Strategic Implications of (Non-)Common Priors I: Belief Potential*

Daisuke Oyama

Graduate School of Economics, Hitotsubashi University
Naka 2-1, Kunitachi, Tokyo 186-8601, Japan
oyama@econ.hit-u.ac.jp

and

Olivier Tercieux

Paris-Jourdan Sciences Economiques (PSE) and CNRS
48, Boulevard Jourdan, Paris, France
tercieux@pse.ens.fr

December 31, 2005; revised May 30, 2006

*We would like to thank Akihiko Matsui, Takashi Ui, Muhamet Yildiz, and especially, Atsushi Kajii and Michihiro Kandori for helpful comments and discussions. This paper has been presented at Hitotsubashi University and the Kyoto Game Theory Workshop in Kyoto. D. Oyama acknowledges Grant-in-Aid for JSPS Fellows. O. Tercieux acknowledges financial support from the French Ministry of Research (Action Concertée Incitative). This research was initiated while O. Tercieux was visiting the Institute of Economic Research, Kyoto University, whose hospitality is gratefully acknowledged.

Web page: www.econ.hit-u.ac.jp/˜oyama/papers/bPotNCP.html.

Abstract

This paper studies the impact of a small probability event for strategic behavior in incomplete information games with non-common priors. It is shown that the impact of a small probability event has an upper bound that is an increasing function of the measure of discrepancy from the common prior assumption. In particular, the impact can be arbitrarily large under non-common priors, but is bounded from above under common priors. This result quantifies the difference between common prior and non-common prior models in terms of implications on the infinite hierarchies of beliefs. *Journal of Economic Literature* Classification Numbers: C72, D82.

KEYWORDS: common prior assumption; higher order belief; rationalizability; contagion; belief potential.

Contents

1 Introduction

While controversial, the common prior assumption (hereafter, CPA) is used in most models of incomplete information in game theory and economics. This assumption says that the beliefs of all players are generated from a single prior, updated by Bayes' rule, so that differences in their beliefs are due solely to differences in information that they receive. It is well known that the CPA is crucial for many results in incomplete information games (e.g., Aumann's (1976) result on agreeing to disagree and no trade theorems by Milgrom and Stokey (1982)). The purpose of this paper is to clarify the restrictions that we implicitly impose on strategic behavior in game theoretic models when we accept the CPA. Specifically, we focus on "contagion" effects that a small amount of payoff uncertainty has on strategic behavior and (ex ante) payoffs through players beliefs about payoffs, their beliefs about others' beliefs, and so on, i.e., the hierarchy of beliefs.

It has been known that once we depart from common knowledge of payoffs introducing a small amount of incomplete information, strategic behavior may change dramatically through higher order beliefs. Rubinstein (1989), Carlsson and van Damme (1993), and Morris, Rob, and Shin (1995), among others, show how a small probability event can have a large impact on strategic behavior (under common priors). To see the logic behind, suppose that player 1 is known to take a certain action at some information set which has a very small *ex ante* probability. If player 2 puts high *conditional* probability on that event at his information sets where the first information set is thought possible, this knowledge might imply a unique best response by player 2 at these information sets. This, in turn, implies how player 1 responds to that knowledge at larger information sets, and so on. If this iterative argument results in a unique action profile played anywhere on the state space, then we have a contagion of this action profile. Then the question is when it is the case that a certain action profile being chosen at some event (which, again, may have a very small probability) implies that this action profile is chosen everywhere on the state space: in other words, when is an action profile contagious?

To answer this question, Morris, Rob, and Shin (1995) measure the impact of an event by the notion of *belief potential*. First, say that an event E has impact p at a state ω if the statement of the form "player 1 believes with probability at least p that 2 believes with probability at least p that 1 believes ... that the true state is in the original event E" is true at state ω. Then, the belief potential of event E is the largest probability p such that E has impact p at *all* states in the state space. Morris, Rob, and Shin (1995) demonstrate that this notion has a close relation to p-dominance, which is a way to measure the "strength" of a Nash equilibrium.

In Section 2, we first demonstrate with an example that under heterogeneous priors, there exists an information system where small probability events have an arbitrarily large impact. In the example, we show that for any strict Nash equilibrium, there is a dominant solvable incomplete information game with non-common priors in which the unique rationalizable strategy profile is to play this equilibrium everywhere. This implies that any strict Nash equilibrium can be contagious under heterogeneous prior beliefs.

We then provide the upper bound on the impact of small probability events by an increasing function of the degree of the discrepancy from the CPA. Here, we measure the discrepancy in two ways: first, it is measured by the supremum of the ratios between the players' prior probabilities over the states in the state space (whenever it is well defined); we also measure discrepancy using the standard notion of distance between probability measures: the supremum of the (absolute value of the) difference between prior probabilities over the events in the state space. This result implies, first, that for a small probability event to have a maximum impact, not only heterogeneous priors are required, but also these priors must be very far from each other. Second, it shows that the impact of small probability event is bounded from above under the CPA. The latter result actually quantifies the implications of the CPA on the infinite hierarchy of beliefs.

Lipman (2003, 2005) considers the implications of the CPA for finite order of beliefs. He shows that given any state in a partition model where players may have heterogenous priors but with common support, there is a corresponding state in another partition model with a common prior that is close to the original model with respect to product topology in the universal type space (Mertens and Zamir (1985) and Brandenburger and Dekel (1993)). That is, the CPA does not have any significant implication on finite order beliefs, if one is interested only in local properties of the beliefs (i.e., properties at a given state). Appealing to this result, Yildiz (2004) shows, in particular, that for any strict Nash equilibrium a^*, there exist nearby types (with respect to product topology) from models with common prior such that a^* is the unique rationalizable strategy outcome for these types. However, Lipman's (2003, 2005) results say nothing about the restrictions imposed on global properties of the whole state space. Indeed, we show that under the CPA, the set of states at which a small probability event has a large impact has an arbitrarily small probability with respect to the common prior distribution. In relation to Yildiz' (2004) result, this implies that for some games and for some strict Nash equilibria a^* (e.g., the risk-dominated equilibrium in a 2×2 coordination game), the ex ante probability of the set of types for which a^* is not uniquely rationalizable vanishes, if we require that perturbations be vanishingly small with respect to the common prior. Thus, our results quantify the difference with respect to impacts to strategic behavior on the whole state space between common and non-common prior models.

Our example in Section 2 shows that in 2×2 coordination games, where there are two strict Nash equilibria, the risk-dominated equilibrium can be contagious if the players are allowed to have heterogenous priors. On the other hand, Kajii and Morris (1997) have shown that under the CPA, this is not possible. In their terminology, the risk-dominant equilibrium is robust to incomplete information. In a companion paper (Oyama and Tercieux (2005)), we show that in generic games, a Nash equilibrium is robust to incomplete information under heterogeneous priors if and only if it is a unique action profile that survives iterative elimination of strictly dominated actions.

The remainder of the paper is organized as follows. Section 2 provides an example which illustrates why without the CPA, every strict Nash equilibrium can be contagious and how it is related to the discrepancy from the CPA. The analyses in the subsequent sections are summarized by using this example. Section 3 introduces the concept of belief potential and states our results relating the distance between priors and the belief potential of small probability events. Section 4 compares the local and the global impacts of small probability events, relating our result to the results by Lipman (2003, 2005) and Yildiz (2004). Section 5 considers an extension to the many player case.

2 Example

In this section, we illustrate the analyses in the subsequent sections with a simple example. Consider the following 2×2 coordination game with complete information which we denote by $\mathbf{g}$. There are two players, 1 and 2, each of whom has two actions L and R. Throughout the paper, for $i = 1, 2$ we write $-i$ for player $j \neq i$. The payoffs are given by

	L	R
L	p, p	$0, 0$
R	$0, 0$	$1-p, 1-p$

where $p \in (1/2, 1)$, so that (L, L) is (both Pareto-dominant and) risk-dominant. We will say that (L, L) is a strict $(1-p)$-dominant equilibrium while (R, R) is a strict p-dominant equilibrium (see Definition 3.4). As p becomes close to one, the strict Nash equilibrium (R, R) becomes "weaker".

Now, we ask the following question: For each strict Nash equilibrium $a^* = (L, L), (R, R)$ of $\mathbf{g}$, are there "perturbations" arbitrarily "close" to $\mathbf{g}$ in which a^* is played as a unique rationalizable strategy outcome? The question, of course, is not well defined unless what we mean by "perturbations" being "close" to $\mathbf{g}$ is specified.

2.1 Incomplete Information Game Perturbations

Here, as perturbations of $\mathbf{g}$ we consider incomplete information games with an information partition structure as well as the same sets of players and actions as $\mathbf{g}$, where we allow the players to have different prior beliefs. The complete information game $\mathbf{g}$ is considered as a degenerate incomplete information game. We regard a perturbed incomplete information game to be close to $\mathbf{g}$ if the event that both players know that their payoffs are given by $\mathbf{g}$ has probability close to one with respect to both players' prior distributions.

To address the question, we consider the following class of perturbations of $\mathbf{g}$. The state space Ω is given by $\{1,2\} \times \mathbb{Z}_+$. Player $i = 1, 2$ observes information partition $\mathcal{Q}_i$ which consists of (i) the event $\{(-i, 0)\}$ and (ii) all the events of the form $\{(i, k-1), (-i, k)\}$ for $k \geq 1$. Observe that this partition structure is of the same type as that in the electronic mail game of Rubinstein (1989).

The players may have different prior beliefs. For $r \in [1, \infty)$ and $\varepsilon \in (0, 1)$, let player i's prior P_i be defined by

$$\begin{aligned} P_i(i,k) &= \frac{r}{r+1} \cdot \varepsilon(1-\varepsilon)^k, \\ P_i(-i,k) &= \frac{1}{r+1} \cdot \varepsilon(1-\varepsilon)^k. \end{aligned}$$

The players have a common prior if and only if $r = 1$. Observe that for all $\omega \in \Omega$, $P_i(\omega)/P_{-i}(\omega) = r$ if $\omega = (i,k)$, while $P_{-i}(\omega)/P_i(\omega) = r$ if $\omega = (-i,k)$. As r tends to be large, the priors become distant. Hence, parameter r measures the degree of discrepancy from the CPA.

Finally, let $E_i = \{(-i, 0)\}$ and $E = E_1 \cup E_2$. The payoffs of each player i are given by $\mathbf{g}$ at all states in $\Omega \setminus E_i$, while a_i^* is a strictly dominant action for player i on event E_i, where $a^* = (a_1^*, a_2^*)$ will be (L, L) or (R, R). Verify that $P_i(E) = \varepsilon$ for each i. Let us denote this incomplete information game by $\mathcal{U}(r, \varepsilon; a^*)$.

(1) Common prior case ($r = 1$): As demonstrated by Morris, Rob, and Shin (1995),[1] if L is a dominant action for each player i at state $(-i, 0)$, then however small $\varepsilon > 0$ is, the incomplete information game $\mathcal{U}(1, \varepsilon; (L, L))$ has a unique rationalizable strategy profile, where (L, L) is played at all $\omega \in \Omega$: that is, we have a "contagion" of the risk-dominant action. On the other hand, as established by Kajii and Morris (1997), even if R is a dominant action for each player i at state $(-i, 0)$, the incomplete information game $\mathcal{U}(1, \varepsilon; (R, R))$ has a Bayesian Nash equilibrium in which (L, L) is played with high (ex ante) probability whenever ε is sufficiently small. We may say that under a common prior, the event E, however small its (ex ante)

[1] Kajii and Morris (1997) extend their argument to the countable state space case.

probability is, has an impact large enough to make the risk-dominant action contagious, but not large enough to make the risk-dominated one contagious.

(2) Non-common prior case ($r > 1$): We can show that for r sufficiently large, each action is contagious: for each equilibrium a^*, if for each player i, a_i^* is a dominant action at state $(-i, 0)$, then there exists $\bar{r}$ such that for all $r > \bar{r}$ and all $\varepsilon \in (0,1)$, the incomplete information game $\mathcal{U}(r, \varepsilon; a^*)$ has a unique rationalizable strategy profile, where a^* is played at all $\omega \in \Omega$. To see this, suppose that for each player i, R is a dominant action at $(-i, 0)$. Observe that

$$P_i(\{(i, k-1)\} | \{(i, k-1), (-i, k)\}) = \frac{r}{r + 1 - \varepsilon} \tag{2.1}$$

for all $k \geq 1$. Now, given $p \in (1/2, 1)$, let $\bar{r} = p/(1-p)$ (> 1), and take any $r \geq \bar{r}$ and $\varepsilon \in (0,1)$. Then, if player $-i$ plays R at $(i, k-1)$ in any rationalizable strategy, then it implies that player i plays R at $(-i, k)$ in any rationalizable strategy, since i assigns a probability $r/(r + 1 - \varepsilon) > p$ to the event $-i$ plays R, which makes R the unique best response. We may hence say that under non-common priors, the event E, however small its (ex ante) probability is, may have an impact large enough that any strict Nash equilibrium is contagious. The key to this result is that by increasing the value of r, we can have the relevant conditional probabilities, $P_i(\{(i, k-1)\} | \{(i, k-1), (-i, k)\})$, be as close to one as possible. The supremum of such conditional probabilities relevant to the contagion argument will be called the *belief potential* of the event E (see Definition 3.1 for the precise definition). In this particular information system with given r and ε, the belief potential of E is $r/(r + 1 - \varepsilon)$, as given by (2.1). But it will turn out that this is the "best case", in which a small probability event has the largest impact. We will show that given values of discrepancy measure, r, and small probability, ε, the value $r/(r + 1 - \varepsilon)$ is the maximum of the belief potential of a small probability event *over information systems* (see Theorem 3.4 for the precise statement). This implies that an event can have a larger impact on higher order beliefs under non-common prior than under common prior.

2.2 Belief Type Perturbations

Yildiz (2004) considers the same question from a different viewpoint. He identifies the complete information game $\mathbf{g}$ with a point $[\mathbf{g}]$ in the universal type space, i.e., the hierarchy of degenerate beliefs, and considers as "perturbations" being "close" to $\mathbf{g}$ types in the universal type space that are convergent to $[\mathbf{g}]$ with respect to product topology. His results imply, in particular, that for any strict Nash equilibrium a^* of complete information game $\mathbf{g}$, there exists a sequence of types converging to $[\mathbf{g}]$ each of which plays a^* as a unique rationalizable strategy outcome. Moreover, by appealing to Lipman's (2003) result, he shows that those converging types can be taken from models (i.e., belief-closed subspaces) with common prior. We will discuss in Section 4 the relationship between Lipman's (2003, 2005) and Yildiz' (2004) results and ours.

The latter result can be stated in our framework as follows: for any strict Nash equilibrium a^* of complete information game $\mathbf{g}$, there exists a sequence of perturbed incomplete information games $\mathcal{U}^k$ with common prior and states ω^k such that any rationalizable strategy profile of $\mathcal{U}^k$ plays a^* at ω^k, where in each $\mathcal{U}^k$, a_i^* is a strictly dominant action for player i on an event E_i^k, and at each ω^k, players know up to kth order that the payoffs are given by $\mathbf{g}$. To see this in our example, let $a^* = (R, R)$. Modifying the incomplete information game in the previous subsection with given $p \in (1/2, 1)$, $\mathcal{U}^k$ can be constructed as follows (common for all k). The state space Ω and the information partitions $\mathcal{Q}_i$ are the same as previously. Define the common prior P by

$$P(1, k) = P(1, k) = \frac{1}{2}\alpha \left(\frac{1-\varepsilon}{r}\right)^k$$

for $k \geq 0$, where

$$\alpha = 1 - \frac{1-\varepsilon}{r} \ (\geq \varepsilon),$$

r is such that $r \geq p/(1-p)$ as previously. The payoffs of each player i are given by $\mathbf{g}$ at all states in $\Omega \setminus E_i$, while a_i^* is a strictly dominant action for player i on event E_i. Finally, let $\omega^k = (1, k)$ (or $(2, k)$). Then, observing that the relevant posteriors are given by

$$P(\{(i, k-1)\}|\{(i, k-1), (-i, k)\}) = \frac{r}{r+1-\varepsilon}$$

for all $k \geq 1$, the same argument in the previous subsection shows that any rationalizable strategy plays R in every state in Ω.

Observe that in case where p is close to one, r must be large so that the probability of the event where either player has a dominant action, $P(E_1 \cup E_2)$ $(= \alpha)$, must be close to one accordingly. Thus let us consider the following modification. Let the state space be $\bar{\Omega} = \Omega \cup \{\infty\}$, and the information partition for each player i be $\bar{\mathcal{Q}}_i = \mathcal{Q}_i \cup \{\{\infty\}\}$. Define the common prior $\bar{P}$ by

$$\bar{P}(1, k) = \bar{P}(2, k) = \frac{1}{2}\varepsilon \left(\frac{1-\varepsilon}{r}\right)^k$$

for $k \geq 0$ and

$$\bar{P}(\infty) = 1 - \frac{r}{r-(1-\varepsilon)}\varepsilon,$$

where r is such that $r \geq p/(1-p)$ as previously. Note that we need to add a state, denoted ∞, in order for $\bar{P}$ to sum up to one. The payoffs of each player i are given by $\mathbf{g}$ at all states in $\bar{\Omega} \setminus E_i$, while a_i^* is a strictly dominant action for player i on event E_i. Then again, the relevant posteriors are given by

$$\bar{P}(\{(i,k-1)\}|\{(i,k-1),(-i,k)\}) = \frac{r}{r+1-\varepsilon}$$

for all $k \geq 1$, so that any rationalizable strategy plays R in every state in $\bar{\Omega} \setminus \{\infty\}$.

Now, if we require that $\bar{P}(E_1 \cup E_2)$ $(= \varepsilon)$ vanish along the sequence, then $\bar{P}(\bar{\Omega} \setminus \{\infty\})$ must vanish accordingly, which is the ex ante probability of the event that R is played as a unique rationalizable strategy action. In fact, as we will argue in Section 4, this is the case not only in this particular construction of incomplete information games, but also in any such construction. This is to be contrasted with the non-common prior case in the previous subsection, where any strict Nash equilibrium can be contagious over the state space. In this sense, if one is interested in strategic behavior on the whole state space, rather than local behavior (i.e., behavior at a particular state), then models with common priors may be significantly different from those with non-common priors.

3 Belief Potential

3.1 Information Systems and Belief Potential

An *information system* is the structure $(\Omega, (P_i)_{i=1,2}, (\mathcal{Q}_i)_{i=1,2})$, where Ω is a countable set of states, P_i is the prior distribution on Ω for player $i = 1, 2$, and $\mathcal{Q}_i$ is the partition of Ω representing the information of player i. We write $Q_i(\omega)$ for the element of $\mathcal{Q}_i$ containing ω. We assume that $P_i(Q_i(\omega)) > 0$ for all $i = 1, 2$ and $\omega \in \Omega$. Under this assumption, the conditional probability of ω' given $Q_i(\omega)$, $P_i(\omega'|Q_i(\omega))$, is well-defined by $P_i(\omega'|Q_i(\omega)) = P_i(\omega')/P_i(Q_i(\omega))$. Given an information system, we define the measure of discrepancy from the common prior case in two ways.

First, define ρ by

$$\rho\big((P_i)_{i=1,2}\big) = \max_{i \neq j} \sup_{\omega \in \Omega} \frac{P_i(\omega)}{P_j(\omega)}$$

with a convention that $q/0 = \infty$ for $q > 0$, and $0/0 = 1$. Note that $\rho\big((P_i)_{i=1,2}\big) < \infty$ only if $(P_i)_{i \in \mathcal{I}}$ have common support. The information system satisfies the CPA if and only if $\rho((P_i)_{i=1,2}) = 1$. Second, we will also consider the usual distance d_0,

$$d_0\big((P_i)_{i=1,2}\big) = \sup_{E \subset \Omega} |P_1(E) - P_2(E)|.$$

The information system satisfies the CPA if and only if $d_0((P_i)_{i=1,2}) = 0$.

We use the notion of p-belief as defined by Monderer and Samet (1989). For $p \in (0,1]$, the p-belief operator for player $i = 1,2$, $B_i^p\colon 2^\Omega \to 2^\Omega$, is defined by

$$B_i^p(E) = \{\omega \in \Omega \mid P_i(E|Q_i(\omega)) \geq p\}.$$

That is, $B_i^p(E)$ is the set of states where player i believes E with probability at least p (with respect to his own prior P_i). We define the operator $H_i^p\colon 2^\Omega \to 2^\Omega$ by

$$H_i^p(E) = B_i^p(B_{-i}^p(E)) \cup E.$$

We denote $(H_i^p)^0(E) = E$ and for $k \geq 1$, $(H_i^p)^k(E) = H_i^p((H_i^p)^{k-1}(E))$.

Denote $(H_i^p)^\infty(E) = \bigcup_{k=1}^\infty (H_i^p)^k(E)$. We follow Morris, Rob, and Shin (1995) to measure the impact of an event by the notion of *belief potential*. The belief potential of an event E is the largest probability p such that a statement of the form "player i believes with probability at least p that player $-i$ believes with probability at least p that i believes ... that the true state is in E" is true at every state in Ω.

Definition 3.1. The *belief potential of event* E, $\sigma(E)$, is

$$\sigma(E) = \max_{i=1,2} \sigma_i(E),$$

where

$$\sigma_i(E) = \sup\{p \in [0,1] \mid (H_i^p)^\infty(E) = \Omega\}.$$

Similarly, we measure the impact of an event at a given state in the following way. Event E is said to have impact p on a state ω if a statement of the form "player i believes with probability a least p that player $-i$ believes with probability at least p that i believes ... that the true state is in E" is true at ω.

Definition 3.2. Event E is said to have *impact p at state ω* if $\omega \in (H_1^p)^\infty(E) \cup (H_2^p)^\infty(E)$.

The *belief potential of event E at state ω*, $\sigma(\omega|E)$, is

$$\sigma(\omega|E) = \sup\{p \in [0,1] \mid E \text{ has impact } p \text{ at } \omega\}.$$

To illustrate these concepts, consider the information system and the event $E = \{(1,0),(2,0)\}$ in Subsection 2.1. Note that this information system satisfies $\rho\big((P_i)_{i=1,2}\big) = r$. Observe first that for each $i = 1,2$, $B_i^p(E) = \{(-i,0)\} \cup \{(i,0),(-i,0)\}$ if $p \leq r/(r+1-\varepsilon)$, and $B_i^p(E) = \{(-i,0)\}$ otherwise. Thus,

$$(H_i^p)^K(E) = \{(-i,0)\} \cup \bigcup_{k=1}^K \{(i,k-1),(-i,k)\}$$

and therefore $(H_i^p)^\infty(E) = \Omega$ if $p \le r/(r+1-\varepsilon)$, and $(H_i^p)^\infty(E) = \{(-i,0)\}$ otherwise. This implies that for this information system,

$$\sigma(E) = \frac{r}{r+1-\varepsilon}.$$

In Subsection 3.3, we will show that, given $r \ge 1$ and $\varepsilon > 0$, this is the maximum value of the belief potential of an event with probability ε over the information systems such that $\rho\big((P_i)_{i=1,2}\big) = r$.

3.2 Incomplete Information Games and p-Dominance

To relate the impact of a small probability event to the contagion of Nash equilibria played at that event (as demonstrated in Section 2), we consider incomplete information games. An incomplete information game is represented by $\mathcal{U} = (IS, (A_i)_{i=1,2}, (u_i)_{i=1,2})$, where IS is an information system as described above, A_i is the set of actions for player i, and $u_i\colon A \times \Omega \to \mathbb{R}$ is the payoff function for player i, We denote $A = A_1 \times A_2$. We assume that players know their own payoffs, i.e., for each i and every $a \in A$, $u_i(a,\cdot)$ is measurable with respect to $\mathcal{Q}_i$. A (mixed) strategy for player i is a function $s_i\colon \Omega \to \Delta(A_i)$ that is measurable with respect to $\mathcal{Q}_i$, where $\Delta(A_i)$ is the set of probability distributions over A_i. Denote by Σ_i the set of player i's strategies. For player $i = 1,2$ and action $a_i \in A_i$, we write the expected payoff against a conjecture $\nu_i \in \Delta(\Omega \times A_{-i})$ as

$$U_i(a_i,\nu_i) = \sum_{\omega\in\Omega}\sum_{a_{-i}\in A_{-i}} \nu_i(\omega,a_{-i})\, u_i(a_i,a_{-i},\omega).$$

The set of i's (pure) best responses against $\nu_i \in \Delta(\Omega \times A_{-i})$ is denoted by

$$BR_i(\nu_i) = \arg\max_{a_i\in A_i} U_i(a_i,\nu_i).$$

As the solution concept, we employ correlated interim rationalizability. For each $i = 1,2$, let $R_i^0[Q_i] = A_i$ for all $Q_i \in \mathcal{Q}_i$. Then, for each $i = 1,2$, and for $Q_i \in \mathcal{Q}_i$ and for $k = 1,2,\ldots$, define $R_i^k[Q_i]$ recursively by

$$R_i^k[Q_i] = \left\{ a_i \in A_i \;\middle|\; \begin{array}{l} \exists\, \nu_i \in \Delta(\Omega \times A_{-i}): \\ \nu_i\big(\{(\omega,a_{-i}) \mid a_{-i} \in R_{-i}^{k-1}[Q_{-i}(\omega)]\}\big) = 1; \\ \operatorname{marg}_\Omega \nu_i = P_i(\cdot|Q_i); \\ a_i \in BR_i(\nu_i) \end{array} \right\}.$$

Let $R_i^\infty[Q_i] = \bigcap_{k=1}^\infty R_i^k[Q_i]$.

Definition 3.3. A strategy $s_i \in \Sigma_i$ is a *rationalizable strategy* of player i in $\mathcal{U}$ if

$$s_i(\omega)(a_i) > 0 \Rightarrow a_i \in R_i^\infty[Q_i(\omega)]$$

for all $a_i \in A_i$ and $\omega \in \Omega$.

We also restate the definition of strict p-dominant equilibrium as defined in Morris, Rob, and Shin (1995).[2]

Definition 3.4. Let $p \in [0,1)$. Action profile $a^* \in A$ is a *strict p-dominant equilibrium* at a state ω if for each $i = 1,2$ and all $a_i \neq a_i^*$,

$$u_i(a_i^*, \pi_i, \omega) > u_i(a_i, \pi_i, \omega)$$

holds for all $\pi_i \in \Delta(A_{-i})$ with $\pi_i(a_{-i}^*) > p$.

The following proposition is a variant of the result by Morris, Rob, and Shin (1995, Theorem 5.1). Roughly, it states that if E has a belief potential equal to σ, then any p-dominant equilibrium with $p < \sigma$ can be contagious.

Proposition 3.1. *Consider an incomplete information game $\mathcal{U}$ and an event E such that $E = E_1 \cup E_2$ for some $E_i \in \mathcal{F}_i$ for each $i = 1,2$. Suppose that* (1) *E has belief potential $\sigma > 0$,* (2) *(a_1^*, a_2^*) is a strict p-dominant equilibrium at every state for some $p < \sigma$, and* (3) *for each player i, a_i^* is a strictly dominant action at each $\omega \in E_i$. Then, playing (a_1^*, a_2^*) everywhere is the unique rationalizable strategy profile of $\mathcal{U}$.*

Proof. See Appendix. ▮

3.3 Upper Bound of Belief Potential

Now we want to characterize the upper bound of the belief potential of small probability events over information systems with a given value of the discrepancy from the CPA (i.e., $\rho((P_i)_{i=1,2})$ or $d_0((P_i)_{i=1,2})$). Given an information system, write $\mathcal{F}_i$ for the sigma algebra generated by $\mathcal{Q}_i$, and denote

$$\mathcal{F}_1 \oplus \mathcal{F}_2 = \{E \subset \Omega \mid E = E_1 \cup E_2 \text{ for some } E_i \in \mathcal{F}_i \text{ for each } i = 1,2\}.$$

For $p \in (0,1]$ and $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$, we define

$$H_*^p(E) = B_1^p(E) \cup B_2^p(E).$$

We denote $(H_*^p)^k(E) = H_*^p((H_*^p)^{k-1}(E))$ for $k \geq 1$, where $(H_*^p)^0(E) = E$, and $(H_*^p)^\infty(E) = \bigcup_{k=1}^\infty (H_*^p)^k(E)$. Verify that $(H_1^p)^\infty(E) \cup (H_2^p)^\infty(E) = (H_*^p)^\infty(E)$, so that if $(H_i^p)^\infty(E) = \Omega$, then $(H_*^p)^\infty(E) = (H_i^p)^\infty(E)$. It is thus sufficient to characterize the (ex ante) probability of $(H_*^p)^\infty(E)$. The following result is the "conjugate" of Proposition 5.2 in Oyama and Tercieux (2005), where the upper bound for $P_j([(H_*^p)^\infty(E^c)]^c)$ is obtained for the many-player case. For its proof, we thus report only crucial steps in the Appendix.

[2] Here we follow the formulation of Kajii and Morris (1997, Definition 5.4).

Lemma 3.2. *For any $r \geq 1$, if $p > r/(1+r)$, then in any information system with $\rho((P_i)_{i=1,2}) = r$, any event $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ satisfies*

$$P_i((H_*^p)^\infty(E)) \leq \frac{p}{(1+r)p - r} \max\{P_1(E), P_2(E)\}$$

for all $i = 1, 2$.

Proof. See Appendix. ▮

The analogue of this result using the distance d_0 is the following.

Lemma 3.3. *Fix $\xi \geq 0$ and $\varepsilon > 0$. If $p > 1/2$, then in any information system with $d_0((P_i)_{i=1,2}) \leq \xi$, any event $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ such that $P_i(E) \leq \varepsilon$ for each $i = 1, 2$ satisfies*

$$P_i((H_*^p)^\infty(E)) \leq \frac{\xi}{2p-1} + \frac{p}{2p-1}\varepsilon$$

for all $i = 1, 2$.

Proof. See Appendix. ▮

The following two theorems are the main results of this section, which show that the belief potential of small probability events has an upper bound that is an increasing function of the discrepancy from the CPA.

Theorem 3.4. *For any $r \geq 1$ and any information system with $\rho((P_i)_{i=1,2}) = r$, if $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ and $P_i(E) \leq \varepsilon$ for each $i = 1, 2$, then*

$$\sigma(E) \leq \frac{r}{1+r-\varepsilon}.$$

Proof. Take any $q > r/(1+r-\varepsilon)$ $(> r/(1+r))$. If $\max\{P_1(E), P_2(E)\} \leq \varepsilon$, then by Lemma 3.2, for each $i = 1, 2$,

$$P_i((H_*^q)^\infty(E)) \leq \frac{q}{(1+r)q - r}\varepsilon < 1,$$

meaning that $(H_*^q)^\infty(E) \neq \Omega$, and hence $(H_i^q)^\infty(E) \neq \Omega$. This implies that $\sigma(E) \leq r/(1+r-\varepsilon)$, as claimed. ▮

Note that the upper bound given above is tight: it is attained by the event E in the information system considered in Subsection 2.1.

Theorem 3.5. *Fix $\xi \geq 0$, $\varepsilon > 0$ and any information system with $d_0(P_1, P_2) \leq \xi$, if $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ and $P_i(E) \leq \varepsilon$ for each $i = 1, 2$, then $\sigma(E) \leq \frac{1+\xi}{2-\varepsilon}$.*

Proof. Assume $\sigma_1(E) \geq \sigma_2(E)$ thus $\sigma(E) = \sigma_1(E)$. Take any p such that $(H^p_*)^\infty(E) = \Omega$. We want to show that $p \leq \frac{1+\xi}{2-\varepsilon}$. Assume that $p > 1/2$ (the case $p \leq 1/2$, is clear). By Lemma 3.3, we have: $P_i((H^p_*)^\infty(E)) \leq \frac{\xi}{2p-1} + \frac{p}{2p-1}\varepsilon$. Hence, if $p > \frac{1+\xi}{2-\varepsilon}$, we obtain that $P_i((H^p_*)^\infty(E)) < 1$ and so $P_i((H^p_1)^\infty(E)) < 1$ which yields $(H^p_1)^\infty(E) \neq \Omega$, a contradiction. ∎

One can show that the upper bound given above is asymptotically tight, while we have not been able to prove or disprove whether this bound is actually attained.

4 Common Prior vs. Non-Common Prior

Lipman (2003, 2005) shows that given any partition model with common support (and tail consistency in the case of infinite state space) and any state in the model, for any finite $N > 0$ there is a partition model with a common prior and a state in that model at which all the same facts about the world are true and all the same statements about beliefs and knowledge of order less than N are true. That is, the common prior assumption does not impose any significant restriction on finite order beliefs.

On the other hand, if one is interested in global properties of the whole state space, models with non-common priors may be quite far from any model with a common prior. In this section, we formalize this observation with the notions of global and local impact as well as with the universal type space setting.

4.1 Local vs. Global Impact of an Event

First, let us note the following lemma which is proved in a constructive way using the example in Section 2.

Lemma 4.1. *Fix a state space Ω, an event $E \subset \Omega$ such that $|E| \geq 2$. Let $p \in [0,1)$ and $\varepsilon > 0$. There exists an information system $(\Omega, (\mathcal{Q})_{i\in\mathcal{I}}, (P_i)_{i\in\mathcal{I}})$ that satisfies common support, $\rho((P_i)_{i\in\mathcal{I}}) < \infty$ and $P_i(E) \leq \varepsilon$ for all $i \in \mathcal{I}$ such that*
(1) for all $\omega \in \Omega$, $\sigma(\omega \mid E) \geq p$; and
(2) for all $N > 0$, $\bigcap_{n=1}^N (K_*)^n(E^c) \neq \emptyset$.

Proof. (1) Since Ω is countable, there exists an injection, $g\colon \Omega \to \{1,2\} \times \mathbb{Z}_+$ such that $g(E) \supset \{(1,0) \cup (2,0)\}$ (use $|E| \geq 2$). Denoting $(\hat{\Omega}, (\hat{\mathcal{Q}}_{i\in\mathcal{I}}), (\hat{P}_i)_{i\in\mathcal{I}})$ the information system in Subsection 2.1 where $r \geq 1$ is chosen so that $\sigma((\{(1,0) \cup (2,0)\}) = p$, define $(\Omega, (\mathcal{Q}_{i\in\mathcal{I}}), (P_i)_{i\in\mathcal{I}})$ as follows: for all $\omega \in \Omega$, $Q_i(\omega) = \hat{Q}_i(g(\omega))$ and $P_i(\omega) = \hat{P}_i(g(\omega))$. Note that $\sigma(E) \geq \sigma(g^{-1}(\{(1,0) \cup (2,0)\}) = \hat{\sigma}(\{(1,0) \cup (2,0)\}) = p$. This implies that for all $\omega \in \Omega$, $\sigma(\omega \mid E) \geq p$.

(2) is easy to check. ∎

This lemma implies the following weaker result.

Proposition 4.2. *Fix a state space Ω, an event $E \subset \Omega$ such that $|E| \geq 2$ and $\omega \in \Omega$. Let $p \in [0,1)$ and $\varepsilon > 0$. For all $N > 0$, there exists an information system $(\Omega, (\mathcal{Q})_{i\in\mathcal{I}}, (P_i)_{i\in\mathcal{I}})$ that satisfies common support, $\rho((P_i)_{i\in\mathcal{I}}) < \infty$ and $P_i(E) \leq \varepsilon$ for all $i \in \mathcal{I}$ such that*
(1) $\sigma(\omega \mid E) \geq p$; and
(2) $\omega \in \bigcap_{n=0}^{N} (K_*)^n(E^c)$.

Roughly the above results show that, under heterogenous priors, given any state space and any (small probability) event E, there are a partition and priors so that E can have an arbitrarily large impact on *all states* in the state space. Importantly, this is true even on states where it is known at an arbitrarily large order that E did not occur. We say in this case that E has a large global impact. It is clear that if an event has a large global impact, then it must have a large local impact, i.e., a large impact on all states in the state space. However, we will see that the converse is not true. More specifically, under common priors, given any state space, any state of the world, ω, in this state space, any (small probability) event E can have an arbitrarily large impact on the given state even if we require that at ω it is known at arbitrarily high order that E did not occur. But as we have seen, the global impact of E cannot be arbitrarily large under a common prior. The following results aim to clarify this point and the relationship between local and global impact of an event.

Our point that under common prior, the local impact of any small probability event can be arbitrarily large draws on a result by Lipman (2003) mentioned above. To discuss this point, let us first state a proposition on the local impact of small probability events under the common prior assumption.

Proposition 4.3. *Fix a state space Ω and an event $E \subset \Omega$. Let $p \in [0,1)$ and $\varepsilon > 0$. For all $N > 0$, there exist an information system $(\Omega, (\mathcal{Q}_i)_{i\in\mathcal{I}}, P)$ that satisfies the CPA and $P(E) \leq \varepsilon$, and $\omega \in \Omega$ such that*
(1) $\sigma(\omega \mid E) \geq p$; and
(2) $\omega \in \bigcap_{n=1}^{N} (K_*)^n(E^c)$.

Remark 4.1. Note that, we could have proved this result using Lipman's main theorem. However, while Lipman's theorem would have been sufficient, it is not necessary for our purpose. Let us briefly explain this point. Given an information system without common prior, Lipman builds an information system with a common prior where beliefs of players (at a finite order) about *any* event are *exactly* the same as in the information system without common prior. However, for our purpose we only need to match beliefs of players (at a finite order) that rely on the small probability event E. In addition, we do not need beliefs to be exactly matched (see the definition of the operator $(H_i^p)^k(E)$). These two differences allow us to provide a simple proof and to change only priors keeping the state space Ω contrary to what we would have obtained using Lipman's result.

However, the main point of this section is that under common prior, the set of states on which a given small probability event has a "large" impact is small with respect to prior probabilities. The following lemma formalizes this point.

Lemma 4.4. *Let $r \geq 1$ and $p > r/(r+1)$. For any $\delta > 0$, there exists $\varepsilon > 0$ such that for any information system IS with $\rho((P_i)_{i\in\mathcal{I}}) = r$ and any event $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ such that $P_i(E) \leq \varepsilon$ for all $i \in \mathcal{I}$, we have*

$$P_i(\{\omega \in \Omega \mid \sigma(\omega|E) \geq p\}) \leq \delta$$

for all $i \in \mathcal{I}$.

Proof. Given $p > r/(r+1)$ and $\delta > 0$, set $\varepsilon = \delta\{(1+r)p - r\}/p$. Then by Lemma 3.2, we have for each $i = 1, 2$,

$$P_i((H_i^p)^\infty(E)) \leq P_i((H_*^p)^\infty(E)) \leq \frac{p}{(1+r)p - r}\varepsilon \leq \delta,$$

as claimed. ▮

Obtained as a corollary of the previous lemma, we have the following main result of this subsection.

Proposition 4.5. *For any $p > 1/2$ and any $\delta > 0$, there exists $\varepsilon > 0$ such that for any information system $\overline{IS}$ that satisfies the CPA and any event $\bar{E} \in \bar{\mathcal{F}}_1 \oplus \bar{\mathcal{F}}_2$ such that $\bar{P}(\bar{E}) \leq \varepsilon$, we have*

$$\bar{P}(\{\omega \in \bar{\Omega} \mid \sigma(\omega|\bar{E}) \geq p\}) \leq \delta.$$

In terms of contagion of Nash equilibria, while any strict Nash equilibrium at a small probability event can spread in some partition model with non-common priors, it may not be the case for partition models with a common prior. Indeed, in 2×2 coordination games, the risk-dominated equilibrium cannot spread from a small probability event when we assume the existence of a common prior, as shown by Kajii and Morris (1997).

However, our latter result on the local impact of small probability events under the common prior assumption allows us to show the following. Assume some strict Nash equilibrium is played at a small probability event E. Then, given any arbitrarily large number N, we can find an information system and a state of the world ω such that at ω, it is mutually known at order N that E did not occur, but the strict Nash equilibrium is played at any rationalizable strategy profile at ω.

To understand this point, let us state the following proposition.

Proposition 4.6. *Fix a state space Ω, an event $E \subset \Omega$, and $p \in [0,1)$. Let $(A_i)_{i=1,2}$ and $(u_i)_{i=1,2}$, $u_i\colon A \times \Omega \to \mathbb{R}$, be such that (1) $(a_1^*, a_2^*) \in A$ is a p-dominant equilibrium at each state ω; (2) (a_1^*, a_2^*) is strictly dominant at each state $\omega \in E$. Then, for all $N > 0$ and $\varepsilon > 0$, there exists an information system IS such that any rationalizable strategy profile σ of the incomplete information game $\mathcal{U} = (IS, (A_i)_{i=1,2}, (u_i)_{i=1,2})$ satisfies $\sigma(\omega) = a^*$ for all $\omega \in \bigcap_{n=1}^{N'} K_*(E^c)$ for any $N' \leq N$.*

To summarize, under heterogeneous priors, both the local and the global impact of any small probability event can be arbitrarily large, whereas under common prior, only the local impact can be arbitrarily large.

This distinction will allow us to shed light on a result of Yildiz (2004) which shows that for any type in the universal type space, there exists arbitrarily close types where rationalizability yields a unique action profile, and moreover, such a type can always be taken from a model with a common prior.

We will first see the connection between this point and the impact of an event. Then, we will see that if one is interested in such a statement on a whole model (and not only on a specific type) dropping the common prior assumption is necessary (and sufficient).

4.2 Embedding our Results in the Universal Type Space

In this subsection, we embed our results in the universal type space. This allows us to compare our results with those of Yildiz (2004). Yildiz (2004) has shown that, under standard assumptions, for any type in the universal type space, there exists nearby types where a unique rationalizable action profile is played. This result is obtained irrespective of whether one assumes that players share a common prior. We claim that if one is interested in the behavior of players not only at a given type but on a whole model, then dropping the common prior assumption is crucial to obtain Yildiz' (2004) type of statement.

4.2.1 The Universal Type Space Setting

Let Θ be a compact metric space of payoff-relevant parameters θ. To make things simple, we assume that $\Theta = \Theta_1 \times \Theta_2$ where $\Theta_i = [0,1]^A$. We write $\Delta(X)$ for the set of probability measures on the Borel field of any topological space X. When X is a set of probability measures, it will be endowed with the weak* topology.

Define recursively $X_0 = \Theta$, $X_1 = \Delta(X_0)$, $X_2 = \Delta(X_0 \times X_1)$ .... Let us now describe a type in the setting of the universal type space. A type of a player i is an infinite hierarchy of beliefs $t_i = (t_i^1, t_i^2, \ldots)$ where $t_i^1 \in X_1$ is a probability distribution on Θ, representing the (first order) beliefs of player i about Θ, $t_i^2 \in X_2$ is a probability distribution representing the (second order) beliefs of player i, i.e., his beliefs about Θ as well as his beliefs about the other player's beliefs over Θ, and so on. We also assume that it is common knowledge that the beliefs are coherent.[3] Denote the set of all such types by T_i^*, and let $T^* = T_1^* \times T_2^*$.

For each type t_i, let $\kappa_{t_i} \in \Delta(\Theta \times T_{-i}^*)$ be the unique probability distribution that represents the beliefs of t_i about (θ, t_{-i}). Mertens and Zamir (1985) have shown that the mapping $t_i \mapsto \kappa_{t_i}$ is a homeomorphism. A set $T \subset T^*$ is said to be *model* if for each $t_i \in T_i$, $\kappa_{t_i}(\Theta \times T_{-i}) = 1$.

We now describe the connection between the partition model setting and the universal type space setting. In particular, we show how a partition model together with a state of the world *induces* a type in the universal type space. Since a parameter space (namely Θ) has been added to the description of the basic uncertainty, we will now need to refer to it in the definition of a partition model. Hence a partition model $\mathcal{M}$ is consists of an information system $[\Omega, ((Q_i)_{i\in I}), (P_i)_{i\in I})]$ together with a function $f\colon \Omega \to \Theta$ where $f(\omega)$ is the value of the unknown parameter at state ω.

Any partition model together with any state ω in that model uniquely identifies a particular type in the universal type space denoted $t[\omega]$ by the the so-called unravelling procedure.[4] Let us briefly describe the procedure. Given a world ω in a partition model, we can identify each player's first order beliefs at ω. Denote player i's first order beliefs at world ω by $t_i^1[\omega] \in X_1$. For each measurable set $B \subset \Theta$, we define

$$t_i^1[\omega](B) = P_i(f^{-1}(B) \mid Q_i(\omega)).$$

In the same way, player i's second order beliefs at ω, say, $t_i^2[\omega] \in X_2$ is defined by

$$t_i^2[\omega](B) = P_i(\{\omega' \mid (f(\omega'), t_{-i}^1[\omega']) \in B\} \mid Q_i(\omega))$$

for each measurable set $B \subset \Theta \times \Delta(\Theta)$. Continuing recursively, we can define $t_i^n[\omega]$ for every n. Let $t_i[\omega] = (t_i^1[\omega], t_i^2[\omega], ...)$ and $t[\omega] = (t_1[\omega], t_2[\omega])$.

4.2.2 Product Topology

The usual topology used in the universal type space, and in particular the one used by Yildiz (2004), is the product topology. Let us review convergence of a sequence of types in the universal type space with respect to product topology. For our purpose, it is sufficient to restrict our attention to convergence toward a complete information type. For a complete information game $\mathbf{g}$, we will write $t_{\mathbf{g}}$ for the complete information type where $\mathbf{g}$

[3] A type of player i, t_i is coherent if for every $n \geq 2$, $\text{marg}_{X_{n-2}} t_i^n = t_i^{n-1}$.

[4] The converse of this statement (i.e., that any type in the universal type space can be constructed using the unravelling procedure from some partition model) is true as long as we allow for a less restrictive class of partitions models from those we use in this paper. See for instance Brandenburger and Dekel (1993).

is common knowledge. Since complete information types, considered as singleton sets, can also be seen as complete information models, this will allow us to compare convergence of a sequence of types toward $t_{\mathbf{g}}$ and convergence of sequences of models toward $t_{\mathbf{g}}$.

Definition 4.1. Let $t_{\mathbf{g}}$ be a complete information type where $\mathbf{g}$ is common knowledge. $t_m \to t_{\mathbf{g}}$ as $m \to \infty$ if for each i and k, $t^k_{i,m} \to t^k_{i,\mathbf{g}}$ as $m \to \infty$.

Product topology has a natural interpretation. Indeed, it captures the idea that we cannot observe infinite hierarchy of beliefs. Suppose that we consider a complete information type $t_{\mathbf{g}}$ in the universal type space. Suppose also that for some k we have made some noisy observation about the first k order of beliefs, and for each $k' \leq k$, we find an open neighborhood $\mathcal{N}^k$ of $t^{k'}_{\mathbf{g}}$ (with respect to weak topology on probability distributions). Then, the open neighborhoods of $t_{\mathbf{g}}$ are those sets of types where for some k, first k order of beliefs are in these $\mathcal{N}^{k'}$ for $k' \leq k$.

To relate the statements of the previous subsection to statements in the universal type space as in Yildiz (2004), we will use the following simple observation.

Observation 4.7. *Consider a sequence $(\mathcal{M}^k, \omega^k)$ of partition models and states in this model. If, for all $N > 0$, there exists K such that for all $k \geq K$ $(\mathcal{M}^k, \omega^k)$ satisfies $\omega^k \in \bigcap_{n=0}^{N} (K_*)^n (f^{-1}(\mathbf{g}))$, then we have $t[\omega^k] \to t_{\mathbf{g}}$.*

We will also use a notion of convergence of models toward complete information models. First, we provide a definition that will allow us to extract a (set of) measure(s) from a given model.

Definition 4.2. Let T be a model. A profile of priors over T, $(P_i)_{i \in \mathcal{I}}$, is said to be *belief consistent* with T if for all i and all type t_i, $\kappa_{t_i} = P_i(\cdot \mid \{t_i\} \times T_{-i})$.

We only consider models that are irreducible.

Definition 4.3. Let $\{t_{\mathbf{g}}\}$ be a complete information model where $\mathbf{g}$ is common knowledge. $T_m \to \{t_{\mathbf{g}}\}$ as $m \to \infty$ if for all $\delta > 0$,

$$P_{i,m}(\{t \in T_m \mid |t^1_{i,m} - t^1_g| < \delta\}) \to 1 \quad \text{as } m \to \infty$$

for any sequence of profile of belief consistent priors $(P_{i,m})_{i \in \mathcal{I}}$ of T_m.

To discuss the relationship between our results and Yildiz' result on generic uniqueness, we first define the impact of an event in the universal type space which is the natural corresponding definition of belief potential in the universal type space.

Consider $\hat{\Theta} \subset \Theta$. Define recursively the sequence of set of distributions:

$$\Pi^p_1(\hat{\Theta}) = \{\pi \in X_1 \mid \pi(\hat{\Theta}) \geq p\}$$

$$\Pi_2^p(\hat{\Theta}) = \{\pi \in X_2 \mid \text{marg}_{X_1}\, \pi(\Pi_1^p(\hat{\Theta})) \geq p\}$$

and for all $k \geq 2$,

$$\Pi_k^p(\hat{\Theta}) = \{\pi \in X_k \mid \text{marg}_{X_{k-1}}\, \pi(\Pi_{k-1}^p(\hat{\Theta})) \geq p\}.$$

Definition 4.4. $\hat{\Theta}$ has impact p on type t if there exist $i \in \{1, 2\}$ and K such that

$$\text{marg}_{X_{K-1}}\, t_i^K(\Pi_{K-1}^p(\hat{\Theta})) \geq p.$$

We say that $\hat{\Theta}$ has impact p on the model T if it has impact p on any $t \in T$.

Proposition 4.8. *Fix any $\mathbf{g} \in \Theta$, $p \in (0, 1)$ and $\hat{\Theta} \subset \Theta$. There exists a sequence of couples $\{(T^k, t^k)\}_{k=0}^{\infty}$ where for each k, T^k is a model satisfying the common prior assumption and t^k is a type in T^k such that $t_k \to t_{\mathbf{g}}$ and where for each k, $\hat{\Theta}$ has impact p on t_k.*

Proof. By Proposition 4.3 together with Observation 4.7. ∎

Note that it is easy to show that T^k can indeed be chosen to be finite. This claim is equivalent to the following statement. Fix $\hat{\Theta}$ and any $p < 1$. For any open neighborhood of any complete information type $t_{\mathbf{g}}$, there exists a type t coming from a model with common prior so that $\hat{\Theta}$ has impact p on t.

However as we have claimed earlier, in the model to which t belongs, the set of types where $\hat{\Theta}$ has impact $p > 1/2$ is assigned probability close to zero by the common prior as long as this prior assigns a small probability to $\hat{\Theta}$.

However allowing for heterogeneous priors enables us to obtain a result that explicitly refers to models. We want to underline that when the object of interest for a modeler is a model, then a statement of the type of proposition 4.8 holds only when we allow for heterogeneous priors.

Proposition 4.9. *Fix any $g \in \Theta$, $p \in (0, 1)$ and $\hat{\Theta} \subset \Theta$. There exists a sequence of models $\{T^k\}_{k=0}^{\infty}$ where for each k: $T_k \to \{t_{\mathbf{g}}\}$ and $\hat{\Theta}$ has impact p on T_k.*

5 Many-Player Extension

In this section, we briefly discuss an extension of belief potential to the case of many players. We denote by $\mathcal{I} = (1, 2, \ldots, I)$ the finite set of players. As previously, an information system $(\Omega, (P_i)_{i\in\mathcal{I}}, (\mathcal{Q}_i)_{i\in\mathcal{I}})$ consists of a countable state space Ω, the prior distribution P_i and the information partition $\mathcal{Q}_i$ for each player $i \in \mathcal{I}$. Denote by $\mathcal{F}_i$ the sigma algebra generated by $\mathcal{Q}_i$.

Let $\mathbf{E}$ be a profile $(E_1, \ldots, E_I)$ where $E_i \in \mathcal{F}_i$. Define

$$\left(\widehat{\mathbf{B}}^p\right)_i(\mathbf{E}) = B_i^p\left(\bigcap_{j\neq i} E_j\right),$$

and $\widehat{\mathbf{B}}^p(\mathbf{E}) = \left((\widehat{\mathbf{B}}^p)_i(\mathbf{E})\right)_{i\in\mathcal{I}}$.

Then, define $\left\{(\widehat{\mathbf{H}}^p)^k(\mathbf{E})\right\}_{k=0}^{\infty}$ recursively by $(\widehat{\mathbf{H}}^p)_i^0(\mathbf{E}) = E_i$ and for $k \geq 1$,

$$\begin{aligned}(\widehat{\mathbf{H}}^p)_i^k(\mathbf{E}) &= B_i^p\left(\bigcap_{j\neq i}(\widehat{\mathbf{H}}^p)_j^{k-1}(\mathbf{E})\right) \cup (\widehat{\mathbf{H}}^p)_i^{k-1}(\mathbf{E})\\ &= (\widehat{\mathbf{B}}^p)_i\left((\widehat{\mathbf{H}}^p)^{k-1}(\mathbf{E})\right) \cup (\widehat{\mathbf{H}}^p)_i^{k-1}(\mathbf{E}).\end{aligned}$$

Definition 5.1. Let $\mathbf{E} = (E_1, \ldots, E_I)$ where $E_i \in \mathcal{F}_i$. The *belief potential of event profile* $\mathbf{E}$, $\sigma(\mathbf{E})$, is

$$\sigma(\mathbf{E}) = \max_{i\in\mathcal{I}} \min_{j\neq i} \sigma_j(\mathbf{E}),$$

where

$$\sigma_i(\mathbf{E}) = \sup\left\{p \in [0,1] \mid \bigcup_{k=0}^{\infty}(\widehat{\mathbf{H}}^p)_i^k(\mathbf{E}) = \Omega\right\}.$$

We want to relate the belief potential to the p-dominance of Nash equilibria. Incomplete information games $\mathcal{U} = (IS, (A_i)_{i\in\mathcal{I}}, (u_i)_{i\in\mathcal{I}})$ are defined analogously to the two player case, where IS is an information system as described above. We denote $A = \prod_{i\in\mathcal{I}} A_i$ and $A_{-i} = \prod_{j\neq i} A_j$.

Definition 5.2. Let $p \in [0, 1)$. Action profile $a^* \in A$ is a *strict p-dominant equilibrium* at a state ω if for each $i \in \mathcal{I}$ and all $a_i \in A_i$,

$$u_i(a_i^*, \pi_i, \omega) > u_i(a_i, \pi_i, \omega)$$

holds for all $\pi_i \in \Delta(A_{-i})$ with $\pi_i(a_{-i}^*) > p$.

We have the following.

Proposition 5.1. *Let $\mathbf{E} = (E_1, \ldots, E_I)$ where $E_i \in \mathcal{F}_i$. Suppose that* (1) $\mathbf{E}$ *has belief potential $\sigma > 0$,* (2) *a^* is a strict p-dominant equilibrium at every state for some $p < \sigma$, and* (3) *for each player $i \in \mathcal{I}$, a_i^* is a strictly dominant action at every $\omega \in E_i$. Then, playing a^* everywhere is the unique rationalizable strategy outcome.*

Appendix

A.1 Proof of Proposition 3.1

Given (a_1^*, a_2^*), denote $\Omega_i = \{\omega \in \Omega \mid R^{\infty}[Q_i(\omega)] = \{a^*\}\}$ for each $i = 1, 2$.

Lemma A.1.1. *Consider an incomplete information game $\mathcal{U}$ and an event E such that $E = E_1 \cup E_2$ for some $E_i \in \mathcal{F}_i$ for each $i = 1, 2$. Suppose that* (i) *(a_1^*, a_2^*) is a strict p-dominant equilibrium at every state, and* (ii) *for each player i, a_i^* is a strictly dominant action at each $\omega \in E_i$. Then, for any $q > p$, $(H_i^q)^{\infty}(E) \subset \Omega_i \cup E_{-i}$ for each i.*

Proof. Fix $q > p$ and $i = 1, 2$. We show by induction that $(H_i^q)^k(E) \subset \Omega_i \cup E_{-i}$ for all k. This is true for $k = 0$ since $E_i \subset \Omega_i$ by assumption (ii). Assume now that it is true for $k-1$, that is, $(H_i^q)^{k-1}(E) \subset \Omega_i \cup E_{-i}$. Then we have

$$B_{-i}^p((H_i^q)^{k-1}(E)) \subset B_{-i}^p(\Omega_i \cup E_{-i}) = B_{-i}^p(\Omega_i) \cup E_{-i},$$

where the equality follows from $E_{-i} \in \mathcal{F}_{-i}$. Since $B_{-i}^p(\Omega_i) \subset \Omega_{-i}$ and $E_{-i} \subset \Omega_{-i}$ by assumptions (i) and (ii), respectively, it follows that $B_{-i}^p((H_i^q)^{k-1}(E)) \subset \Omega_{-i}$. Again by (i) and (ii) as well as the induction hypothesis, we have

$$\begin{aligned}(H_i^q)^k(E) &= B_i^q(B_{-i}^p((H_i^q)^{k-1}(E))) \\ &\subset B_{-i}^p((H_i^q)^{k-1}(E)) \subset \Omega_i \cup (\Omega_i \cup E_{-i}) = \Omega_i \cup E_{-i},\end{aligned}$$

as desired. ▮

Proof of Proposition 3.1. Let $\sigma(E) = \sigma_i(E)$. By Lemma A.1.1, assumptions (2) and (3) imply that $(H_i^\sigma)^\infty(E) \subset \Omega_i \cup E_{-i}$. But (1) implies that $(H_i^\sigma)^\infty(E) = \Omega$, so that $\Omega_i \cup E_{-i} = \Omega$. But since $\Omega \setminus \Omega_i \in \mathcal{F}_i$, we have $B_i^\sigma(\Omega \setminus \Omega_i) = \Omega \setminus \Omega_i$, and therefore using (2) together with $\Omega \setminus \Omega_i \subset E_{-i}$, we have $\Omega \setminus \Omega_i \subset \Omega_i$, which implies $\Omega_i = \Omega$ (and $\Omega \setminus \Omega_i = \emptyset$). Also, by (2) it must be that $\Omega_{-i} = \Omega$. ▮

A.2 Proof of Lemma 3.2

We first note the following, which is essentially equivalent to Lemma A in Kajii and Morris (1997).

Lemma A.2.1. *Let $p > 0$. For any event E and player i, if $F_i \in \mathcal{F}_i$ and $F_i \subset B_i^p(E)$, then $P_i(F_i \setminus E) \leq ((1-p)/p)P_i(F_i \cap E)$.*

Fix $r \geq 1$, and consider any information system with $\rho((P_i)_{i=1,2}) = r$ and any event $E = E_1 \cup E_2$, each $E_i \in \mathcal{F}_i$. In the following, we want to obtain an upper bound for $P_j((H_*^p)^K(E))$.

Let $E_i^0 = E_i$ and $E^0 = E_1^0 \cup E_2^0$. Given $K \geq 1$ and $p \in (0, 1]$, define $\{E_1^k, E_2^k, E^k\}_{k=1}^{K+1}$ recursively by

$$E_i^k = B_i^p(E^{k-1}), \quad E^k = E_1^k \cup E_2^k.$$

Then, $(H_*^p)^K(E) = E^K$. Let $D_i^0 = E_i^0$ and $D_i^k = E_i^k \setminus E_i^{k-1}$ for $k = 1, \ldots, K+1$. Observe that $\{D_i^k\}_{k=0}^{K+1}$ is a partition of Ω, which is coarser than $\mathcal{Q}_i$.

For $i, j = 1, 2$, let $x_i(j, 0) = 0$ and

$$x_i(j, k) = \sum_{\ell=1}^{k} P_i(D_j^\ell \setminus E^{\ell-1})$$

for $k = 1, \ldots, K$. Note that

$$P_i((H_*^p)^K(E)) \leq P_i(E) + x_i(1, K) + x_i(2, K).$$

By using Lemma A.2.1, we have the following.

Lemma A.2.2. *For all $k = 1, \ldots, K$ and $i = 1, 2$,*

$$x_i(i, k) \leq \frac{1-p}{p} x_i(-i, k-1) + \frac{1-p}{p} P_i(E_{-i}^0 \setminus E_i^0).$$

Now, $\rho((P_i)_{i=1,2}) = r$ implies that $x_j(i, k) \leq r x_{-j}(i, k)$. Thus by Lemma A.2.2, we have the following. Let $x_i(k) = x_i(1, k) + x_i(2, k)$.

Lemma A.2.3. *For all $k = 1, \ldots, K$ and $i = 1, 2$,*

$$x_i(k) \leq \frac{r(1-p)}{p} x_{-i}(k-1) + \frac{r(1-p)}{p} P_{-i}(E^0).$$

By recursively using Lemma A.2.3, we obtain the upper bound of $P_i((H_*^p)^K(E))$.

Lemma A.2.4. *In any information system with $\rho((P_i)_{i=1,2}) = r$, any event $E \in \mathcal{F}_1 \oplus \mathcal{F}_2$ satisfies*

$$P_i((H_*^p)^K(E)) \leq \max\{P_1(E), P_2(E)\} \sum_{k=0}^{K} \left\{ \frac{r(1-p)}{p} \right\}^k \tag{A.1}$$

for all $i = 1, 2$.

We are now in a position to prove Lemma 3.2. It remains to consider the limit of the right hand side of (A.1) as $K \to \infty$. This is where the assumption that $p > r/(1+r)$ is used.

Proof of Lemma 3.2. If $p > r/(1+r)$, or $r(1-p)/p < 1$, then the right hand side of (A.1), $\sum_{k=0}^{K} \{r(1-p)/p\}^k$, converges to

$$\frac{1}{1 - \dfrac{r(1-p)}{p}} = \frac{p}{(1+r)p - r}$$

as $K \to \infty$. Hence, by Lemma A.2.4 we have the desired inequality. ∎

A.3 Proof of Lemma 3.3

The proof of the present Lemma is closely related to the proof of Lemma 3.2. We state it for completeness. Let us restate the labelling of state space we have used in the proof of Lemma 3.2.

Let $E_i^0 = E_i$ and $E^0 = E_1^0 \cup E_2^0$. Define $\{E_1^k, E_2^k, E^k\}_{k=1}^{\infty}$ recursively by $E_i^k = B_i^p(E^{k-1})$, $E^k = E_1^k \cup E_2^k$. Set also $E_i^{-1} = \emptyset$ and let $D_i^k = E_i^k \setminus E_i^{k-1}$. Note again that for all $k \geq 0$, $(H_*^p)^k(E) = E^k$.

Proof of Lemma 3.3.

$$
\begin{aligned}
P_i\left((H^p_*)^\infty(E)\right) &= P_i\left(\bigcup_{i\in\{1,2\}}\bigcup_{l=0}^{\infty}\left(D^l_i\backslash E^{l-1}\right)\right) \text{ by Lemma ??;}\\
&\le P_i\left(\bigcup_{l=1}^{\infty}\left(D^l_i\backslash E^{l-1}\right)\right)+P_i\left(\bigcup_{l=1}^{\infty}\left(D^l_{-i}\backslash E^{l-1}\right)\right)+P_i(E^0)\\
&= \sum_{l\ge 1}P_i\left(D^l_i\backslash E^{l-1}\right)+\sum_{l\ge 1}P_i\left(D^l_{-i}\backslash E^{l-1}\right)+P_i(E^0)\\
\text{(since for each } i\text{, and } l &\ne l',\ D^l_i\backslash E^{l-1} \text{ and } D^{l'}_i\backslash E^{l'-1} \text{ are disjoints)}\\
&\le \frac{1-p}{p}\left(\sum_{l\ge 1}P_i\left(D^l_i\cap E^{l-1}\right)+\sum_{l\ge 1}P_{-i}\left(D^l_{-i}\cap E^{l-1}\right)\right)+\xi+\varepsilon\\
&\le \frac{1-p}{p}\left(\sum_{l\ge 1}P_i\left(D^l_i\cap E^{l-1}\right)+\sum_{l\ge 1}P_i\left(D^l_{-i}\cap E^{l-1}\right)\right)+\frac{\xi}{p}+\varepsilon\\
&= \frac{1-p}{p}\left(P_i\left(\bigcup_{l=1}^{\infty}(D^l_i\cap E^{l-1})\right)+P_i\left(\bigcup_{l=1}^{\infty}(D^l_{-i}\cap E^{l-1})\right)\right)+\frac{\xi}{p}+\varepsilon\\
\text{(since for each } i\text{, and } l &\ne l',\ D^l_i\cap E^{l-1} \text{ and } D^{l'}_i\cap E^{l'-1} \text{ are disjoints)}\\
&= \frac{1-p}{p}P_i\left(\bigcup_{i\in\{1,2\}}\bigcup_{l=1}^{\infty}\left(D^l_i\cap E^{l-1}\right)\right)+\frac{\xi}{p}+\varepsilon \text{ by Lemma ??}\\
&\le \frac{1-p}{p}P_i\left((H^p_*)^\infty(E)\right)+\frac{\xi}{p}+\varepsilon \text{ by Remark ??.}
\end{aligned}
$$

This yields

$$P_i((H^p_*)^\infty(E)) \le \frac{\xi}{2p-1}+\frac{p}{2p-1}\varepsilon$$

and completes the proof.

▮

References

Aumann, R. J. (1976). “Agreeing to Disagree,” *Annals of Statistics* **4**, 1236-1239.

Brandenburger, A. and E. Dekel (1993). “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory* **59**, 189-198.

Carlsson, H. and E. van Damme (1993). “Global Games and Equilibrium Selection,” *Econometrica* **61**, 989-1018.

Kajii, A. and S. Morris (1997). "The Robustness of Equilibria to Incomplete Information," *Econometrica* **65**, 1283-1309.

Lipman, B. L. (2003). "Finite Order Implications of Common Priors," *Econometrica* **71**, 1255-1267.

Lipman, B. L. (2005). "Finite Order Implications of Common Priors in Infinite Models," mimeo.

Mertens, J.-F. and S. Zamir (1985). "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory* **14**, 1-29.

Milgrom, P. and N. Stokey (1982). "Information, Trade, and Common Knowledge," *Journal of Economic Theory* **26**, 17-27.

Monderer, D. and D. Samet (1989). "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior* **1**, 170-190.

Morris, S., R. Rob, and H. S. Shin (1995). "p-Dominance and Belief Potential," *Econometrica* **63**, 145-157.

Oyama, D. and O. Tercieux (2005). "Strategic Implications of (Non-)Common Priors II: Robustness of Equilibria," mimeo.

Rubinstein, A. (1989). "The Electronic Mail Game: Strategic Behavior under 'Almost Common Knowledge'," *American Economic Review* **79**, 385-391.

Yildiz, M. (2004). "Generic Uniqueness of Rationalizable Actions," mimeo.