; m+1)}$, for each $m \geq 0$. Consider the event $E = \bigcup_{i \in \mathbb{Z}_+} \{i\}$. By construction $P[E] = 1 - \epsilon^k$. Observe that $B_x^{p_i:k}(E) = \sum_{(i; m) \in \mathcal{G}} P^k(i; m)$. The proof is by induction. This is obviously true for $n = 0$. Suppose it is true for $n \leq 1$. Now

$$\begin{aligned} \frac{P^k[\bigcup_{(j; m) \in \mathcal{G}} \{j; m\} \setminus \bigcup_{(i; n)} \{i; n\}]}{P[\bigcup_{(j; m) \in \mathcal{G}} \{j; m\}]} &= \frac{\epsilon^k (1 - \epsilon)^k p_i}{\epsilon^k (1 - \epsilon)^k p_i + \epsilon^k (1 - \epsilon)^{k-1} (1 - p_i)} \\ &= \frac{(1 - \epsilon)^k p_i}{(1 - \epsilon)^k p_i + 1 - p_i} \\ &< p_i \end{aligned}$$

Thus $\sum_{(i; n) \in \mathcal{G}} B_i^{p_i:k}(\{j; m\}_{j \in \mathbb{Z}_+; (i; n+1)}) = B_i^{p_i:k} B_x^{p_i:k}(\{i\})$. This is true for each i , so $B_x^{p_i:k}(E) = \sum_{(i; m) \in \mathcal{G}} P^k(i; m)$ and so $C^{p_i:k}(E) = 0$.

Proposition 9.8 didn't require that P^k converged pointwise to some limit distribution. What happens if we add this requirement?

Proposition 9.10. Fix a state space structure $\langle \Omega, \mathcal{G} \rangle$ and consider a sequence of probability distributions P^k and event E such that $P^k \rightarrow P^1$ pointwise and $P^1[E] = 1$. (1) If $\liminf_{i \rightarrow \infty} p_i < 1$, then $P^k C^{p_i:k}(E) \rightarrow 1$. (2) Conversely, if $\liminf_{i \rightarrow \infty} p_i > 1$, it is possible to construct a state space structure $\langle \Omega, \mathcal{G} \rangle$, a sequence of probability distributions P^k and an event E such that $P^k \rightarrow P^1$ pointwise and $P^1[E] = 1$ but $C^{p_i:k}(E) = 0$ for all k .

Note that (1) follows from proposition 4.3 except in the critical case where $\prod_{i \in I} p_i = 1$. To see what is so special about $\prod_{i \in I} p_i = 1$, consider again example 9.9. We had $P^k[E] = 1$ and $C^{P^k}(E) = \emptyset$ for all k . But the sequence P^k did not have a well-defined limit. It is not possible to modify the example to have a well defined limit and still have $C^{P^k}(E) = \emptyset$. We can prove (2) by the following construction.

Example 9.11. Suppose $\prod_{i \in I} p_i > 1$. Let $\Omega = \{0, 1\}^I \in Z_+$. Let Q_1 consist of $f(0; k) : k \in Z_+$; $(j; 0)_{j \in I}$; and all information sets of the form $(i; m); (j; m+1)_{j \in I}$, for $m \in Z_+$. For $i \in I$, let Q_i consist of $(j; 0)_{j \in I}$; and all information sets of the form $(i; m); (j; m+1)_{j \in I}$, for $m \in Z_+$.

$$P^k(i; m) = \begin{cases} \frac{(1-q)^m p_i}{c q^k + p_j} & \text{if } i \in I \\ \frac{(1-q)^m c}{c q^k + p_j} & \text{if } i = 0 \text{ and } m > k \\ 0 & \text{if } i = 0 \text{ and } m \leq k \end{cases}$$

where

$$\frac{1}{\prod_{j \in I} p_j} < q < 1 \tag{9.1}$$

$$c > \max_{i \in I} \frac{p_i}{\prod_{j \in I} p_j} \tag{9.2}$$

P^k converges pointwise to P^1 , where

$$P^k(i; m) = \begin{cases} \frac{(1-q)^m p_i}{p_j} & \text{if } i \in I \\ 0 & \text{if } i = 0 \end{cases}$$

Let $E = \{f(i; k) : i \in I\}$ and $k \in Z_+$; now $P^1[E] = 1$; E is simple; and $B_1^{P^1}(E) = E$. Consider information set $(i; m); (j; m+1)_{j \in I}$ of some player $i \in I$; i 's posterior probability of event E

at this information set is 1, if $m \leq k$; otherwise, it is $\frac{q^m p_i + q^{m+1} p_j}{q^m p_i + q^{m+1} p_j + q^{m+1} c}$, which, by (9.1)

and (9.2), is strictly less than p_i . Thus $B_{\alpha}^P(E) = f(1; k)g [f(i; m) : m < kg$. Now consider information set $(1; k); (j; k + 1)_{j \in I \setminus \{i\}}$ of player 1; player 1's posterior probability of event $B_{\alpha}^P(E)$ is $\frac{q^k p_i}{q^k p_i + q^{k+1} p_j}$, which is less than p_i , by (9.1). Now an iterative argument establishes that $C^P(E) = \dots$.

Thus to prove proposition 9.10, it remains only to show that (1) holds when $\sum_{i \in I} p_i = 1$. To do this, we will require three lemmas. The first concerns the matrix R (see equation 4.2 on page 21).

Lemma 9.12. If $\sum_{i \in I} p_i = 1$, $R^k 1^0 = \frac{1}{1 + \max_{j \in I} (p_j)} (1^0 + p^0)$ for $k = 0; 1; \dots$

Proof. By the argument of lemma 4.12, $R(1^0 + p^0) = (1^0 + p^0)$ and thus $R^k(1^0 + p^0) = (1^0 + p^0)$ for all $k = 0; 1; \dots$. Now $R^k 1^0 = \frac{1}{1 + \max_{j \in I} (p_j)} R^k(1^0 + p^0) = \frac{1}{1 + \max_{j \in I} (p_j)} (1^0 + p^0)$ for all $k = 0; 1; \dots$

Lemma 9.13. In any information system, any simple event E satisfies:

$$P [B_{\alpha}^P]^{k+1}(E) / [B_{\alpha}^P]^k(E) = (1 + P[E]) \frac{1 + \min_{i \in I} (p_i)}{1 + \max_{i \in I} (p_i)}$$

Sketch of Proof: The proof follows that of lemma 4.10. Now we seek to maximize

$$P [B_{\alpha}^P]^{k+1}(E) / [B_{\alpha}^P]^k(E) = \max_{f(n) \geq j \min(n) = Kg} X \quad \frac{1}{n}$$

instead of

$$1 + P [B_{\alpha}^P]^k(E) = \max_{f(n) \geq j \min(n) = Kg} X \quad \frac{1}{n}$$

Essentially the same critical path argument goes through, and we see that

$$P [B_{\alpha}^P]^{k+1}(E) / [B_{\alpha}^P]^k(E) = (1 + P[E])^{3^k} (p; K)$$

where $3^k (p; K) = \max_{i \in I} R^k 1^0_i$. By lemma 9.12, this is no more than $\frac{1 + \min_{i \in I} (p_i)}{1 + \max_{i \in I} (p_i)}$.

Lemma 9.14. Suppose $P_{\alpha}^k \rightarrow P^1$ pointwise. Then for all $\epsilon > 0$, there exists $\delta > 0$ such that for all k sufficiently large, $P^k \rightarrow P^1$ pointwise. Then for all $\epsilon > 0$, there exists $\delta > 0$ such that for all k sufficiently large, $P^k \rightarrow P^1$ pointwise.

Proof. Fix any $\epsilon > 0$. Since Ω is countable, there is a finite set Ω^ϵ such that $P^1[\Omega^\epsilon] > 1 - \frac{\epsilon}{2}$ and $P^1(\omega) > 0$ for all $\omega \in \Omega^\epsilon$. Since Ω^ϵ is finite, P^k converges uniformly to P^1 on Ω^ϵ , so $P^k[\Omega^\epsilon] > 1 - \epsilon$, for all k sufficiently large. Also, we can choose $\delta > 0$ such that for all $\omega \in \Omega^\epsilon$, $P^1(\omega) > \delta$ and $P^k(\omega) > \delta$ for all sufficiently large k . But now $\mu : P^k(\omega) < \delta - \epsilon$ for all sufficiently large k , hence $P^k(\omega) < \delta - \epsilon < 1 - \epsilon$. \square

Proof of final step of Proposition 9.10: Fix the sequence $P^k \rightarrow P^1$ with $P^1[E] = 1$. Fix $q < 1$. By lemma 9.14, there exists $\delta > 0$ such that for all k sufficiently large, $P^k(\omega) < \delta < 1 - q$. Since $P^k \rightarrow P^1$ and $P^1[E] = 1$, then $(1 - P^1[E]) \frac{1 - \min(p_i)}{1 - \max(p_i)} < \delta$ for all k sufficiently large. This implies (by lemma 9.13) that $\mu : P^k(\omega) < \delta$ and $B_i^{p,k}(\omega) \rightarrow B_i^p(\omega)$ for all i, ω . Thus $\mu : P^k(\omega) < \delta$. So $P^k(B_i^{p,k}(E)) = 1 - P^k(B_i^{p,k}(E)) > 1 - \delta > q$ for all sufficiently large k . Thus $P^k(B_i^{p,k}(E)) \rightarrow 1$, and the robustness of p -dominant equilibria, for $p_i > 1$, is proved by the standard argument using lemma 5.2. \square

This small difference in the common p -belief result gives a small difference between the notions of robustness. Consider the symmetric simple coordination game (example 9.6 on page 42). Neither of the two pure strategy equilibria is robust to incomplete information. But the sequence of ϵ -elaborations used to prove this did not have a well-defined limit probability distribution. In fact, both pure strategy equilibria are robust to elaboration sequences, since $P^k(B_i^{p,k}(\omega)) \rightarrow 1$ if $P^k \rightarrow P^1$ pointwise and $P^1[\omega] = 1$.

9.4. Extensions and the relation to Carlsson and van Damme (1993a)

In this section, we will speculate rather loosely on the relation on our results to the work of Carlsson and van Damme (1993a). To do this, we will first speculate how our results would change under certain assumptions.

9.4.1. Uncountable State Spaces

First of all, notice that we have no problem dealing with uncountable state spaces as long as each player has at most a countable number of possible signals. Complications arise with uncountable signals for two reasons.

First, the indeterminacy of conditional probability on zero probability information sets implies that belief operators defined on sets will not be well defined. However, we have shown elsewhere (Kajii and Morris (1994b)) that if belief operators are defined in a natural measure theoretic way (so that they operate on equivalence classes of events), it is possible to replicate (as probability 1

statements) all the main results from the countable case; and the proof of proposition 4.3, relying on the division of the state space into countable boxes, would presumably go through essentially unchanged.

Second, existence of equilibrium in incomplete information games with correlated uncountable signals remains an open question. Of course, we could still develop characterizations conditional on existence.

9.4.2. Knowing your own Payoffs

We required an ϵ -elaboration to have the property that $P[-U] = 1 - \epsilon$. This entailed the assumption that not only is there a high probability that payoffs are given by G , but also there is a high probability that everyone knows their own payoffs. A weaker assumption would let $-U = f(\theta) - \sum_i u_i(a; \theta) = g_i(a)$ for all $a \in A, i \in I$ and require an ϵ -elaboration to have the property that $P[-U] = 1 - \epsilon$. Clearly, $-U \leq -U_0$, so $P[-U] = 1 - \epsilon \implies P[-U_0] \leq 1 - \epsilon$. With this change, it would be necessary to bound the size of the payoffs in the elaborations: otherwise, larger and larger payoffs outside $-U_0$ could dominate the results. But with the bound on elaboration payoffs, we believe our main results would continue to hold (although the change in definition would presumably change significantly which non-strict Nash equilibria were robust).

9.4.3. Local Robustness

If we give our results a sequence interpretation and index information systems by the probability distribution P^k , then we consider sequences of incomplete information games embedding G with $P^k \rightarrow P^1$ (see appendix section 9.3). For any sequence of strategy profiles σ^k , write 1_E^k for the probability distribution over actions contingent on event E being true, given P^k and σ^k . Write $C^{1;\epsilon}(-U)$ for the set of states where $-U$ is common $1-\epsilon$ -belief under probability distribution P^1 .

Now say that action distribution σ^1 is locally robust if, for every sequence of elaborations, there is a sequence of equilibria σ^k , with $1_{C^{1;\epsilon}(-U)}^k \rightarrow \sigma^1$. Our original definition of robustness requires local robustness under the additional restriction that $P^1[-U] = 1$ (which guarantees that $C^{1;\epsilon}(-U) = \Omega$). Presumably our sufficient conditions for robustness would also be sufficient for local robustness.

9.4.4. Relation to Carlsson and van Damme (1993a)

Putting the above results together, we conjecture - leaving aside existence problems - a result of the following form.

Fix a complete information game $G = (I; f; A_i; g_i; \sigma; F; f; F_i; g_i; \sigma; f; u_i; g_i; \sigma)$ and consider any sequence of incomplete information games embedding G , $U^k = (I; f; A_i; g_i; \sigma; P^k; F; f; F_i; g_i; \sigma; f; u_i; g_i; \sigma)$, where

3
 $\Omega; P^k; F$ is a measure space and F_i is the sub σ -field representing i 's information; write $U^1 = (I; fA_i; g_{i21}; \dots; P^1; F, fF_i; g_{i21}; fU_i; g_{i21})$, for the limit of the incomplete information games where $P^k \rightarrow P^1$ weakly. Kajii and Morris (1994a) provides a definition for common 1-belief in this setting.

Conjecture 9.15. Suppose that μ^* is either a p -dominant equilibrium of G for some p with $p_i < 1$ or the unique correlated equilibrium of G . Then it is possible to find a sequence of equilibria μ^k such that $\mu^k \xrightarrow{C^{1,1}(-0)} \mu^*$.

Now we are in a position to compare our result with Carlsson and van Damme (1993a). They show that if an equilibrium of a two player, two action game is risk dominant then, for every sequence of incomplete information games satisfying their "global uncertainty" property, every sequence of equilibria μ^k has $\mu^k \xrightarrow{C^{1,1}(-0)} \mu^*$, where μ^* puts probability 1 on the risk dominant equilibrium.

Thus our result implies the existence of a sequence of equilibria which have the risk dominant outcome played as global uncertainty disappears. Carlsson and van Damme also obtain uniqueness. Intuitively, the reason they obtain uniqueness is because they have restricted attention to a class of critical perturbations (i.e. those satisfying "global uncertainty").