

Robust equilibria under non-common priors

Daisuke Oyama^{a,*}, Olivier Tercieux^b

^a Graduate School of Economics, Hitotsubashi University, Naka 2-1, Kunitachi, Tokyo 186-8601, Japan

^b Paris School of Economics and CNRS, 48 Boulevard Jourdan, 75014 Paris, France

Received 25 March 2009; accepted 15 September 2009

Available online 24 October 2009

Abstract

This paper considers the robustness of equilibria to a small amount of incomplete information, where players are allowed to have heterogeneous priors. An equilibrium of a complete information game is *robust to incomplete information under non-common priors* if for every incomplete information game where each player's prior assigns high probability on the event that the players know at arbitrarily high order that the payoffs are given by the complete information game, there exists a Bayesian Nash equilibrium that generates behavior close to the equilibrium in consideration. It is shown that for generic games, an equilibrium is robust under non-common priors if and only if it is the unique rationalizable action profile. Set-valued concepts are also introduced, and for generic games, a smallest robust set is shown to exist and coincide with the set of a posteriori equilibria.

© 2009 Elsevier Inc. All rights reserved.

JEL classification: C72; D82

Keywords: Incomplete information; Robustness; Common prior assumption; Higher order belief

1. Introduction

One important research program in game theory has been to examine the robustness of Nash equilibria of a given complete information game to incomplete information, i.e., whether the predictions generated by Nash equilibria are still valid in “nearby” incomplete information games obtained by perturbing the complete information game (see, e.g., Fudenberg, Kreps, and

* Corresponding author.

E-mail addresses: oyama@econ.hit-u.ac.jp (D. Oyama), tercieux@pse.ens.fr (O. Tercieux).

URLs: <http://www.econ.hit-u.ac.jp/~oyama> (D. Oyama), <http://www.pse.ens.fr/tercieux> (O. Tercieux).

Levine [8] and Kajii and Morris [10]). There, most existing approaches (Kajii and Morris [10], Ui [22], and Morris and Ui [17], among others) assume that players share a common prior belief in perturbed incomplete information games, as do most work in other fields in game theory and information economics. In this paper, we characterize equilibria that remain valid in perturbed incomplete information games dropping the common prior assumption (CPA, henceforth), i.e., allowing players to have heterogeneous subjective prior beliefs. This enables us to assess the role of the CPA in examining the robustness of equilibria to incomplete information.

The results in this paper are summarized as follows. First, we show that an equilibrium of a generic complete information game is robust under non-common priors if and only if it is the unique rationalizable action profile. Second, we consider set-valued robustness and show that, for generic games, a smallest robust set exists and coincides with the set of a posteriori equilibria.¹ These imply that the robustness test that allows non-common priors is substantially stronger than the one with the CPA as considered by Kajii and Morris [10]. We also examine the effect of imposing an upper bound on the degree of heterogeneity of the prior beliefs in incomplete information perturbations, and relate our results to those by Weinstein and Yildiz [24], who obtain an analogous characterization for robustness to perturbation of interim beliefs, but show that their characterization remains unchanged even with the CPA. We demonstrate, in contrast, that unbounded heterogeneity in priors is essential for our characterization of ex ante robustness.

To explain our framework, consider an analyst who plans to model some strategic situation by a particular complete information game \mathbf{g} . While he believes that the environment is well described by this game, he is also aware that each player faces a small amount of payoff uncertainty, so that players may play some incomplete information game in which their ex ante subjective payoffs are close to their payoffs in \mathbf{g} , where he does not assume that the players share a common prior. We want to ask whether the analyst's prediction based on the complete information game is not different from the ex ante average equilibrium behavior of the real incomplete information game.²

Our key assumption to formalize closeness between incomplete information games and the complete information game \mathbf{g} is that the analyst is restricted to incomplete information games where with high ex ante (subjective) probability, players know that the game is \mathbf{g} up to arbitrary but finite orders of knowledge. The ideal situation with complete information assumes that it is common knowledge among players that the game played is \mathbf{g} . Intuitively speaking, this says that everyone knows that the game is \mathbf{g} (the game is mutual knowledge), everyone knows that everyone knows that the game is \mathbf{g} (the game is mutual knowledge at order two), and so on. In our setting, in contrast, the analyst does not know the entire hierarchy of knowledge. Indeed, he is confident in his model up to a certain extent so that he believes that with a high probability players mutually know that the real game is \mathbf{g} up to a finite level (possibly very large). To be specific, an incomplete information game is an (ε, N) -perturbation of \mathbf{g} if the action sets are same as those of \mathbf{g} and each player's prior assigns probability at least $1 - \varepsilon$ on the event that the players know at order N that the payoffs are given by \mathbf{g} . In (ε, N) -perturbations of \mathbf{g} with small ε , ex ante payoffs (computed according to each player's prior) are close to payoffs in \mathbf{g} . An equilibrium of \mathbf{g} is *robust to incomplete information under non-common priors* in \mathbf{g} if there

¹ A posteriori equilibrium is a refinement of subjective correlated equilibrium introduced and studied in Aumann [1] and Brandenburger and Dekel [4].

² This "ex ante" perspective is the one adopted in the purification literature (Harsanyi [9]). A standard interpretation of correlated equilibrium also follows this perspective (see, e.g., Aumann [2]).

exist $\varepsilon > 0$ and $N \geq 0$ such that every (ε, N) -perturbation of \mathbf{g} has a Bayesian Nash equilibrium³ under which the ex ante probability that each player assigns to any action profile is approximately given by the equilibrium of \mathbf{g} . This guarantees that under this Bayesian Nash equilibrium, the ex ante (subjective) payoff of each player is approximately given by the equilibrium payoff in \mathbf{g} .⁴

Our *first main result* shows that for generic games, an action distribution is robust under non-common priors in \mathbf{g} if and only if the game is dominance solvable and the distribution assigns probability one on the unique rationalizable action profile of \mathbf{g} . Its sufficiency follows from the assumption that in incomplete information perturbations close to \mathbf{g} , the real game is mutually known to be \mathbf{g} up to high enough order (at least the number of the dominance iteration rounds needed to reach the single action profile in \mathbf{g}). To show the necessity, which is the main part of this paper, we obtain a contagion result for a posteriori equilibria: for any a posteriori equilibrium of a generic game and for any $\varepsilon > 0$ and $N \geq 0$, we construct a dominance solvable (ε, N) -perturbation whose unique rationalizable strategy profile generates an action distribution that can be arbitrarily close to this a posteriori equilibrium.⁵ From the result by Brandenburger and Dekel [4], we know that if more than one action profile is rationalizable, then there are several a posteriori equilibria. Hence, if more than one actions survive iterative elimination of actions that are never best responses, then our result shows that no action profile is robust.

Brandenburger and Dekel [4] show that for any a posteriori equilibrium of a given complete information game, one can add payoff-irrelevant types to have an incomplete information game with non-common priors whose Bayesian Nash equilibrium generates the distribution of the a posteriori equilibrium. In contrast, our contagion result used for our necessity result shows that (in generic games) when we allow for payoff-relevant (“crazy”) types that have vanishingly small prior probability, the Bayesian Nash equilibrium that generates this distribution can indeed be the unique rationalizable strategy profile of a dominance solvable incomplete information game. We note that it is crucial for our result as well as for the result of Brandenburger and Dekel [4] to drop the CPA.

Our *second main result* concerns set-valued robustness. Since many games have more than one rationalizable outcomes and hence have no robust equilibrium under non-common priors, it is natural to ask if a set of action distributions is robust. Indeed, Kohlberg and Mertens [11] propose making set of equilibria the object of a theory of equilibrium refinements. Following their program as well as Morris and Ui’s [17], we also investigate the robustness of set of equilibria. A set of equilibria of a complete information game is *robust to incomplete information under non-common priors* if there exist $\varepsilon > 0$ and $N \geq 0$ such that any (ε, N) -perturbation has a Bayesian Nash equilibrium whose behavior can be approximated by some action distribution in this set.⁶ If a robust set is a singleton, then its element is a robust equilibrium in the previous sense. A set of action distributions is called a smallest robust set if it is robust and is contained in any robust set. We show that for generic games, a smallest robust set exists and coincides with the set of a posteriori equilibria. When the smallest robust set is a singleton, the condition reduces to

³ Our results would remain unchanged when the solution concept considered is given by any non-empty refinement of interim correlated rationalizability.

⁴ We choose a formulation of our robustness test in terms of (subjective) equilibrium action distributions rather than in terms of equilibrium payoffs for comparison with the previous literature (in particular with Kajii and Morris [10]).

⁵ The construction is more involved but analogous in essence to Rubinstein’s [20] email game.

⁶ As noted previously for the point-valued test, this robustness test can be written in terms of sets of equilibrium payoffs rather than sets of equilibrium action distributions.

the uniqueness of a posteriori equilibrium, so that the result on set-valued robustness covers our result on point-valued robustness.

Kajii and Morris [10] introduce the notion of robustness of equilibria to incomplete information under common prior. They consider incomplete information perturbations of a given complete information where the players share a common prior. They show in particular that a p -dominant equilibrium⁷ with p sufficiently small is robust to incomplete information under common prior. Following Kajii and Morris [10], papers by Ui [22], Morris and Ui [17], and Oyama and Tercieux [19] provide sufficient conditions for a Nash equilibrium to be robust to incomplete information under common prior. Our result shows that when we relax the CPA, none of the existing sufficient conditions implies robustness under non-common priors.

Weinstein and Yildiz [24] consider a notion of *interim robustness*.⁸ A Nash equilibrium a^* is interim robust in \mathbf{g} if for some $N \geq 0$ and for any incomplete information game with (or without) common prior where the action sets are same as those of \mathbf{g} , there exists a Bayesian Nash equilibrium, say σ , such that in any state of the world at which it is mutually known up to order N that \mathbf{g} is the true game, a^* is played under σ . They show that for generic games, a Nash equilibrium is interim robust in \mathbf{g} if and only if it is the unique rationalizable action profile of \mathbf{g} . Contrary to that for our robustness concept, this characterization remains unchanged even if we restrict our attention to incomplete information games with common prior. This result follows from a result of Lipman [12,13], which says that given any partition model with non-common priors (under certain conditions) and any state of the world in the model, for any finite N one can construct a partition model with a common prior such that there is a state in that model at which all the same facts about the world as well as all the same statements about beliefs and knowledge of order less than N are true. Thus, in an interim context where the analyst has to make a prediction given interim beliefs of the players, imposing the CPA does not alter the set of robust predictions. On the other hand, we conclude that in an ex ante context in which the analyst has no information about the players' interim beliefs and is interested in the ex ante behavior so that he may need to know the average behavior over the state space, the CPA has a real bite and allowing for models with heterogeneous priors has important strategic consequences.

To prove their main result, Weinstein and Yildiz [24] show that for any complete information type in the universal type space (see Mertens and Zamir [14] and Brandenburger and Dekel [5])⁹ and any rationalizable action profile a^* of this game, there exist a dominance solvable incomplete information game and a sequence of types drawn from this game such that (1) these types are arbitrarily close to the complete information type (i.e., this sequence converges to it with respect to the product topology in the universal type space) and (2) each type of the sequence plays a^* .¹⁰ Roughly speaking, the former condition requires that changes of *interim* beliefs be small. To establish our results, on the contrary, we construct a dominance solvable incomplete information game such that (1') changes of ex ante beliefs are arbitrarily small and (2') the profile of ex ante subjective payoffs of the unique rationalizable strategy profile is arbitrarily close to the profile of

⁷ See Morris, Rob, and Shin [16] and Kajii and Morris [10].

⁸ See also their working paper version, Weinstein and Yildiz [23].

⁹ A complete information type is a (degenerate) type in the universal type space where it is common knowledge that payoffs are given by the complete information game. Note that Weinstein and Yildiz [24] do not necessarily restrict their attention to complete information types.

¹⁰ One can construct such a dominant solvable incomplete information so that it satisfies the CPA. As shown in Oyama and Tercieux [18], however, in this case it need not be an (ε, N) -perturbation for ε small and N large.

expected payoffs. Hence, we may say that the same type of statement as that by Weinstein and Yildiz [24] is obtained by our *ex ante* approach, provided that the CPA is dropped.

The point behind our contagion argument used in the proofs is that, under non-common priors, a small (*ex ante*) probability event can have a larger impact on higher order (interim) beliefs than under common prior. The “critical path result” of Kajii and Morris [10, Proposition 4.2] shows that, under common prior, small changes in prior beliefs impose some restrictions on interim beliefs. This implies that the impact of a small probability event is not large enough in the sense that, in some games and for some strict Nash equilibrium, a small amount of payoff uncertainty cannot induce this equilibrium to be played everywhere on the state space (i.e., it is not contagious). For instance, in 2×2 coordination games, the risk-dominated equilibrium cannot be contagious, and indeed the risk-dominant one is robust under common prior. In a companion paper (Oyama and Tercieux [18]), we demonstrate, in contrast, that with non-common priors, any strict Nash equilibrium can be contagious. In that paper, for two-player incomplete information games with non-common priors, we study the strategic impact of an event using the notion of belief potential (Morris, Rob, and Shin [16]). We find a measure of discrepancy from the CPA such that the belief potential of a small probability event has an upper bound that is an increasing function of this measure. Indeed, in order for any strict Nash equilibrium to be contagious, this measure of discrepancy has to be large. In the present paper, we extend this observation and show that for any a posteriori equilibrium of any complete information game to be induced by a unique rationalizable strategy of some nearby dominance solvable incomplete information perturbations, this measure of discrepancy from the CPA in these perturbations need to be arbitrarily large, due to the fact that a bound on the measure would impose restrictions on interim beliefs.¹¹

The remainder of the paper is organized as follows. Section 2 presents our notions of nearby incomplete information games and robustness. Section 3 states and proves our characterization of robust equilibria, while robust sets are studied in Section 4. Section 5 discusses alternative notions of robustness, in particular as those studied by Kajii and Morris [10] and Weinstein and Yildiz [24].

2. Framework

This section collects the notations and definitions that will be used in the subsequent analysis. The reader at first reading may proceed directly to the simple example in Section 3.1 which illustrates our main results and when necessary refer to the relevant notation/definition in this section.

2.1. Complete information games

A complete information game consists of the set of players, $\mathcal{I} = \{1, 2, \dots, I\}$, the finite set of actions, A_i , for each player $i \in \mathcal{I}$, and the payoff function, $g_i : A \rightarrow \mathbb{R}$, for each player $i \in \mathcal{I}$. Throughout our analysis, we fix a complete information game, simply denoted by $\mathbf{g} = (g_i)_{i \in \mathcal{I}}$.

For any at most countable set S , we denote by $\Delta(S)$ the set of all probability measures on S . We call elements in $\Delta(A)$ action distributions. For $a \in A$, we write $[a]$ for the element in $\Delta(A)$ that assigns weight one to a . For $\xi \in \Delta(A)$ and $a_i \in A_i$, we denote $\xi(a_i) = \sum_{a_{-i} \in A_{-i}} \xi(a_i, a_{-i})$,

¹¹ If we first fix a given a posteriori equilibrium of a complete information game, then we can find a finite upper bound of the measure for the incomplete information perturbations to generate the a posteriori equilibrium. Note that the same comment applies to the result of Brandenburger and Dekel [4].

and if $\xi(a_i) > 0$, we define $\xi(\cdot|a_i) \in \Delta(A_{-i})$ by $\xi(a_{-i}|a_i) = \xi(a_i, a_{-i})/\xi(a_i)$. We measure the distance between any two elements $\xi, \xi' \in \Delta(A)$ by $|\xi - \xi'| = \max_{a \in A} |\xi(a) - \xi'(a)|$. For $\delta > 0$, we denote $V_\delta(\xi) = \{\xi' \in \Delta(A) \mid |\xi' - \xi| < \delta\}$ for $\xi \in \Delta(A)$ and $V_\delta(\mathcal{E}) = \{\xi' \in \Delta(A) \mid |\xi' - \xi| < \delta \text{ for some } \xi \in \mathcal{E}\}$ for $\mathcal{E} \subset \Delta(A)$. With abuse of notation, we also write $|\mu - \mu'| = \max_{i \in \mathcal{I}} |\mu_i - \mu'_i|$ for $\mu = (\mu_i)_{i \in \mathcal{I}}, \mu' = (\mu'_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ and $|\pi_i - \pi'_i| = \max_{a_{-i} \in A_{-i}} |\pi_i(a_{-i}) - \pi'_i(a_{-i})|$ for $\pi_i, \pi'_i \in \Delta(A_{-i})$.

Given \mathbf{g} , let $br_i : \Delta(A_{-i}) \rightarrow A_i$ be the best response correspondence in pure actions for player $i \in \mathcal{I}$:

$$br_i(\pi_i) = \arg \max_{a_i \in A_i} g_i(a_i, \pi_i)$$

for $\pi_i \in \Delta(A_{-i})$, where $g_i(a_i, \cdot)$ is extended to $\Delta(A_{-i})$ in the usual way. We define *correlated rationalizability* (e.g., Brandenburger and Dekel [4]). For each $i \in \mathcal{I}$, set $S_i^0[\mathbf{g}] = A_i$. Then, for $k = 1, 2, \dots$, define $S_i^k[\mathbf{g}]$ recursively by

$$S_i^k[\mathbf{g}] = \{a_i \in A_i \mid a_i \in br_i(\pi_i) \text{ for some } \pi_i \in \Delta(S_{-i}^{k-1}[\mathbf{g}])\},$$

where we denote $S_{-i}^{k-1}[\mathbf{g}] = \prod_{j \neq i} S_j^{k-1}[\mathbf{g}]$. The set of all rationalizable actions for player $i \in \mathcal{I}$ is $S_i^\infty[\mathbf{g}] = \bigcap_{k=0}^\infty S_i^k[\mathbf{g}]$. We denote $S^\infty[\mathbf{g}] = \prod_{i \in \mathcal{I}} S_i^\infty[\mathbf{g}]$ as well as $S^k[\mathbf{g}] = \prod_{i \in \mathcal{I}} S_i^k[\mathbf{g}]$ for $k \geq 1$. We say that \mathbf{g} is dominance solvable if $S^\infty[\mathbf{g}]$ is a singleton set. We also define the set of actions that survive iterative elimination of actions that are never *strict* best response.¹² For each $i \in \mathcal{I}$, set $W_i^0[\mathbf{g}] = A_i$. Then, for $k = 1, 2, \dots$, define $W_i^k[\mathbf{g}]$ recursively by

$$W_i^k[\mathbf{g}] = \{a_i \in A_i \mid a_i \in br_i(\pi_i) \text{ for some } \pi_i \in \Delta(W_{-i}^{k-1}[\mathbf{g}])\},$$

where we denote $W_{-i}^{k-1}[\mathbf{g}] = \prod_{j \neq i} W_j^{k-1}[\mathbf{g}]$. Finally, let $W_i^\infty[\mathbf{g}] = \bigcap_{k=0}^\infty W_i^k[\mathbf{g}]$. We denote $W^\infty[\mathbf{g}] = \prod_{i \in \mathcal{I}} W_i^\infty[\mathbf{g}]$ as well as $W^k[\mathbf{g}] = \prod_{i \in \mathcal{I}} W_i^k[\mathbf{g}]$ for $k \geq 1$. Note that $S^\infty[\mathbf{g}]$ is always non-empty, while $W^\infty[\mathbf{g}]$ may be empty (consider, e.g., games where the payoff functions are constant). But the set of normal form games \mathbf{g} for which these sets coincide, $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$, is generic in the set of finite games. Our main result will be proved for this generic class of games.

We also use the following notions due to Aumann [1] and Brandenburger and Dekel [4]. First, let us review the definition of subjective correlated equilibrium.

Definition 2.1. A profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is a *subjective correlated equilibrium* of \mathbf{g} if for all $i \in \mathcal{I}$ and all $a_i \in A_i$,

$$\mu_i(a_i) > 0 \implies a_i \in br_i(\mu_i(\cdot|a_i)).$$

As in Brandenburger and Dekel [4], our analysis employs the refinement of subjective correlated equilibrium called a posteriori equilibrium.

Definition 2.2. For non-negative integer N , a profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is an *N-subjective correlated equilibrium* of \mathbf{g} if it is a subjective correlated equilibrium of \mathbf{g} and $\mu_i(S^N[\mathbf{g}]) = 1$ for all $i \in \mathcal{I}$.

A profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is an *a posteriori equilibrium* of \mathbf{g} if it is a subjective correlated equilibrium of \mathbf{g} and $\mu_i(S^\infty[\mathbf{g}]) = 1$ for all $i \in \mathcal{I}$.

¹² To the best of our knowledge, this notion has been first defined by Weinstein and Yildiz [23].

Denote by $\mathcal{E}^N[\mathbf{g}]$ the set of N -subjective correlated equilibria of \mathbf{g} and by $\mathcal{E}[\mathbf{g}]$ the set of a posteriori equilibria of \mathbf{g} . Observe that $\mathcal{E}^N[\mathbf{g}]$ and $\mathcal{E}[\mathbf{g}]$ are product sets ($\mathcal{E}^N[\mathbf{g}] = \prod_{i \in \mathcal{I}} \mathcal{E}_i^N[\mathbf{g}]$ with each $\mathcal{E}_i^N[\mathbf{g}] \subset \Delta(A_i)$) and closed sets in $(\Delta(A))^I$.

We introduce further refinements of a posteriori equilibrium.

Definition 2.3. (a) For non-negative integer N , a profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is an *undominated N -subjective correlated equilibrium* of \mathbf{g} if it is an N -subjective correlated equilibrium such that $\mu_i(W^N[\mathbf{g}]) = 1$ for all $i \in \mathcal{I}$.

(b) $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is an *undominated a posteriori equilibrium* of \mathbf{g} if it is a posteriori equilibrium such that $\mu_i(W^\infty[\mathbf{g}]) = 1$ for all $i \in \mathcal{I}$.

(c) $(\mu_i)_{i \in \mathcal{I}}$ is a *strict a posteriori equilibrium* if it is an undominated a posteriori equilibrium such that for all $i \in \mathcal{I}$ and all $a_i \in A_i$,

$$\mu_i(a_i) > 0 \quad \Rightarrow \quad \{a_i\} = br_i(\mu_i(\cdot|a_i)).$$

We denote by $\mathcal{E}^u[\mathbf{g}]$ the set of undominated a posteriori equilibrium of \mathbf{g} , which is again a product set. If $W^\infty[\mathbf{g}] \neq \emptyset$, then $\mathcal{E}^u[\mathbf{g}] \neq \emptyset$. For generic games where $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$, we have $\mathcal{E}[\mathbf{g}] = \mathcal{E}^u[\mathbf{g}]$.

2.2. Incomplete information perturbations

We would like to consider incomplete information games that are close to complete information game \mathbf{g} .

An incomplete information game \mathcal{U} consists of the set of players, \mathcal{I} ; their action sets, A_1, \dots, A_I ; a countable state space, Ω ; ¹³ a prior probability measure on the state space, P_i , for each player $i \in \mathcal{I}$; a partition of the state space, Q_i , for each $i \in \mathcal{I}$; and a bounded payoff function, $u_i: A \times \Omega \rightarrow \mathbb{R}$, for each $i \in \mathcal{I}$. The incomplete information game $\mathcal{U} = (\Omega, (P_i)_{i \in \mathcal{I}}, (Q_i)_{i \in \mathcal{I}}, (u_i)_{i \in \mathcal{I}})$ is said to be an *incomplete information perturbation* of \mathbf{g} (recall that \mathcal{U} shares the player set and the action sets with the complete information game $(\mathcal{I}, (A_i)_{i \in \mathcal{I}}, \mathbf{g})$). In order to incorporate an outsider analyst’s viewpoint, we also explicitly consider his prior belief. A pair (\mathcal{U}, P_0) of an incomplete information game as described above, \mathcal{U} , and a probability measure on the state space for the analyst, P_0 , is called an *incomplete information elaboration* of \mathbf{g} .

Subsets of Ω are called events. For each $i \in \mathcal{I}$, we write \mathcal{F}_i for the sigma algebra generated by Q_i , i.e., the set of unions of events in Q_i together with the empty set. We say that an event $E \subset \Omega$ is *simple* if $E = \bigcap_{i \in \mathcal{I}} E_i$ where each $E_i \in \mathcal{F}_i$. We write $Q_i(\omega)$ for the element of Q_i containing ω . We assume that $P_i(Q_i(\omega)) > 0$ for all $i \in \mathcal{I}$ and $\omega \in \Omega$. Under this assumption, the conditional probability of any ω' given $Q_i(\omega)$, $P_i(\omega'|Q_i(\omega))$, is well defined by $P_i(\omega'|Q_i(\omega)) = P_i(\omega')/P_i(Q_i(\omega))$ whenever $\omega' \in Q_i(\omega)$.

We sometimes impose restrictions of possible priors as in Lipman [12,13].

Definition 2.4. $\{P_i\}_{i \in \mathcal{I}}$ is said to have *common support* if $\text{supp}(P_i) = \text{supp}(P_j)$ for all $i, j \in \mathcal{I}$.

¹³ Assuming countability allows us to avoid measurability issues, in particular regarding existence of Bayesian Nash equilibria (cf. Simon [21]).

Definition 2.5. For $\mathcal{L} \geq 1$, $\{P_i\}_{i \in \mathcal{I}}$ is said to satisfy \mathcal{L} -tail consistency if

$$\frac{1}{\mathcal{L}} \leq \frac{P_i(\omega)}{P_j(\omega)} \leq \mathcal{L}$$

for all $i, j \in \mathcal{I}$ and all $\omega \in \Omega$ with $P_j(\omega) > 0$.

By a slight abuse of language, we say that an incomplete information game \mathcal{U} satisfies common support or \mathcal{L} -tail consistency.

Definition 2.6. \mathcal{U} is said to satisfy the CPA if $P_i = P_j$ for all $i, j \in \mathcal{I}$.

We now define the solution concepts we use for incomplete information games. Given an incomplete information game \mathcal{U} , a (behavioral) strategy for player i is a \mathcal{Q}_i -measurable function $\sigma_i : \Omega \rightarrow \Delta(A_i)$. Denote by Σ_i the set of player i 's strategies, and let $\Sigma = \prod_{i \in \mathcal{I}} \Sigma_i$ and $\Sigma_{-i} = \prod_{j \neq i} \Sigma_j$. We write $\sigma_i(a_i | \omega)$ for the probability that action $a_i \in A_i$ is chosen at $\omega \in \Omega$ under $\sigma_i \in \Sigma_i$, and denote $\sigma_{-i}(a_{-i} | \omega) = \prod_{j \neq i} \sigma_j(a_j | \omega)$ for $\sigma_{-i} \in \Sigma_{-i}$ and $a_{-i} \in A_{-i}$ as well as $\sigma(a | \omega) = \prod_{i \in \mathcal{I}} \sigma_i(a_i | \omega)$ for $\sigma \in \Sigma$ and $a \in A$. For $\sigma \in \Sigma$ and $P_i \in \Delta(\Omega)$, we write $\sigma_{P_i} \in \Delta(A)$ for the induced action distribution with respect to P_i , i.e., $\sigma_{P_i}(a) = \sum_{\omega \in \Omega} P_i(\omega) \sigma(a | \omega)$ for $a \in A$.

For player $i \in \mathcal{I}$ and action $a_i \in A_i$, we write the expected payoff against a conjecture $v_i \in \Delta(\Omega \times A_{-i})$ as

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

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

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

For $i \in \mathcal{I}$ and $\sigma_{-i} \in \Sigma_{-i}$, we denote by $\sigma_{-i}^{Q_i} \in \Delta(\Omega \times A_{-i})$ the induced conjecture at $Q_i \in \mathcal{Q}_i$:

$$\sigma_{-i}^{Q_i}(\omega, a_{-i}) = P_i(\omega | Q_i) \sigma_{-i}(a_{-i} | \omega).$$

Note that $\text{marg}_{\Omega} \sigma_{-i}^{Q_i} = P_i(\cdot | Q_i)$.

Definition 2.7. A strategy profile σ is a Bayesian Nash equilibrium of \mathcal{U} if for all $i \in \mathcal{I}$,

$$\sigma_i(a_i | \omega) > 0 \Rightarrow a_i \in BR_i(\sigma_{-i}^{Q_i(\omega)})$$

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

We also define interim correlated rationalizability. For each $i \in \mathcal{I}$, let $R_i^0[Q_i] = A_i$ for all $Q_i \in \mathcal{Q}_i$. Then, for each $i \in \mathcal{I}$, and for $Q_i \in \mathcal{Q}_i$ and for $k = 1, 2, \dots$, define $R_i^k[Q_i]$ recursively by

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

where we denote $R_{-i}^{k-1}[\omega] = \prod_{j \neq i} R_j^{k-1}[Q_j(\omega)]$. Let $R_i^{\infty}[Q_i] = \bigcap_{k=0}^{\infty} R_i^k[Q_i]$.

Definition 2.8. A strategy $\sigma_i \in \Sigma_i$ is a *rationalizable strategy* of player i in \mathcal{U} if

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

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

This definition states that player i 's strategy is rationalizable if it is in the convex hull of $R_i^\infty[Q_i]$ for all $Q_i \in \mathcal{Q}_i$. While this is weaker than the standard definitions (Battigalli and Siniscalchi [3], Dekel, Fudenberg, and Morris [7]), our results would remain valid under any stronger notion.

Note that a Bayesian Nash equilibrium is a rationalizable strategy profile. We say that incomplete information game \mathcal{U} is *dominance solvable* if $R_i^\infty[Q_i]$ is a singleton set for all $i \in \mathcal{I}$ and $Q_i \in \mathcal{Q}_i$.

We then restate the standard definition of knowledge operator which is used in defining our main concept of robustness. Fix the information system part of an incomplete information game, $(\Omega, (P_i)_{i \in \mathcal{I}}, (Q_i)_{i \in \mathcal{I}})$. The knowledge operator for player i , $K_i : 2^\Omega \rightarrow 2^\Omega$, is defined by

$$K_i(E) = \{\omega \in \Omega \mid Q_i(\omega) \subset E\}.$$

That is, $K_i(E)$ is the set of states where player i knows that event E is true. Let $K_*(E) = \bigcap_{i \in \mathcal{I}} K_i(E)$ be the set of states where it is mutual knowledge that event E is true, i.e., where every player knows that event E is true. At a state ω , an event E is said to be mutual knowledge at order N if $\omega \in \bigcap_{n=1}^N [K_*]^n(E)$, where $[K_*]^n(\cdot)$ is defined recursively by $[K_*]^n(E) = K_*([K_*]^{n-1}(E))$. Finally, at state ω , an event E is said to be common knowledge if $\omega \in \bigcap_{n=1}^\infty [K_*]^n(E)$.

2.3. Robustness

In this subsection, we introduce our concept of robustness of equilibria to incomplete information under non-common priors. Given an incomplete information perturbation \mathcal{U} of \mathbf{g} , let $\Omega_{\mathbf{g}}^i$ be the set of states where the payoffs of player $i \in \mathcal{I}$ are given by g_i and he knows his payoff:

$$\Omega_{\mathbf{g}}^i = \{\omega \in \Omega \mid u_i(\cdot, \omega') = g_i(\cdot) \text{ for all } \omega' \in Q_i(\omega)\}.$$

Denote $\Omega_{\mathbf{g}} = \bigcap_{i \in \mathcal{I}} \Omega_{\mathbf{g}}^i$.

Definition 2.9. An incomplete information game \mathcal{U} is an (ε, N) -perturbation of \mathbf{g} if $P_i(\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}})) \geq 1 - \varepsilon$ for all $i \in \mathcal{I}$.

Observe that since $K_*(E) \subset E$ for any event E , we have that $[K_*]^N(\Omega_{\mathbf{g}})$ is decreasing in N and thus $\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}}) = [K_*]^N(\Omega_{\mathbf{g}}) \subset \Omega_{\mathbf{g}}$. Note also that if $\varepsilon' \leq \varepsilon$ and $N' \geq N$, then an (ε', N') -perturbation is an (ε, N) -perturbation.

We define our robustness concept for action distribution profiles, where each action distribution in a profile is generated by the corresponding player's prior.

Definition 2.10. A profile of action distributions $\mu = (\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is *N-robust to incomplete information under non-common priors*, or simply, *N-robust*, if for all $\delta > 0$, there exists $\varepsilon > 0$ such that any (ε, N) -perturbation of \mathbf{g} , \mathcal{U} , has a Bayesian Nash equilibrium σ such that $|\mu_i - \sigma_{P_i}| \leq \delta$ for all $i \in \mathcal{I}$.

A profile of action distributions $\mu \in (\Delta(A))^I$ is *robust to incomplete information under non-common priors*, or simply, *robust*, if there exists $N \geq 0$ such that μ is N -robust.

Observe that if μ is N -robust, then it is N' -robust for all $N' \geq N$.

This concept is most relevant in the following situation. Imagine an analyst who considers an equilibrium of a particular complete information game. He is interested in the profile of equilibrium payoffs of this game (e.g., because of some welfare criterion he cares about). This analyst has a lack of confidence in his model. Hence, he would like to check whether the equilibrium payoff profile he considers is not sensitive to the assumption of common knowledge of payoffs. If the profile is robust in this sense, then ex ante (subjective) expected payoffs of each player in nearby incomplete information games will not change significantly from the complete information game situation. We do not define directly robustness for (subjective) ex ante expected payoff profiles since the ex ante payoff of each player i is immediately obtained from the action distribution μ_i by $\sum_{a \in A} \mu_i(a)g_i(a)$ (whenever ε is vanishingly small).

We also propose another robustness concept, which incorporates the analyst's possible priors.

Definition 2.11. A pair (\mathcal{U}, P_0) of an incomplete information game \mathcal{U} and a prior distribution P_0 on Ω is an (ε, N) -elaboration of \mathbf{g} if \mathcal{U} is an (ε, N) -perturbation of \mathbf{g} and $P_0(\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}})) \geq 1 - \varepsilon$.

Definition 2.12. An action distribution $\xi \in \Delta(A)$ is N -robust to incomplete information under non-common priors, or simply, N -robust, if for all $\delta > 0$, there exists $\varepsilon > 0$ such that for any (ε, N) -elaboration of \mathbf{g} , (\mathcal{U}, P_0) , \mathcal{U} has a Bayesian Nash equilibrium σ such that $|\xi - \sigma_{P_0}| \leq \delta$.

An action distribution $\xi \in \Delta(A)$ is robust to incomplete information under non-common priors, or simply, robust, if there exists $N \geq 0$ such that ξ is N -robust.

This concept is relevant in a situation where the analyst is interested in ex ante expected behavior of the players, but the expectation is taken with respect to his own prior distribution, which is not necessarily equal to the priors the players may have.¹⁴

We will show that in generic games, a profile of action distributions $(\mu_i)_{i \in \mathcal{I}}$ (an action distribution ξ , resp.) of \mathbf{g} is robust to incomplete information under non-common priors if and only if $(\mu_i)_{i \in \mathcal{I}}$ (ξ , resp.) consists of the unique rationalizable action profile of \mathbf{g} .¹⁵ We want to underline that our main result will stay unchanged if we modify these robustness notions in various directions. In particular, since the non-trivial result is the “only if” part, we want to show that we can weaken this concept in many respects keeping our characterization.

Remark 2.1. In the definition of robustness, we use the notion of Bayesian Nash equilibrium to be consistent with that by Kajii and Morris [10] except for dropping the CPA. One might find it questionable to use Bayesian Nash equilibrium when players do not share a common prior (Dekel, Fudenberg, and Levine [6]). However, our results would be unchanged if we changed the solution concept to the weaker concept of interim correlated rationalizability.¹⁶ Indeed, all the lemmata that are used to prove our main result are stated with rationalizable strategies.

¹⁴ Kajii and Morris [10] offer a motivating story of this type for their robustness concept under common prior, where the analyst shares a common prior with the players.

¹⁵ The two robustness concepts a priori have no logical link and indeed are distinct if we consider their set-valued extensions, as we will see in Section 4. In games that have a unique rationalizable action profile, both versions of robust sets collapse to a singleton, and therefore the two point-valued concepts share the same characterization, showing their equivalence.

¹⁶ Our results would be actually unchanged if we use any non-empty refinement of interim correlated rationalizability.

Remark 2.2. In Definition 2.9, we could have defined an (ε, N) -perturbation to be an incomplete information perturbation such that $P_i(\bigcap_{n=1}^N [B_*^{1-\varepsilon}]^n(\Omega_{\mathbf{g}})) \geq 1 - \varepsilon$ for all $i \in \mathcal{I}$, where the mutual knowledge operator K_* is replaced with the mutual $(1 - \varepsilon)$ -belief operator $B_*^{1-\varepsilon}$ (see Section 5.2 for its formal definition). The formulation of Kajii and Morris [10] would then become equivalent to ours *with* common priors (due to Lemma B of Kajii and Morris [10]), while this equivalence breaks down under non-common priors. Our results would remain unchanged with this formulation.

If, in addition, we dropped the measurability of the set $\Omega_{\mathbf{g}}$ with respect to the information partitions and defined it by $\Omega_{\mathbf{g}} = \{\omega \in \Omega \mid u_i(\cdot, \omega) = g_i(\cdot) \text{ for all } i \in \mathcal{I}\}$, then our characterizations of robustness, which apply to a generic class of games, could extend to the set of *all* finite games (with the assumption that payoffs are uniformly bounded over all incomplete information perturbations).

Remark 2.3. Another way to weaken our test would be to restrict our attention to specific (ε, N) -elaborations of \mathbf{g} . Natural restrictions would be to require that each (ε, N) -elaboration \mathcal{U} satisfies common support and $\mathcal{L}^{\mathcal{U}}$ -tail consistency for some $\mathcal{L}^{\mathcal{U}} > 0$. Here again, our results would stay unchanged.

2.4. Preliminary results

We conclude this section with two preliminary observations.

Lemma 2.1. *For any $N \geq 1$, any rationalizable strategy profile σ of any incomplete information perturbation of \mathbf{g} satisfies*

$$\sum_{a_i \in S_i^N[\mathbf{g}]} \sigma_i(a_i | \omega) = 1 \quad \text{for all } i \in \mathcal{I} \text{ and } \omega \in [K_*]^{N-1}(\Omega_{\mathbf{g}}).$$

Proof. We prove by induction that for all $n = 1, \dots, N$,

$$R_i^n[Q_i(\omega)] \subset S_i^n[\mathbf{g}] \quad \text{for all } i \in \mathcal{I} \text{ and all } \omega \in [K_*]^{n-1}(\Omega_{\mathbf{g}}), \tag{*n}$$

where $[K_*]^0(\Omega_{\mathbf{g}}) = \Omega_{\mathbf{g}}$. First, since for all $i \in \mathcal{I}$, $u_i(\cdot, \omega) = g_i(\cdot)$ for all $\omega \in \Omega_{\mathbf{g}}$, $(*_1)$ is true.

Assume $(*_n)$. Consider any $i \in \mathcal{I}$ and any $\omega \in [K_*]^n(\Omega_{\mathbf{g}})$, and take any conjecture of i , $v_i \in \Delta(\Omega \times A_{-i})$, such that $v_i(\{(\omega', a_{-i}) \mid a_{-i} \in R_{-i}^n[\omega']\}) = 1$ and $\text{marg}_{\Omega} v_i = P_i(\cdot | Q_i(\omega))$. Note that $[K_*]^{n-1}(\Omega_{\mathbf{g}}) \subset \Omega_{\mathbf{g}}$ since $\Omega_{\mathbf{g}}$ is a simple event. By the assumption of $(*_n)$, for all $\omega' \in [K_*]^{n-1}(\Omega_{\mathbf{g}})$, $R_{-i}^n[\omega'] \subset S_{-i}^n[\mathbf{g}]$. Since $\omega \in K_i([K_*]^{n-1}(\Omega_{\mathbf{g}}))$, we have $(\text{marg}_{A_{-i}} v_i)(S_{-i}^n[\mathbf{g}]) = 1$. Since $U_i(a_i, v_i) = g_i(a_i, \text{marg}_{A_{-i}} v_i)$, it follows that $BR_i(v_i) = br_i(\text{marg}_{A_{-i}} v_i) \subset S_i^{n+1}[\mathbf{g}]$, implying $(*_{n+1})$. \square

This lemma has the following implication.

Lemma 2.2. *Fix $N \geq 0$. Suppose that $\varepsilon^k \rightarrow 0$ as $k \rightarrow \infty$, that each \mathcal{U}^k is an (ε^k, N) -perturbation of \mathbf{g} , and that each $(\sigma_{P_i}^k)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is a profile of equilibrium action distributions of \mathcal{U}^k . Then, a subsequence of $\{(\sigma_{P_i}^k)_{i \in \mathcal{I}}\}_k$ converges to some $(N + 1)$ -subjective correlated equilibrium of \mathbf{g} .*

Proof. It is simple to show that a subsequence of $\{(\sigma_{p_i}^k)_{i \in \mathcal{I}}\}_k$ converges to some subjective correlated equilibrium of \mathbf{g} (see, e.g., Kajii and Morris [10, Lemma 3.4 and Corollary 3.5]). Lemma 2.1 completes the proof. \square

3. Point-valued robustness

In this section, we present and prove our first main result. For $a \in A$, we denote by $([a])^I$ the profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ such that $\mu_i = [a]$ for all $i \in \mathcal{I}$.

Theorem 3.1. *Suppose that $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$. Then, \mathbf{g} has a robust equilibrium if and only if \mathbf{g} is dominance solvable.*

When \mathbf{g} is dominance solvable and a^* is the unique rationalizable action profile (i.e., $S^\infty[\mathbf{g}] = \{a^*\}$), $([a^*])^I$ ($[a^*]$, resp.) is the robust action distribution profile (robust action distribution, resp.) in \mathbf{g} .

Dominance solvability is obviously a very strong condition. For instance, the theorem does not guarantee that a unique Nash equilibrium is robust. Indeed, as proved by Kajii and Morris [10], there exists an open set of games with a unique Nash equilibrium that is not robust.¹⁷ On the other hand, it is well known that in supermodular games, a unique Nash equilibrium is necessarily a unique rationalizable action profile, and hence a robust equilibrium under non-common priors.

In the first subsection that follows, we discuss a simple two-player two-action example to illustrate our results. We then prove the sufficiency and the necessity parts of Theorem 3.1 respectively in Sections 3.2 and 3.3. While the former follows from Lemma 2.1 in a straightforward way, in proving the latter we will utilize our key lemma on contagion of a posteriori equilibria (Lemma 3.4). This lemma will also be crucial in proving the result on set-valued robustness (Theorem 4.2).

3.1. Example

In this subsection, we illustrate the necessity result in Theorem 3.1 by using a simple example of complete information matching pennies game. In particular, we sketch our key construction of dominant solvable (ε, N) -perturbation provided in the proof of Lemma 3.4.

The game \mathbf{g} is given by

	H_2	T_2
H_1	1, -1	-1, 1
T_1	-1, 1	1, -1

where each player $i \in \mathcal{I} = \{1, 2\}$ plays H_i or T_i so that $A = \{H_1, T_1\} \times \{H_2, T_2\}$. This game has a unique Nash equilibrium which is the unique (objective) correlated equilibrium. By Kajii and Morris [10], we know that this equilibrium is robust under the assumption that players share a common prior in nearby incomplete information perturbations. However, this game is clearly not

¹⁷ While the robustness concept introduced by Kajii and Morris [10] is different from ours, one can easily show that their example goes through if we use our formulation of robustness.

dominance solvable, and hence, according to our Theorem 3.1 we claim that no equilibrium is robust once the CPA is dropped. To see this, it will be sufficient to show that any strict a posteriori equilibrium (recall Definition 2.3) can be played as a *unique* rationalizable strategy in some nearby incomplete information perturbations.

Let us fix any strict a posteriori equilibrium $(\mu_1, \mu_2) \in \Delta(A) \times \Delta(A)$ with full support. Verify that, being strict and having full support, it must satisfy

$$\begin{aligned} 0 < \mu_1(H_1, T_2) < \mu_1(H_1, H_2) < 1, \quad 0 < \mu_1(T_1, H_2) < \mu_1(T_1, T_1) < 1, \\ 0 < \mu_2(H_1, H_2) < \mu_2(T_1, H_2) < 1, \quad 0 < \mu_2(T_1, T_2) < \mu_2(H_1, T_2) < 1, \end{aligned} \tag{3.1}$$

and thus $\mu_1 \neq \mu_2$. Pick any $\varepsilon > 0$ and $N \geq 0$. We construct a dominance solvable (ε, N) -perturbation such that the strict a posteriori equilibrium (μ_1, μ_2) becomes “contagious”: that is, under the unique rationalizable strategy profile, say σ , the ex ante probability that each player $i \in \mathcal{I}$ assigns to any action profile $a \in A$ is given by $\mu_i(a)$, i.e., $P_i(\{\omega \mid \sigma(a|\omega) = 1\}) = \mu_i(a)$.

The construction is as follows.¹⁸ The state space Ω is given by $\Omega = \mathcal{I} \times \mathbb{Z}_+ \times A = \{1, 2\} \times \{0, 1, 2, \dots\} \times \{H_1, T_1\} \times \{H_2, T_2\}$. Denote $\tilde{\varepsilon} = 1 - (1 - \varepsilon)^{1/(N+1)}$. For each player $i \in \mathcal{I}$, define the prior probability distribution $P_i \in \Delta(\Omega)$ by

$$P_i(i, k, a) = \tilde{\varepsilon}(1 - \tilde{\varepsilon})^k \mu_i(a)$$

and

$$P_i(-i, k, a) = 0$$

where $a \in \{H_1, T_1\} \times \{H_2, T_2\}$.¹⁹ Let each information partition \mathcal{Q}_i consist of the events

$$\begin{aligned} E_{iH_i}^0 &= \{(-i, 0, H_i, H_{-i}), (-i, 0, H_i, T_{-i})\}, \\ E_{iT_i}^0 &= \{(-i, 0, T_i, H_{-i}), (-i, 0, T_i, T_{-i})\}, \end{aligned}$$

and all the events of the form

$$\begin{aligned} E_{iH_i}^k &= \{(i, k - 1, H_i, H_{-i}), (i, k - 1, H_i, T_{-i}), (-i, k, H_i, H_{-i}), (-i, k, H_i, T_{-i})\}, \\ E_{iT_i}^k &= \{(i, k - 1, T_i, H_{-i}), (i, k - 1, T_i, T_{-i}), (-i, k, T_i, H_{-i}), (-i, k, T_i, T_{-i})\}, \end{aligned}$$

for each $k \geq 1$. Finally, at any state $\omega \in E_{iH_i}^0$ ($\omega \in E_{iT_i}^0$, resp.), player i 's (ex post) payoffs are given by a game where playing H_i (T_i resp.) is a strictly dominant action, while at any other state, (ex post) payoffs are given by the complete information game \mathbf{g} , and thus $\Omega_{\mathbf{g}} = \{(i, k, a) \in \Omega \mid k \geq 1, i \in \mathcal{I}, a \in A\}$. One can check that $\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}}) = \{(i, k, a) \in \Omega \mid k \geq N + 1, i \in \mathcal{I}, a \in A\}$ and $P_i(\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}})) = (1 - \tilde{\varepsilon})^{N+1} = 1 - \varepsilon$ for each $i \in \mathcal{I}$, so that this incomplete information game is an (ε, N) -perturbation of \mathbf{g} (recall Definition 2.9).

Now let σ be any rationalizable strategy profile of this incomplete information game. We want to show that for all $i \in \mathcal{I}$ and each action a_i , $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^k$ and all $k \geq 0$. Note that this implies that for all $a \in A$, $\sigma(a|(i, k, a)) = 1$ for all $i \in \mathcal{I}$ and all $k \geq 0$. Hence, $P_i(\{\omega \mid \sigma(a|\omega) = 1\}) = P_i(\bigcup_{j \in \mathcal{I}} \bigcup_{k=0}^{\infty} \{(j, k, a)\}) = \sum_{k=0}^{\infty} P_i(i, k, a) = \mu_i(a)$ as claimed.

¹⁸ Contagion of strict (a fortiori, pure strategy) Nash equilibria can be obtained with a simpler construction. See the companion paper Oyama and Tercieux [18] for such a construction in an example of 2×2 coordination games.

¹⁹ For simplicity, we here assume extreme heterogeneity in the prior distributions P_1 and P_2 : to each state, one player assigns strictly positive probability, while the other assigns probability 0. This is not necessary for the result to hold, and one may in fact perturb the priors so that they have common support. See the proof of Lemma 3.4.

First, by construction, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^0$ and all $i \in \mathcal{I}$. Let us then check the claim for $k = 1$. Consider any $\omega \in E_{ia_i}^1$. By construction of the state space and by definition of P_i , we have

$$\begin{aligned} P_i((i, 0, a_i, a_{-i})|Q_i(\omega)) &= \frac{P_i((i, 0, a_i, a_{-i}))}{P_i((i, 0, a_i, H_{-i})) + P_i((i, 0, a_i, T_{-i}))} \\ &= \frac{\tilde{\varepsilon} \times \mu_i(a_i, a_{-i})}{\tilde{\varepsilon} \times \mu_i(a_i, H_{-i}) + \tilde{\varepsilon} \times \mu_i(a_i, T_{-i})} \\ &= \mu_i(a_{-i}|a_i). \end{aligned}$$

Thus, at state $\omega \in E_{ia_i}^1$, for any rationalizable strategy, player i assigns probability $\mu_i(H_{-i}|a_i)$ to his opponent playing H_{-i} and $\mu_i(T_{-i}|a_i)$ to his opponent playing T_{-i} . Therefore, by definition of strict a posteriori equilibrium, we must have $\sigma_i(a_i|\omega) = 1$. In this way, one can indeed show by induction that for all $k \geq 1$, $P_i((i, k, a_i, a_{-i})|Q_i(\omega)) = \mu_i(a_{-i}|a_i)$ for $\omega \in E_{ia_i}^{k+1}$.

Hence, in the (ε, N) -perturbation above, the unique rationalizable strategy profile σ is such that $P_i(\{\omega \mid \sigma(a|\omega) = 1\}) = \mu_i(a)$ for all $i \in \mathcal{I}$ and all $a \in A$.²⁰ Since ε and N have been fixed arbitrarily, this shows that any action distribution profile other than (μ_1, μ_2) is not a robust equilibrium. Clearly, this game has multiple strict a posteriori equilibria (in fact, there are a continuum of distributions satisfying (3.1)), so that the above construction can be done for $(\mu'_1, \mu'_2) \neq (\mu_1, \mu_2)$, which shows in particular that (μ_1, μ_2) too is not robust. Hence, we conclude that this game \mathbf{g} has no robust equilibrium.

For general games, we know by Brandenburger and Dekel [4] that whenever a game has several rationalizable outcomes, there are several a posteriori equilibria. In the generic class of games \mathbf{g} such that $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$, one can show that multiplicity of rationalizable outcomes implies multiplicity of *strict* a posteriori equilibria. Lemma 3.4 performs a similar construction as the one above to show that a generic game with multiple a posteriori equilibria has no robust equilibrium, which concludes the proof of the necessity part in Theorem 3.1.

3.2. Sufficiency

In this subsection, we show the sufficiency part of Theorem 3.1: that if an equilibrium is a unique rationalizable action profile, then it is robust to incomplete information under non-common priors. By the finiteness of A , there exists $N^* \geq 0$ such that $S^n[\mathbf{g}] = S^{N^*}[\mathbf{g}]$ for all $n \geq N^*$. Recall that if action distribution profile μ (or action distribution ξ) is $(N - 1)$ -robust, then it is N' -robust for all $N' \geq N - 1$. Hence, it suffices to show the following.

Proposition 3.2. *Let $N^* \geq 1$ be such that $S^n[\mathbf{g}] = S^{N^*}[\mathbf{g}]$ for all $n \geq N^*$. If $S^{N^*}[\mathbf{g}] = \{a^*\}$, then $([a^*])^J$ ($[a^*]$, resp.) is $(N^* - 1)$ -robust in \mathbf{g} .*

Thus, in order for a unique rationalizable outcome a^* to be robust, mutual knowledge of order N^* about the event “the payoffs are given by \mathbf{g} ” is needed, where N^* is the number of necessary elimination iteration rounds to reach the singleton $\{a^*\}$.

²⁰ This may be seen as “iterative dominance purification” of the a posteriori equilibrium (μ_1, μ_2) . See Corollary 3.5.

Proof of Proposition 3.2. Suppose that $S^{N^*}[\mathbf{g}] = \{a^*\}$. Fix any $\delta \in (0, 1)$. Now take $\varepsilon > 0$ such that $\varepsilon \leq \delta$. Consider any $(\varepsilon, N^* - 1)$ -perturbation \mathcal{U} and any Bayesian Nash equilibrium of \mathcal{U}, σ . Then, we have for all $i \in \mathcal{I}$,

$$\begin{aligned} \sigma_{P_i}(a^*) &\geq P_i([K_*]^{N^*-1}(\Omega_{\mathbf{g}})) \\ &= P_i\left(\bigcap_{n=1}^{N^*-1} [K_*]^n(\Omega_{\mathbf{g}})\right) \geq 1 - \varepsilon \geq 1 - \delta, \end{aligned}$$

where the first inequality follows from Lemma 2.1 and the second last inequality from the definition of $(\varepsilon, N^* - 1)$ -perturbation. A similar reasoning holds when considering $(\varepsilon, N^* - 1)$ -elaboration. \square

3.3. Necessity

In this subsection, we show the necessity part of Theorem 3.1: the dominance solvability is a necessary condition for a complete information game to have a robust equilibrium under non-common priors.

This shows that replacing common knowledge by mutual knowledge at an arbitrary high (but finite) level has far reaching consequences in particular when we drop the CPA. Indeed, under the CPA, several wider classes of games have been known in which a robust equilibrium exists (see the references cited in the Introduction). The result below shows that all these results heavily depend on the CPA.

Proposition 3.3. *Suppose that $W^\infty[\mathbf{g}] \neq \emptyset$. If $\mu^* \in (\Delta(A))^I$ ($\xi^* \in \Delta(A)$, resp.) is robust in \mathbf{g} , then $\mu^* = ([a^*])^I$ ($\xi^* = [a^*]$, resp.) for some $a^* \in A$ such that $W^\infty[\mathbf{g}] = \{a^*\}$.*

The following lemma is sufficient to prove the proposition. The proof of the lemma relies on a contagion argument for rationalizable action profiles or, to be more specific, for a set having the (strict) best response property (rather than for a single strict Nash equilibrium as often performed in the literature). Using this technique, in Corollary 3.5 we prove a result on what we call *iterative dominance purification* of (undominated) a posteriori equilibria, which allows us to prove the necessity part. It will also play a central role in the proofs for the set-valued robustness results in Section 4.

Recall from Definition 2.3 that a profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is a strict a posteriori equilibrium if for all $i \in \mathcal{I}$, $\mu_i(W^\infty[\mathbf{g}]) = 1$, and all $a_i \in A_i$,

$$\mu_i(a_i) > 0 \quad \Rightarrow \quad \{a_i\} = br_i(\mu_i(\cdot|a_i)).$$

Lemma 3.4. *Let $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ be a strict a posteriori equilibrium such that $\text{supp}(\mu_i) = \text{supp}(\mu_j)$ for all $i, j \in \mathcal{I}$. Then, for any $\varepsilon > 0$ and $N \geq 0$ there exists an (ε, N) -perturbation of \mathbf{g} such that there is a unique rationalizable strategy profile σ and it satisfies $\sigma_{P_i} = \mu_i$ for all $i \in \mathcal{I}$.*

Proof. Let $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ be as above. Denote $S = \text{supp}(\mu_i) (\subset W^\infty[\mathbf{g}])$ and $S_i = \{a_i \mid \mu_i(a_i) > 0\}$.

Given $\varepsilon \in (0, 1)$ and $N \geq 0$, let $\tilde{\varepsilon} = 1 - (1 - \varepsilon)^{1/(N+1)}$. For each $i \in \mathcal{I}$ and $a_i \in S_i$, there exists $\eta_i(a_i) > 0$ such that for all $\pi_i \in \Delta(A_{-i})$, if $|\pi_i - \mu_i(\cdot|a_i)| \leq \eta_i(a_i)$, then $\{a_i\} = br_i(\pi_i)$, which is

well defined by the continuity of g_i . Let $\eta = \min_{i \in \mathcal{I}} \min_{a_i \in S_i} \eta_i(a_i)$. Then, take any $r \geq 1$ large enough so that $r / \{r + (I - 1)(1 - \tilde{\varepsilon})\} \geq 1 - \eta$.

We now construct an (ε, N) -perturbation $\mathcal{U}^{\varepsilon, N}$ as follows. Let $\Omega = \mathcal{I} \times \mathbb{Z}_+ \times S$ and define $P_i \in \Delta(\Omega)$ for each $i \in \mathcal{I}$ by

$$P_i(i, k, a) = \frac{r}{r + I - 1} \tilde{\varepsilon}(1 - \tilde{\varepsilon})^k \mu_i(a)$$

and

$$P_i(j, k, a) = \frac{1}{r + I - 1} \tilde{\varepsilon}(1 - \tilde{\varepsilon})^k \mu_i(a)$$

for $j \neq i$. Let each \mathcal{Q}_i consist of (i) the events

$$E_{ia_i}^0 = \{(j, 0, a_i, a_{-i}) \mid j \neq i, (a_i, a_{-i}) \in S\},$$

and (ii) all the events of the form

$$E_{ia_i}^k = \{(i, k - 1, a_i, a_{-i}), (j, k, a_i, a_{-i}) \mid j \neq i, (a_i, a_{-i}) \in S\}$$

for each $k \geq 1$. Finally, define each $u_i : A \times \Omega \rightarrow \mathbb{R}$ by

$$u_i((a_i, a_{-i}), \omega) = \begin{cases} g_i(a_i, a_{-i}) & \text{if } \omega \notin \bigcup_{a'_i} E_{ia'_i}^0, \\ 1 & \text{if } \omega \in E_{ia_i}^0, \\ 0 & \text{if } \omega \in E_{ia'_i}^0 \text{ for } a'_i \neq a_i. \end{cases}$$

Verify that $\Omega_g = \{(i, k, a) \mid k \geq 1, i \in \mathcal{I} \text{ and } a \in S\}$ and $\bigcap_{n=1}^N [K_*]^n(\Omega_g) = \{(i, k, a) \mid k \geq N + 1, i \in \mathcal{I} \text{ and } a \in S\}$, so that $P_i(\bigcap_{n=1}^N [K_*]^n(\Omega_g)) = (1 - \tilde{\varepsilon})^{N+1} = 1 - \varepsilon$ for all $i \in \mathcal{I}$.

We first show that $\mathcal{U}^{\varepsilon, N}$ constructed above has a unique rationalizable strategy (recall Definition 2.8). For this, we prove by induction that for all $k = 1, 2, \dots$,

$$R_i^k[E_{ia_i}^{k-1}] = \{a_i\} \quad \text{for all } i \in \mathcal{I} \text{ and all } a_i \in S_i. \tag{*k}$$

First, $(*_1)$ holds true by construction.

Assume $(*_k)$. Fix any $i \in \mathcal{I}$ and $a_i \in S_i$, and take any conjecture of i , $v_i \in \Delta(\Omega \times A_{-i})$, such that $v_i(\{(\omega', a_{-i}) \mid a_{-i} \in R_{-i}^k[\omega']\}) = 1$ and $\text{marg}_\Omega v_i = P_i(\cdot | E_{ia_i}^k)$, where $R_{-i}^k[\omega'] = \prod_{j \neq i} R_j^k[\mathcal{Q}_j(\omega')]$. We show that $|\text{marg}_{A_{-i}} v_i - \mu_i(\cdot | a_i)| \leq \eta$.

Consider the states in $E_{ia_i}^k$ of the form $(i, k - 1, a_i, (a_j)_{j \neq i})$. By construction, such a state belongs to $E_{ja_j}^{k-1}$ for all $j \neq i$. Since $R_j^k[E_{ja_j}^{k-1}] = \{a_j\}$ by the induction hypothesis, the conjecture v_i must satisfy $v_i((i, k - 1, a_i, a_{-i}), a'_{-i}) = 0$ for all $a'_{-i} \neq a_{-i}$. Hence, we have

$$(\text{marg}_{A_{-i}} v_i)(a_{-i}) = v_i((i, k - 1, a_i, a_{-i}), a_{-i}) + \sum_{a'_{-i}} \sum_{j \neq i} v_i((j, k, a_i, a'_{-i}), a_{-i}), \tag{3.2}$$

and

$$\begin{aligned} v_i((i, k - 1, a_i, a_{-i}), a_{-i}) &= (\text{marg}_\Omega v_i)(i, k - 1, a_i, a_{-i}) \\ &= P_i((i, k - 1, a_i, a_{-i}) | E_{ia_i}^k). \end{aligned} \tag{3.3}$$

Now, by the construction of the state space,

$$\begin{aligned}
 & P_i((i, k - 1, a_i, a_{-i}) | E_{ia_i}^k) \\
 &= \frac{P_i((i, k - 1, a_i, a_{-i}))}{\sum_{a'_{-i}} P_i((i, k - 1, a_i, a'_{-i})) + \sum_{a'_{-i}} \sum_{j \neq i} P_i((j, k, a_i, a'_{-i}))} \\
 &= \frac{r \mu_i(a_i, a_{-i})}{\{r + (I - 1)(1 - \tilde{\varepsilon})\} \sum_{a'_{-i}} \mu_i(a_i, a'_{-i})} = c \mu_i(a_{-i} | a_i), \tag{3.4}
 \end{aligned}$$

and hence $v_i((i, k - 1, a_i, a_{-i}), a_{-i}) = c \mu_i(a_{-i} | a_i)$ by (3.3), where we denote $c = r / \{r + (I - 1)(1 - \tilde{\varepsilon})\}$. Since $\sum_{a'_{-i}} v_i((i, k - 1, a_i, a'_{-i}), a'_{-i}) = c$, it therefore follows from (3.2) that, for all a_{-i} ,

$$(\text{marg}_{A_{-i}} v_i)(a_{-i}) \leq c \mu_i(a_{-i} | a_i) + (1 - c),$$

so that

$$\begin{aligned}
 |(\text{marg}_{A_{-i}} v_i)(a_{-i}) - \mu_i(a_{-i} | a_i)| &\leq (1 - c)(1 - \mu_i(a_{-i} | a_i)) \\
 &\leq 1 - c \leq \eta,
 \end{aligned}$$

where the last inequality follows from the choice of r . We thus have $|\text{marg}_{A_{-i}} v_i - \mu_i(\cdot | a_i)| \leq \eta$.

Thus by the definition of μ_i and the choice of η , we have $BR_i(v_i) = br_i(\text{marg}_{A_{-i}} v_i) = \{a_i\}$. Since v_i has been taken arbitrarily, it follows that $R_i^{k+1}[E_{ia_i}^k] = \{a_i\}$, hence $(*_{k+1})$.

Finally, if σ is the unique rationalizable strategy of $\mathcal{U}^{\varepsilon, N}$, we have $\sigma_{P_i}(a) = \mu_i(a)$ for all $a \in A$ by construction. \square

Note that the (ε, N) -perturbation constructed above satisfies common support and \mathcal{L} -tail consistency for some $\mathcal{L} \geq 1$ (where \mathcal{L} depends on ε and N).

As a corollary of Lemma 3.4, we obtain the iterative dominance purification result, that we can purify any undominated a posteriori equilibrium by a unique rationalizable strategy profile of a dominance solvable (ε, N) -perturbation.

Corollary 3.5. *Let $(\mu_i)_{i \in \mathcal{I}}$ be any undominated a posteriori equilibrium. For any $\delta > 0$, $\varepsilon > 0$, and $N \geq 0$, there exists an (ε, N) -perturbation of \mathbf{g} such that there is a unique rationalizable strategy profile σ and it satisfies $|\sigma_{P_i} - \mu_i| \leq \delta$ for all $i \in \mathcal{I}$.*

Proof. Let $(\mu_i)_{i \in \mathcal{I}}$ be an undominated a posteriori equilibrium. Fix $\delta > 0$. By the continuity of g_i 's, we can take a strict a posteriori equilibrium $(\mu'_i)_{i \in \mathcal{I}}$ with $\text{supp}(\mu'_i) = W^\infty[\mathbf{g}]$ for all $i \in \mathcal{I}$ such that $|\mu'_i - \mu_i| \leq \delta$. Hence by Lemma 3.4, for any $\varepsilon > 0$, and $N \geq 0$ there exists an (ε, N) -perturbation of \mathbf{g} such that a unique rationalizable strategy profile σ satisfies $\sigma_{P_i} = \mu'_i$ and hence $|\sigma_{P_i} - \mu_i| \leq \delta$ for all $i \in \mathcal{I}$, which completes the proof. \square

Note that this corollary in fact proves that in the generic class of games where $W^\infty[\mathbf{g}] = S^\infty[\mathbf{g}]$, any a posteriori equilibrium can be purified in the previous sense.

We now prove the necessity part for our robustness result.

Proof of Proposition 3.3. Suppose that $W^\infty[\mathbf{g}] \neq \emptyset$ and that $(\mu_i^*)_{i \in \mathcal{I}}$ (ξ^* , resp.) is robust. Observe that the set of strict a posteriori equilibria $(\mu_i)_{i \in \mathcal{I}}$ such that $\text{supp}(\mu_i) = W^\infty[\mathbf{g}]$ for all $i \in \mathcal{I}$ is non-empty. Take any such strict a posteriori equilibrium $(\mu_i)_{i \in \mathcal{I}}$ of \mathbf{g} . Then, by Lemma 3.4, for any $\varepsilon > 0$ and $N \geq 0$ there exists an (ε, N) -perturbation where the unique rationalizable

strategy profile σ satisfies $\sigma_{P_i} = \mu_i$ for all $i \in \mathcal{I}$. Hence, it must be true that $\mu_i = \mu_i^*$ for all $i \in \mathcal{I}$ ($\mu_i = \xi^*$ for all $i \in \mathcal{I}$, resp.). This implies that $(\mu_i^*)_{i \in \mathcal{I}}$ ($(\xi^*)^I$, resp.) is the unique strict a posteriori equilibrium, so that $W^\infty[\mathbf{g}]$ is a singleton set, say, $\{a^*\}$, and thus $\mu_i^* = [a^*]$ for all $i \in \mathcal{I}$ ($\xi^* = [a^*]$, resp.). \square

Let us relate our results to previous studies. Weinstein and Yildiz [24] show that for any complete information type in the universal type space²¹ and any rationalizable action profile a^* of this game, there exist a dominance-solvable incomplete information game and a sequence of types drawn from this game such that (1) this sequence converges to the complete information type (with respect to the product topology in the universal type space) and (2) each type of the sequence plays a^* , while they are not concerned with priors but working directly with interim beliefs. Our construction in Lemma 3.4 shows that, by attaching heterogeneous priors appropriately, such a dominance solvable incomplete information game can be made an (ε, N) -perturbation (where ε can be arbitrarily small and N arbitrarily large).²² In addition, Corollary 3.5 shows that the unique equilibrium of this (ε, N) -perturbation can be fully characterized by using the notion of a posteriori equilibrium. Whenever a^* is a strict Nash equilibrium, the unique equilibrium of this dominance-solvable game will play a^* everywhere. However, when a^* is not a strict Nash equilibrium, this is not possible: a^* cannot be played everywhere (as the proof of Lemma 3.4 demonstrates, action profiles different from a^* have also to be contagious). In proving our complete characterization result, the use of our iterative dominance purification argument that relies on a posteriori equilibrium becomes crucial.

If, as in Kajii and Morris [10], one considers (ε, N) -perturbations of \mathbf{g} that satisfy the CPA, then any equilibrium action distribution of such an (ε, N) -perturbation must be close to some *objective* correlated equilibrium of \mathbf{g} when ε is small (Corollary 3.5 in Kajii and Morris [10]). In contrast, with non-common priors, equilibrium action distributions are close to *subjective* correlated equilibria (Lemma 2.2 in Section 2.4). The former fact allows Kajii and Morris [10] to establish that if \mathbf{g} has a unique objective correlated equilibrium, then it is a robust equilibrium under common priors, while, as our result shows, it is not robust under non-common priors unless \mathbf{g} itself is dominance solvable (or equivalently, \mathbf{g} has a unique a posteriori equilibrium).

Lipman [12,13] shows that given any partition model with non-common priors (and tail consistency) and any state ω in the model, for any finite N one can construct a partition model with a common prior such that there is a state in that model that has the same higher order beliefs up to order N as those at ω . Lipman's [13] construction can be applied as well to our (ε, N) -perturbation in the proof of Lemma 3.4 to have a further incomplete information perturbation with a common prior. One might then ask whether we could use the further perturbation with a common prior so obtained for the purpose of Lemma 3.4. It is, however, not true; for, when the given a posteriori equilibrium is not a robust equilibrium under common prior (as in the example in Section 3.1), it is not possible, for vanishingly small ε , to have an (ε, N) -perturbation such that (1) it has a common prior, (2) it is dominance solvable, and (3) the unique rationalizable strat-

²¹ Recall that a complete information type is a (degenerate) type in the universal type space where it is common knowledge that payoffs are given by the complete information game.

²² While Weinstein and Yildiz [24] use a different setting and language, their construction as well as ours are in essence analogous to the information structure of Rubinstein's [20] email game. We could have worked with type spaces as in Weinstein and Yildiz [24], but chose the current setup (with an infinitely countable state space) for comparison with Kajii and Morris [10].

egy profile generates an action distribution arbitrarily close to that a posteriori equilibrium. The additional constraints that the CPA imposes on perturbations are discussed further in Section 5.

4. Set-valued robustness

Given that many games possess no robust equilibrium, it is natural to consider a set-valued robustness concept. Such an idea can be found for instance in Morris and Ui [17] where the common prior is assumed. In the following, we define robustness for sets of action distribution profiles as well as those of action distributions. We give a separate treatment to these two notions since, contrary to their point-valued versions, they lead to distinct characterizations. We then show our second main result, that for any generic game \mathbf{g} in which $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$, a smallest robust set of action distribution profiles (action distributions, resp.) exists and coincides with the set of a posteriori equilibria of \mathbf{g} (the convex hull of $S^\infty[\mathbf{g}]$, resp.).

4.1. Robust sets of action distribution profiles

Let us first define the robustness of sets of action distribution profiles.

Definition 4.1. A product set of action distribution profiles $M = \prod_{i \in \mathcal{I}} M_i \subset (\Delta(A))^I$ is N -robust to incomplete information under non-common priors, or simply, N -robust, if it is closed, and for all $\delta > 0$, there exists $\varepsilon > 0$ such that any (ε, N) -perturbation of \mathbf{g} has a Bayesian Nash equilibrium σ such that for all $i \in \mathcal{I}$, there exists $\mu_i \in M_i$ with $|\mu_i - \sigma_{P_i}| \leq \delta$.

M is robust to incomplete information under non-common priors, or simply, robust, if there exists $N \geq 0$ such that M is N -robust.

Observe that if M is (N) -robust, then any $M' \supset M$ is (N) -robust. In particular, $(\Delta(A))^I$ is N -robust for all $N \geq 0$ and thus robust. We say that M is a minimal (N) -robust set if it is an (N) -robust set and no proper subset of it is an (N) -robust set; and that M is a smallest (N) -robust set if it is an (N) -robust set and is contained in any (N) -robust set.

The following proposition establishes that minimal robust set is well defined for all finite games, while we will show that a smallest robust set actually exists for generic games.

Proposition 4.1. Any game has a minimal N -robust set for each $N \geq 0$ and a minimal robust set.

Proof. We show the existence of a minimal N -robust set; the existence of a minimal robust set can be proved analogously.

Let (\mathcal{M}, \subset) be the (non-empty) collection of N -robust sets partially ordered by set inclusion. We show that (\mathcal{M}, \subset) has a minimal element. Take any totally ordered subset of \mathcal{M} and denote it by \mathcal{M}' . Let $M^* = \bigcap_{M \in \mathcal{M}'} M$. Since each $M \in \mathcal{M}'$ is non-empty and closed in a compact set $(\Delta(A))^I$, so is M^* . We want to show that M^* is N -robust and therefore is a lower bound of \mathcal{M}' in \mathcal{M} . Then, it follows from Zorn's lemma that \mathcal{M} has a minimal element.

Fix any $\delta > 0$. By the compactness of $(\Delta(A))^I$, we can take an $M' \in \mathcal{M}'$ such that $M' \subset V_{\delta/2}(M^*)$. By definition, there exists $\varepsilon > 0$ such that any (ε, N) -elaboration has a Bayesian Nash equilibrium σ such that for some $(\mu'_i)_{i \in \mathcal{I}} \in M'$, $|\mu'_i - \sigma_{P_i}| \leq \delta/2$ for all $i \in \mathcal{I}$. But we can take $(\mu_i)_{i \in \mathcal{I}} \in M^*$ such that $|\mu_i - \mu'_i| \leq \delta/2$ for all $i \in \mathcal{I}$, and hence $|\mu_i - \sigma_{P_i}| \leq \delta$ for all $i \in \mathcal{I}$. This implies that M^* is N -robust, completing the proof. \square

In characterizing robust sets of action distribution profiles, the concept of a posteriori equilibrium is the key notion. Recall that a profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is an a posteriori equilibrium (N -subjective correlated equilibrium, resp.) of \mathbf{g} if it is a subjective correlated equilibrium of \mathbf{g} and $\mu_i(S^\infty[\mathbf{g}]) = 1$ ($\mu_i(S^N[\mathbf{g}]) = 1$, resp.) for all $i \in \mathcal{I}$, and that $\mathcal{E}[\mathbf{g}]$ ($\mathcal{E}^N[\mathbf{g}]$, resp.) denotes the set of a posteriori equilibria (N -subjective correlated equilibria, resp.) of \mathbf{g} . We show that for generic games, a smallest robust set of action distribution profiles exists and coincides with $\mathcal{E}[\mathbf{g}]$.

Theorem 4.2. *Suppose that $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$. Then, $\mathcal{E}[\mathbf{g}]$ is the smallest robust set of \mathbf{g} .*

The proof proceeds in two steps. We first show that $\mathcal{E}^N[\mathbf{g}]$ is $(N - 1)$ -robust, which implies that $\mathcal{E}[\mathbf{g}]$ is robust. Then we show that any robust set contains the set of undominated a posteriori equilibria, $\mathcal{E}^u[\mathbf{g}]$.

From Lemma 2.2 in Section 2.4, we immediately have the following.²³

Proposition 4.3. *$\mathcal{E}^N[\mathbf{g}]$ is $(N - 1)$ -robust.*

Since $\mathcal{E}^u[\mathbf{g}]$ coincides with $\mathcal{E}[\mathbf{g}]$ whenever $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$, the next proposition is sufficient to complete the proof of Theorem 4.2.

Proposition 4.4. *If M is robust in \mathbf{g} , then $\mathcal{E}^u[\mathbf{g}] \subset M$.*

Proof. Let $M = \prod_{i \in \mathcal{I}} M_i$ be robust in \mathbf{g} . Take any $(\mu_i)_{i \in \mathcal{I}} \in \mathcal{E}^u[\mathbf{g}]$. By Corollary 3.5, for each $i \in \mathcal{I}$, $V_\delta(\mu_i) \cap M_i \neq \emptyset$ for any $\delta > 0$. By the closedness of M , it follows that $\mu_i \in M_i$ for each $i \in \mathcal{I}$. \square

4.2. Robust sets of action distributions

As for the point-valued robustness concept, we consider the following alternative concept.

Definition 4.2. A set of action distributions $\mathcal{E} \subset \Delta(A)$ is N -robust to incomplete information under non-common priors, or simply, N -robust, if it is closed, and for all $\delta > 0$, there exists $\varepsilon > 0$ such that for any (ε, N) -elaboration of \mathbf{g} , (\mathcal{U}, P_0) , \mathcal{U} has a Bayesian Nash equilibrium σ such that there exists $\xi \in \mathcal{E}$ with $|\xi - \sigma_{P_0}| \leq \delta$.

A set of action distributions $\mathcal{E} \subset \Delta(A)$ is robust to incomplete information under non-common priors, or simply, robust, if there exists $N \geq 0$ such that \mathcal{E} is N -robust.

Observe that if \mathcal{E} is (N) -robust, then any $\mathcal{E}' \supset \mathcal{E}$ is (N) -robust. In particular, $(\Delta(A))$ is N -robust for all $N \geq 0$ and thus robust. We say that \mathcal{E} is a minimal (N) -robust set if it is an (N) -robust set and no proper subset of it is an (N) -robust set; and that \mathcal{E} is a smallest (N) -robust set if it is an (N) -robust set and is contained in any (N) -robust set. The existence of minimal robust set can be verified in the same way as in Proposition 4.1.

²³ Lemma 2.2 in fact implies a stronger result that for all $\delta > 0$, there exists $\varepsilon > 0$ such that for any Bayesian Nash equilibrium σ of any (ε, N) -elaboration, there exists $(\mu_i)_{i \in \mathcal{I}} \in \mathcal{E}^N[\mathbf{g}]$ such that $|\mu_i - \sigma_{P_i}| \leq \delta$ for all $i \in \mathcal{I}$ (recall that our robustness test only requires that there exist such a Bayesian Nash equilibrium).

We show that for generic games, a smallest robust set of action distributions exists and coincides with the convex hull of the set of rationalizable action profiles of \mathbf{g} .

Theorem 4.5. *Suppose that $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$. Then, $\Delta(S^\infty[\mathbf{g}])$ is the smallest robust set of \mathbf{g} .*

Note that the elements of the convex hull $\Delta(S^\infty[\mathbf{g}])$ represent the modeler’s predictions over A generated via his own subjective prior P_0 in (ε, N) -elaborations, and thus some elements may not be part of any a posteriori equilibrium, since P_0 may be unrelated to the players’ priors (provided that it assigns probability at least $1 - \varepsilon$ to the event that $\Omega_{\mathbf{g}}$ is mutually known up to order N).

The proof proceeds in two propositions. The first proposition implies that $\Delta(S^\infty[\mathbf{g}])$ is $(N^* - 1)$ -robust, where N^* is such that $\Delta(S^\infty[\mathbf{g}]) = \Delta(S^{N^*}[\mathbf{g}])$, which is well defined by the finiteness of A . The second proves that any robust set contains $\Delta(W^\infty[\mathbf{g}])$.

Proposition 4.6. *For any $N \geq 1$, $\Delta(S^N[\mathbf{g}])$ is $(N - 1)$ -robust.*

Proof. Fix any $N \geq 1$ and any $\delta > 0$. Now take $\varepsilon > 0$ such that $\varepsilon \leq \delta$. Consider any $(\varepsilon, N - 1)$ -elaboration (\mathcal{U}, P_0) and any rationalizable strategy profile of \mathcal{U} , σ . Then, we have

$$\begin{aligned} \sigma_{P_0}(S^N[\mathbf{g}]) &\geq P_0([K_*]^{N-1}(\Omega_{\mathbf{g}})) \\ &= P_0\left(\bigcap_{n=1}^{N-1} [K_*]^n(\Omega_{\mathbf{g}})\right) \geq 1 - \varepsilon \geq 1 - \delta, \end{aligned}$$

where the first inequality follows from Lemma 2.1 and the second last inequality from the definition of $(\varepsilon, N - 1)$ -elaboration. Hence, we have that $|\sigma_{P_0} - \xi| \leq \delta$ for some $\xi \in \Delta(S^N(\mathbf{g}))$, which shows that $\Delta(S^N[\mathbf{g}])$ is N -robust. \square

Proposition 4.7. *If \mathcal{E} is robust in \mathbf{g} , then $\Delta(W^\infty[\mathbf{g}]) \subset \mathcal{E}$.*

Proof. Suppose that \mathcal{E} is a robust set, which is closed by definition, and let $N \geq 0$ be such that \mathcal{E} is N' -robust for all $N' \geq N$. Assume that $W^\infty[\mathbf{g}] \neq \emptyset$ and take any $\xi \in \Delta(W^\infty[\mathbf{g}])$ with $\text{supp}(\xi) = W^\infty[\mathbf{g}]$. We show that for all $\varepsilon > 0$, there exists an (ε, N) -elaboration $(\mathcal{U}^{\varepsilon, N}, P_0)$ such that any rationalizable strategy profile σ of \mathcal{U} satisfies $\sigma_{P_0} = \xi$. Then, this implies that $V_\delta(\xi) \cap \mathcal{E} \neq \emptyset$ for all $\delta > 0$, so that $\xi \in \mathcal{E}$ due to the closedness of \mathcal{E} .

Take any $\varepsilon > 0$, and any strict a posteriori equilibrium $(\mu_i)_{i \in \mathcal{I}}$ such that $\text{supp}(\mu_i) = W^\infty[\mathbf{g}]$ for all $i \in \mathcal{I}$. Then, let $\mathcal{U}^{\varepsilon, N}$ be the (ε, N) -perturbation as in the proof of Lemma 3.4, and let P_0 be defined by

$$P_0(1, k, a) = \frac{r}{r + I - 1} \tilde{\varepsilon}(1 - \tilde{\varepsilon})^k \xi(a)$$

and

$$P_0(j, k, a) = \frac{1}{r + I - 1} \tilde{\varepsilon}(1 - \tilde{\varepsilon})^k \xi(a)$$

for $j \neq 1$. By Lemma 3.4, we know that the unique rationalizable strategy profile σ of $\mathcal{U}^{\varepsilon, N}$ satisfies $\sigma(a|j, k, a) = 1$ for all $(j, k, a) \in \Omega$. By construction, $P_0(\bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}})) \geq 1 - \varepsilon$ and $\sigma_{P_0} = \xi$.

Now we have that the set of distributions with support equal to $W^\infty[\mathbf{g}]$ is contained in \mathcal{E} . But the closure of this set is actually $\Delta(W^\infty[\mathbf{g}])$. Hence, since \mathcal{E} is closed, we must have $\Delta(W^\infty[\mathbf{g}]) \subset \mathcal{E}$. \square

5. Discussion

In this section, we examine the critical assumptions for our results to hold. We also discuss the relationship with other robustness notions as defined in Weinstein and Yildiz [23,24].

5.1. N -robustness

In Section 3.2, we showed that if a^* is a unique rationalizable action profile in \mathbf{g} and if an incomplete information game \mathcal{U} is an $(\varepsilon, N^* - 1)$ -elaboration of \mathbf{g} , where N^* is such that $S^{N^*}[\mathbf{g}] = \{a^*\}$, then in any Bayesian Nash equilibrium of \mathcal{U} , a^* is played with high probability with respect to any player’s prior (whenever ε is small with respect to these priors). In this subsection, we show that mutual knowledge of $\Omega_{\mathbf{g}}$ at high order is also necessary for such an action profile a^* to be played in Bayesian Nash equilibria. To see this, we examine the concept of N -robustness.

Proposition 5.1. *Suppose that $S^N[\mathbf{g}] = W^N[\mathbf{g}]$. Then, $\mathcal{E}^N[\mathbf{g}]$ is the smallest $(N - 1)$ -robust set of \mathbf{g} .*

In the sequel, we say that a profile of action distributions $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is a *strict N -subjective correlated equilibrium* if it is an undominated N -subjective correlated equilibrium and for all $i \in \mathcal{I}$ and all $a_i \in A_i$,

$$\mu_i(a_i) > 0 \quad \Rightarrow \quad \{a_i\} = br_i(\mu_i(\cdot|a_i)).$$

Lemma 5.2. *Fix any $N \geq 1$. Let $(\mu_i^N)_{i \in \mathcal{I}} \in (\Delta(A))^I$ be a strict N -subjective correlated equilibrium. Then, for any $\delta > 0$ and $\varepsilon > 0$, there exists an $(\varepsilon, N - 1)$ -perturbation of \mathbf{g} such that any rationalizable strategy profile σ satisfies $|\sigma_{P_i} - \mu_i| < \delta$ for all $i \in \mathcal{I}$.*

Proof. See Appendix A.1. \square

As in Section 3.3, we can derive the following corollary, which is analogous to Corollary 3.5.

Corollary 5.3. *Fix any $N \geq 1$. Let $(\mu_i)_{i \in \mathcal{I}}$ be an undominated N -subjective correlated equilibrium. For any $\delta > 0$ and $\varepsilon > 0$, there exists an $(\varepsilon, N - 1)$ -perturbation of \mathbf{g} such that any rationalizable strategy profile σ satisfies $|\sigma_{P_i} - \mu_i| \leq \delta$ for all $i \in \mathcal{I}$.*

Using the same argument as in Section 4, we can complete the proof of Proposition 5.1.

Remark that if we consider 0-robustness as an extension to heterogeneous priors of the robustness test defined in Kajii and Morris [10], we obtain that for a generic class of games, a profile of action distribution $(\mu_i)_{i \in \mathcal{I}}$ is robust à la Kajii and Morris [10] without a common prior if and

only if $\mu_i = [a^*]$ for all $i \in \mathcal{I}$ where a_i^* is a strictly dominant action for each player $i \in \mathcal{I}$.²⁴ This result generalizes Proposition 12 in Weinstein and Yildiz [23]. Indeed, as Weinstein and Yildiz [24, Section 8] point out, without common prior, the restriction only on the prior probabilities of $\Omega_{\mathbf{g}}$ has no implication for conditional beliefs beyond second order.

5.2. Uniform bound on posteriors

In Section 3.3, we showed that if \mathbf{g} has more than one rationalizable action profiles, then for each such profile, we can construct an (ε, N) -perturbation whose Bayesian Nash equilibrium plays this profile with probability zero. The crucial point for this result is that relevant posterior probabilities can be arbitrarily close to one *simultaneously* for the players (see (3.4) in the proof of Lemma 3.4) by choosing heterogeneous priors sufficiently different from each other.

In this subsection, we examine how the results change if we require that the elaborations uniformly satisfy \mathcal{L} -tail consistency (uniformly over elaborations of a complete information game). Indeed, when such an assumption is made, given \mathcal{L} there are generic games with multiple rationalizable action profiles that have a robust equilibrium. To extract the effect of this restriction, we consider ε -perturbation (i.e., $(\varepsilon, 0)$ -perturbation), whereas the result below would hold if we considered (ε, N) -perturbations. Given an ε -perturbation, let

$$\rho((P_i)_{i \in \mathcal{I}}) = \max_{i \neq j} \sup_{\omega \in \Omega} \frac{P_i(\omega)}{P_j(\omega)}$$

with a convention that $q/0 = \infty$ for $q > 0$, and $0/0 = 1$. Note that $\rho((P_i)_{i \in \mathcal{I}}) < \infty$ only if $(P_i)_{i \in \mathcal{I}}$ has common support.

Definition 5.1. An action distribution $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ is *r-robust* if for all $\delta > 0$, there exists $\varepsilon > 0$ such that any ε -perturbation of \mathbf{g} such that $\rho((P_i)_{i \in \mathcal{I}}) \leq r$ has a Bayesian Nash equilibrium σ such that $|\mu_i - \sigma_{P_i}| \leq \delta$ for all $i \in \mathcal{I}$.

We use p -belief operators as defined in Monderer and Samet [15]. For any number $p \in (0, 1]$, the p -belief operator for player i , $B_i^p : 2^\Omega \rightarrow 2^\Omega$, 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). Let $B_*^p(E) = \bigcap_{i \in \mathcal{I}} B_i^p(E)$ be the set of states where E is mutually p -believed, i.e., where every player believes E with probability at least p . Finally, an event E is common p -belief at state ω if $\omega \in C^p(E) = \bigcap_{n=1}^\infty [B_*^p]^n(E)$.

Observe that for any event E , we have by definition of knowledge operators that $K_*(E) \subset E$. On the other hand, this inclusion is not necessarily true when replacing knowledge operators by p -belief operators. The following lemma shows that this inclusion remains true for simple events.

Lemma 5.4. For any simple event E and any $p \in (0, 1]$, $B_*^p(E) \subset E$.

²⁴ A same result would hold if we considered robustness of action distributions instead of profiles of action distributions. In the case of sets of action distributions, the proposition analogous to Proposition 5.1 would state that for generic games, the smallest $(N - 1)$ -robust set coincides with $\Delta(S^N[\mathbf{g}])$.

Proof. Let E be a simple event, and $E = \bigcap_{i \in \mathcal{I}} E_i$ where $E_i \in \mathcal{F}_i$ for each $i \in \mathcal{I}$. Observe that $B_i^p(E_i) = E_i$ for any $p > 0$. By the monotonicity of $B_i^p(\cdot)$, we have $B_i^p(E) \subset B_i^p(E_i) = E_i$. It follows that $\bigcap_{i \in \mathcal{I}} B_i^p(E) \subset \bigcap_{i \in \mathcal{I}} E_i$, as claimed. \square

The following result, which corresponds to the critical path result of Kajii and Morris [10, Proposition 4.2] in our context with non-common priors, shows that with a uniform bound on posteriors, the ex ante probability (with respect to any player’s prior) of the event $C^p(E)$ is bounded from below uniformly in all information systems.

Proposition 5.5. *For any $r > 0$, if $p < 1/\{1 + r(I - 1)\}$, then in any information system $IS = [\Omega, (P_i)_{i \in \mathcal{I}}, (Q_i)_{i \in \mathcal{I}}]$ with $\rho((P_i)_{i \in \mathcal{I}}) \leq r$, any simple event E satisfies*

$$P_j(C^p(E)) \geq 1 - \frac{1 - p}{(1 - \{1 + r(I - 1)\}p)} \max_{i \in \mathcal{I}} (1 - P_i(E)) \tag{5.1}$$

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

Proof. See Appendix A.2. \square

Conversely, if $p \geq 1/\{1 + r(I - 1)\}$, then one can find an information system with $\rho((P_i)_{i \in \mathcal{I}}) \leq r$ and a simple event E such that (5.1) does not hold. Indeed, one can show that the information system given in the proof of Lemma 3.4 is such an example.

To give a sufficient condition for r -robustness, we use the notion of p -dominant equilibrium as introduced by Morris, Rob, and Shin [16] and Kajii and Morris [10].

Definition 5.2. Let $p \in [0, 1]$. Action profile $a^* \in A$ is a p -dominant equilibrium in \mathbf{g} if for all $i \in \mathcal{I}$,

$$a_i^* \in br_i(\pi_i)$$

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

We also use the following lemma, a straightforward corollary of Kajii and Morris [10, Lemma 5.2], which relates the notion of common p -belief to that of p -dominance.

Lemma 5.6. *Suppose that action profile a^* is a p -dominant equilibrium of \mathbf{g} . Consider any incomplete information perturbation \mathcal{U} of \mathbf{g} . Then, \mathcal{U} has a Bayesian Nash equilibrium where $\sigma_i(a_i^* | \omega) = 1$ for all $i \in \mathcal{I}$ and $\omega \in C^p(\Omega_{\mathbf{g}})$.*

Proposition 5.7. *Suppose that action profile a^* is a p -dominant equilibrium of \mathbf{g} where $p < 1/(1 + r(I - 1))$. Then, $[a^*]^I$ is r -robust.*

Proof. Let a^* be a p -dominant equilibrium with $p < 1/(1 + r(I - 1))$. Fix any $\delta > 0$. By Proposition 5.5, we can choose $\varepsilon > 0$ such that for any information system and any simple event E , $P_i(E) \geq 1 - \varepsilon$ implies $P_i(C^p(E)) \geq 1 - \delta$. Thus, by the choice of ε , for any ε -perturbation \mathcal{U} of \mathbf{g} with $\rho((P_i)_{i \in \mathcal{I}}) \leq r$, we have $P_i(C^p(\Omega_{\mathbf{g}})) \geq 1 - \delta$. By Lemma 5.6, it follows that there exists a Bayesian Nash equilibrium σ of \mathcal{U} with $\sigma_{P_i}(a^*) \geq P_i(C^p(\Omega_{\mathbf{g}})) \geq 1 - \delta$, meaning that $[a^*]$ is r -robust. \square

5.3. Interim robustness

We discuss the robustness concept due to Weinstein and Yildiz [23, Definition 10].²⁵

Definition 5.3. (\mathcal{U}, ω) is an N -perturbation of \mathbf{g} if $\omega \in \bigcap_{n=1}^N (K_*)^n(\Omega_{\mathbf{g}})$.

Definition 5.4. $a^* \in A$ is interim robust to incomplete information in \mathbf{g} if there exists $N \geq 0$ such that for any N -perturbation of \mathbf{g} , (\mathcal{U}, ω) , such that \mathcal{U} satisfies the CPA, \mathcal{U} has a Bayesian Nash equilibrium σ^* such that $\sigma^*(\omega) = [a^*]$.

Note that it is not required that $P(\Omega_{\mathbf{g}})$ be large in \mathcal{U} , where P is the common prior. Weinstein and Yildiz [23, Proposition 11] showed the following.²⁶

Proposition 5.8. Suppose that $S^\infty[\mathbf{g}] = W^\infty[\mathbf{g}]$. Then, a^* is interim robust in \mathbf{g} if and only if $S^\infty[\mathbf{g}] = \{a^*\}$.

Let us prove that we can derive this result from our previous results.

Lemma 5.9. For any $\varepsilon > 0$ and $N \geq 0$, there exists an (ε, N) -perturbation of \mathbf{g} with a common prior such that for all $a \in W^\infty[\mathbf{g}]$, there exists $\omega \in \bigcap_{n=1}^N (K_*)^n(\Omega_{\mathbf{g}})$ such that $\sigma(\omega) = [a]$ for any rationalizable strategy profile σ .

Proof. Let $(\mu_i)_{i \in \mathcal{I}} \in (\Delta(A))^I$ be a strict a posteriori equilibrium such that $\text{supp}(\mu_i) = W^\infty[\mathbf{g}]$ for all $i \in \mathcal{I}$. Consider the (ε, N) -perturbation $\mathcal{U}^{\varepsilon, N}$ as built in Lemma 3.4 but where $r \geq 1$ is chosen large enough so that for all $i \in \mathcal{I}$,

$$\frac{\mu_i(a_i, a_{-i})}{\mu_i(a_i) + \frac{1-\tilde{\varepsilon}}{r} \sum_{j \neq i} \mu_j(a_i)} \geq (1 - \eta)\mu_i(a_{-i}|a_i)$$

for all $a_i \in W_i^\infty$ and $a_{-i} \in W_{-i}^\infty$ (where η is chosen as in Lemma 3.4). In addition, the common prior over $\Omega = \mathcal{I} \times \mathbb{Z}_+ \times W^\infty[\mathbf{g}]$, P , is defined as follows. For each $i \in \mathcal{I}$ and $a \in W^\infty[\mathbf{g}]$, let

$$P(i, k, a) = \frac{1}{I} \tilde{\varepsilon} \left(\frac{1 - \tilde{\varepsilon}}{r} \right)^k \mu_i(a)$$

for all $k \leq N + 2$. Observe that $\sum_{i \in \mathcal{I}} \sum_{k=0}^{N+2} \sum_{a \in W^\infty[\mathbf{g}]} P(i, k, a) < 1$. Then, for all $k \geq N + 3$, define $P(i, k, a)$ so that $P(\Omega) = 1$.

Now let σ be any rationalizable strategy profile of $\mathcal{U}^{\varepsilon, N}$. We show that for all $i \in \mathcal{I}$ and all $a_i \in W_i^\infty[\mathbf{g}]$, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in \bigcup_{k=0}^{N+2} E_{ia_i}^k$. By construction, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^0$ and all $i \in \mathcal{I}$. Then suppose that for $k \leq N + 1$, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^k$ and all $i \in \mathcal{I}$. Consider any $\omega \in E_{ia_i}^{k+1}$. By the construction of the state space, we have

$$P((i, k, a_i, a_{-i})|Q_i(\omega))$$

²⁵ Note that the exact definition in Weinstein and Yildiz [23, Definition 10] uses $B_*^{1-\varepsilon}$ instead of K_* . One can show that our proofs would work with their formulation.

²⁶ One can also show this result using Lemma 3.4 and Lipman [13, Theorem 1].

$$\begin{aligned}
 &= \frac{P((i, k, a_i, a_{-i}))}{\sum_{a'_{-i}} P((i, k, a_i, a'_{-i})) + \sum_{a'_{-i}} \sum_{j \neq i} P((j, k + 1, a_i, a'_{-i}))} \\
 &= \frac{\mu_i(a_i, a_{-i})}{\mu_i(a_i) + \frac{1-\varepsilon}{r} \sum_{j \neq i} \mu_j(a_i)},
 \end{aligned}$$

so that

$$|P((i, k, a_i, a_{-i}) | Q_i(\omega)) - \mu_i(a_{-i} | a_i)| \leq \eta \mu_i(a_{-i} | a_i) \leq \eta,$$

where the last inequality follows from the choice of r . Thus, by the definition of μ_i and the choice of η as well as the induction hypothesis, we have $\sigma_i(a_i | \omega) = 1$. The proof is completed observing that $(i, N + 1, a) \in E_{ia_i}^{N+2} \cap \bigcap_{j \neq i} E_{ja_j}^{N+1}$, so that a is played at $(i, N + 1, a)$ at any rationalizable strategy profile and that $(i, N + 1, a) \in \bigcap_{n=1}^N (K_*)^n(\Omega_{\mathbf{g}})$. \square

Proof of Proposition 5.8. The sufficiency part follows from Lemma 2.1. To show the necessity, assume that a^* is interim robust to incomplete information. Take any $a \in W^\infty[\mathbf{g}]$. Note that the set of strict a posteriori equilibria $(\mu_i)_{i \in \mathcal{I}}$ such that $\text{supp}(\mu_i) = W^\infty[\mathbf{g}]$ is non-empty. Consider any such $(\mu_i)_{i \in \mathcal{I}}$. Then, by Lemma 5.9, we have that for all $\varepsilon > 0$ and $N \geq 0$, there exists an (ε, N) -perturbation of \mathbf{g} , $\mathcal{U}^{\varepsilon, N}$ such that (1) it satisfies the CPA; and (2) there is some $\omega \in \bigcap_{n=1}^N [K_*]^n(\Omega_{\mathbf{g}})$ where $\sigma(\omega) = [a]$ for any rationalizable strategy profile σ . Since $(\mathcal{U}^{\varepsilon, N}, \omega)$ is an N -perturbation of \mathbf{g} satisfying the CPA, we must have that $|W^\infty[\mathbf{g}]| = 1$ and hence $W^\infty[\mathbf{g}] = \{a^*\}$. \square

In the constructed perturbation with a common prior P , when $P(\Omega_{\mathbf{g}})$ becomes close to one,²⁷ then $\sigma_P(a^*)$ must in some cases (e.g., when a^* is the risk-dominated equilibrium of a 2×2 coordination game) be vanishingly small. We explore this point in detail (in two-player games) in Oyama and Tercieux [18].

6. Conclusion

Following Kajii and Morris [10], the present paper has investigated the question of ex ante robustness of action distributions of complete information games to a small amount of incomplete information. Contrary to previous work in this literature, we postulated that in our incomplete information perturbations, players may have heterogeneous prior beliefs. We demonstrated that dropping the common prior assumption (CPA) has far reaching consequences. Our first result shows that an action distribution of a generic complete information game is robust under non-common priors if and only if the game is dominance solvable and the action distribution assigns weight one on the unique action profile surviving iterative deletion of strictly dominated actions. This implies that the robustness test that allows for incomplete information perturbations without common prior is substantially stronger than the one with common prior as considered by Kajii and Morris [10].

Our approach in this paper is an ex ante one, where the outside analyst has no information about interim beliefs of the players and thus is concerned with the ex ante average behavior of the players. On the contrary, Weinstein and Yildiz [24] consider a similar robustness question with an *interim* approach where the analyst is given a hierarchy of interim beliefs and is concerned

²⁷ Observe that this is the case in the perturbation in the proof of Lemma 5.9.

with the behavior the players may have at this specific hierarchy. They show that an equilibrium is interim robust if and only if it is the unique action profile surviving iterative deletion of strictly dominated strategies. Importantly, by the result of Lipman [12], this characterization remains valid even when players are assumed to share a common prior in incomplete information perturbations. Our characterization, in contrast, does not hold under the CPA. This is due to the fact that, as the critical path result of Kajii and Morris [10] establishes, under the CPA the restrictions on prior beliefs also impose restrictions on interim (higher order) beliefs.

We also investigated the question of robustness of sets of action distributions. Our second result, which generalizes our result on point-valued robustness, shows that in generic games, a smallest robust set exists and coincides with the set of a posteriori equilibria. Given the result by Brandenburger and Dekel [4], this means that even if the analyst uses any refinement of rationalizability such as Nash equilibrium, in order to obtain robust predictions he cannot reject any outcome from the set of rationalizable outcomes when he takes into account the possibility that the players may not share a common prior in nearby incomplete information perturbations. For the set of robust predictions to be sharpened further, the analyst has to have more information about the actual situation, to impose some more restrictions on prior beliefs than the present paper assumes. As an example, we provided a measure of disagreement among the players' prior beliefs such that a bound on this measure may lead to a sharper robust prediction (Proposition 5.7).

Acknowledgments

Earlier versions of this paper were circulated under the title “Strategic Implications of (Non-)Common Priors II: Robustness of Equilibria”. We would like to thank Akihiko Matsui, Stephen Morris, Takashi Ui, Muhamet Yildiz, anonymous referees, and especially Atsushi Kajii for helpful comments and discussions. This paper has been presented at Harvard University, Hitotsubashi University, IHP Paris Game Theory Seminar, Kyoto University, MIT, Paris I-Panthéon Sorbonne, PSE, University of Toronto, the 21st Annual Congress of the European Economic Association in Vienna, the Twelfth Decentralization Conference in Osaka, and TARK XI in Brussels. D. Oyama acknowledges financial support by Grant-in-Aid for JSPS Fellows and from the Japan Economic Research Foundation. 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 and partly conducted while D. Oyama was visiting Paris School of Economics. Their hospitality is gratefully acknowledged.

Appendix A

A.1. Proof of Lemma 5.2

Proof. The proof mimics the proof of Lemma 3.4. Fix $\delta > 0$ and $\varepsilon > 0$. For each $k = 0, \dots, N - 1$, let $(\mu_i^k)_{i \in \mathcal{I}} \in (\Delta(A))^I$ be a strict k -subjective correlated equilibrium such that $\text{supp}(\mu_i^k) = W_i^{k+1}[\mathbf{g}] \times W_{-i}^k[\mathbf{g}]$.²⁸ By continuity of g_i 's, we can take for each $k = 0, \dots, N$, $(\tilde{\mu}_i^k)_{i \in \mathcal{I}} \in (\Delta(A))^I$ such that for each $i \in \mathcal{I}$, (1) $|\tilde{\mu}_i^k - \mu_i^k| < \delta/2$; (2) $\text{supp}(\tilde{\mu}_i^k) = A$; and (3) for

²⁸ Any strict k -subjective correlated equilibrium $(\mu_i^k)_{i \in \mathcal{I}}$ must satisfy $\text{supp}(\mu_i^k) \subset W_i^{k+1}[\mathbf{g}] \times W_{-i}^k[\mathbf{g}]$ for all $i \in \mathcal{I}$.

some $\eta_i > 0$, and all $a_i \in A_i$ with $\mu_i^k(a_i) > 0$, $\{a_i\} = br_i(\pi_i)$ for all $\pi_i \in \Delta(A_{-i})$ satisfying $|\pi_i(a_{-i}) - \tilde{\mu}_i^k(a_{-i}|a_i)| < \eta_i$ for all $a_{-i} \in W_{-i}^k[\mathbf{g}]$.

Set $\eta = \min_{i \in \mathcal{I}} \eta_i > 0$, and take an $r > 0$ such that $1/(1+r) \leq \eta$. Let $\tilde{\varepsilon} \in (0, 1)$, which will be taken to be small depending on ε and r . We now construct an $(\varepsilon, N - 1)$ -perturbation $\mathcal{U}^{\varepsilon, N-1}$ as follows. Let $\Omega = \mathcal{I} \times \{0, \dots, N + 1\} \times A$ and define $P_i \in \Delta(\Omega)$ for each $i \in \mathcal{I}$ and $k = 0, \dots, N - 1, N + 1$ by

$$P_i(i, k, a) = r\tilde{\varepsilon}(I - 1)\tilde{\mu}_i^k(a)$$

and

$$P_i(i, N, a) = (1 - \tilde{\varepsilon}(I - 1)(r(N + 1) + N + 2))\tilde{\mu}_i^N(a).$$

In addition, for each $i \in \mathcal{I}$ and $k = 0, \dots, N + 1$

$$P_i(j, k, a) = \tilde{\varepsilon}\tilde{\mu}_i^{k-1}(a)$$

for $j \neq i$ where by convention, $\tilde{\mu}_i^{-1} = \tilde{\mu}_i^0$.

Let each \mathcal{Q}_i consist of (i) the events

$$E_{ia_i}^0 = \{(j, 0, a_i, a_{-i}) \mid j \neq i, a_{-i} \in A_{-i}\},$$

(ii) all the events of the form

$$E_{ia_i}^k = \{(i, k - 1, a_i, a_{-i}), (j, k, a_i, a_{-i}) \mid j \neq i, a_{-i} \in A_{-i}\}$$

for each $k = 1, \dots, N + 1$, and (iii) the events

$$E_{ia_i}^{N+2} = \{(i, N + 1, a_i, a_{-i}) \mid a_{-i} \in A_{-i}\}.$$

Finally, define each $u_i : A \times \Omega \rightarrow \mathbb{R}$ by

$$u_i((a_i, a_{-i}), \omega) = \begin{cases} g_i(a_i, a_{-i}) & \text{if } \omega \notin \bigcup_{a'_i} E_{ia'_i}^0, \\ 1 & \text{if } \omega \in E_{ia_i}^0, \\ 0 & \text{if } \omega \in E_{ia'_i}^0 \text{ for } a'_i \neq a_i. \end{cases}$$

Verify that $\Omega_{\mathbf{g}} = \{(i, k, a) \mid k \geq 1, i \in \mathcal{I}\}$ and $\bigcap_{n=0}^{N-1} [K_*]^n(\Omega_{\mathbf{g}}) = \{(i, k, a) \mid k \geq N, i \in \mathcal{I} \text{ and } a \in A\}$. Simple algebra shows that for $\tilde{\varepsilon}$ small enough, $P_i(\bigcap_{n=0}^{N-1} [K_*]^n(\Omega_{\mathbf{g}})) \geq 1 - \varepsilon$ for all $i \in \mathcal{I}$.²⁹

Now let σ be any rationalizable strategy profile of $\mathcal{U}^{\varepsilon, N-1}$. We show that for all $k = 0, \dots, N$, all $i \in \mathcal{I}$ and all $a_i \in W_i^k[\mathbf{g}]$, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^k$. By construction, for all $i \in \mathcal{I}$, and $a_i \in W_i^0[\mathbf{g}]$, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^0$. Then suppose that for all $i \in \mathcal{I}$, and all $a_i \in W_i^k[\mathbf{g}]$, $\sigma_i(a_i|\omega) = 1$ for all $\omega \in E_{ia_i}^k$. Consider any $i \in \mathcal{I}$, $a_i \in W_i^{k+1}[\mathbf{g}]$ (i.e., $a_i \in A_i$ such that $\mu_i^k(a_i) > 0$) and $\omega \in E_{ia_i}^{k+1}$. By construction of the state space, we have for $k = 1, \dots, N - 1$

$$\begin{aligned} & P_i((i, k, a_i, a_{-i}) | \mathcal{Q}_i(\omega)) \\ &= \frac{P_i((i, k, a_i, a_{-i}))}{\sum_{a'_{-i}} P_i((i, k, a_i, a'_{-i})) + \sum_{a'_{-i}} \sum_{j \neq i} P_i((j, k + 1, a_i, a'_{-i}))} \\ &= \frac{r\tilde{\mu}_i^k(a_i, a_{-i})}{(r + 1)\sum_{a'_{-i}} \tilde{\mu}_i^k(a_i, a'_{-i})} = \frac{r}{r + 1} \tilde{\mu}_i^k(a_{-i}|a_i), \end{aligned}$$

²⁹ By convention $[K_*]^0(\Omega_{\mathbf{g}}) = \Omega_{\mathbf{g}}$.

so that

$$\begin{aligned} |P_i((i, k, a_i, a_{-i})|Q_i(\omega)) - \tilde{\mu}_i^k(a_{-i}|a_i)| &\leq \frac{1}{r+1} \tilde{\mu}_i^k(a_{-i}|a_i) \\ &\leq \frac{1}{r+1} \leq \eta, \end{aligned}$$

where the last inequality follows from the choice of an appropriate (large enough) r . For $k = N$, again, by the choice of an appropriate (small enough) $\tilde{\varepsilon}$, a same reasoning applies and we obtain

$$|P_i((i, N, a_i, a_{-i})|Q_i(\omega)) - \tilde{\mu}_i^N(a_{-i}|a_i)| \leq \eta.$$

Thus, by the induction hypothesis, the choice of η , as well as the condition (3) in the definition of $\tilde{\mu}_i^k$, we have $\sigma_i(a_i|\omega) = 1$.

Finally, for $\tilde{\varepsilon} > 0$ small enough, we have by construction $|\sigma_{P_i}(a) - \tilde{\mu}_i^N(a)| < \delta/2$ for all $a \in A$ and so $|\sigma_{P_i}(a) - \mu_i^N(a)| < \delta$ for all $a \in A$ by construction. \square

Remark A.1. Three remarks on the above proof are in order. First, the $(\varepsilon, N - 1)$ -perturbation constructed in the proof is not an (ε, N) -perturbation. Second, contrary to the construction in Lemma 3.4, this $(\varepsilon, N - 1)$ -perturbation need not have a unique rationalizable strategy profile. Finally, the proof is rather tricky compared to the proof of Lemma 3.4, in that the prior probability of each player will put almost all its mass only on a small number of states.

A.2. Proof of Proposition 5.5

Fix an information system $[\Omega, (P_i)_{i \in \mathcal{I}}, (Q_i)_{i \in \mathcal{I}}]$ and a simple event $E = \bigcap_{i \in \mathcal{I}} E_i$, each $E_i \in \mathcal{F}_i$. We use the same labeling as in Kajii and Morris [10, Lemma C]. Fix $K \geq 0$, and define inductively $\{E_1^k, \dots, E_I^k, E^k\}_{k=1}^{K+1}$ as follows: $E_i^1 = E_i$, and $E^k = \bigcap_{i \in \mathcal{I}} E_i^k$ and $E_i^{k+1} = B_i^{P_i}(E^k)$. By convention, let $E_i^0 = \Omega$. Then let $D_i^k = E_i^k \setminus E_i^{k+1}$ for $k = 0, 1, \dots, K$, and $D_i^{K+1} = E_i^{K+1}$. Observe that $\{D_i^k\}_{k=0}^{K+1}$ is a partition coarser than Q_i . Writing $\mathbf{n} = (n_1, \dots, n_I)$ for a typical element of $\{0, 1, \dots, K + 1\}^I$, we denote by $\min(\mathbf{n})$ the smallest number in $\{n_1, \dots, n_I\}$. Define $L(\mathbf{n}) = \bigcap_{i \in \mathcal{I}} D_i^{n_i}$ and $\pi_i(\mathbf{n}) = P_i(L(\mathbf{n}))$ for $\mathbf{n} \in \{0, 1, \dots, K + 1\}^I$. Note that for all $k = 0, 1, \dots, K + 1$ and $i \in \mathcal{I}$, $D_i^k = \bigcap_{\mathbf{n}: n_i=k} L(\mathbf{n})$ and $E^k = \bigcap_{\mathbf{n}: n_i \geq k} L(\mathbf{n})$.

Let $x_j(i, 0) = 0$ and

$$x_j(i, k) = \sum_{\mathbf{n}: n_i=\min(\mathbf{n}), 0 < \min(\mathbf{n}) \leq k} \pi_j(\mathbf{n})$$

for $k = 1, \dots, K$, and

$$y_j = \sum_{\mathbf{n}: n_j > 0, \min(\mathbf{n})=0} \pi_j(\mathbf{n}).$$

Lemma A.2.1. For all $i \in \mathcal{I}$ and $k \geq 1$,

$$x_i(i, k) \leq \frac{P_i}{1 - P_i} \sum_{h \neq i} x_i(h, k - 1) + \frac{P_i}{1 - P_i} y_i. \tag{A.1}$$

Proof. For all $i \in \mathcal{I}$ and $k \geq 1$,

$$\begin{aligned}
 x_i(i, k) &= \sum_{\ell=1}^k \sum_{\mathbf{n}: n_i=\ell, \min(\mathbf{n})=\ell} \pi_i(\mathbf{n}) \\
 &\leq \sum_{\ell=1}^k \frac{p_i}{1-p_i} \sum_{\mathbf{n}: n_i=\ell, \min(\mathbf{n})<\ell} \pi_i(\mathbf{n}) \\
 &\leq \frac{p_i}{1-p_i} \sum_{m=0}^{k-1} \sum_{\mathbf{n}: n_i>m, \min(\mathbf{n})=m} \pi_i(\mathbf{n}) \\
 &= \frac{p_i}{1-p_i} \sum_{m=1}^{k-1} \sum_{\mathbf{n}: n_i>m, \min(\mathbf{n})=m} \pi_i(\mathbf{n}) + \frac{p_i}{1-p_i} \sum_{\mathbf{n}: n_i>0, \min(\mathbf{n})=0} \pi_i(\mathbf{n}) \\
 &\leq \frac{p_i}{1-p_i} \sum_{h \neq i} \sum_{m=1}^{k-1} \sum_{\mathbf{n}: n_i>m, n_h=m, \min(\mathbf{n})=m} \pi_i(\mathbf{n}) + \frac{p_i}{1-p_i} \sum_{\mathbf{n}: n_i>0, \min(\mathbf{n})=0} \pi_i(\mathbf{n}) \\
 &\leq \frac{p_i}{1-p_i} \sum_{h \neq i} x_i(h, k-1) + \frac{p_i}{1-p_i} y_i,
 \end{aligned}$$

as claimed. \square

Lemma A.2.2. For all $j \in \mathcal{I}$ and $m \geq 0$,

$$\sum_{\mathbf{n}: n_j>m, \min(\mathbf{n})=m} \pi_j(\mathbf{n}) + r \sum_{i \neq j} \sum_{\mathbf{n}: n_i>m, \min(\mathbf{n})=m} \pi_i(\mathbf{n}) \leq r \sum_{i \neq j} \sum_{\mathbf{n}: \min(\mathbf{n})=m} \pi_i(\mathbf{n}). \tag{A.2}$$

Proof. Let J be a typical element of $2^I \setminus \{\emptyset\}$. For all $j \in \mathcal{I}$ and $m \geq 0$,

$$\begin{aligned}
 &\sum_{\mathbf{n}: n_j>m, \min(\mathbf{n})=m} \pi_j(\mathbf{n}) + r \sum_{i \neq j} \sum_{\mathbf{n}: n_i>m, \min(\mathbf{n})=m} \pi_i(\mathbf{n}) \\
 &= \sum_{J \neq \emptyset: j \notin J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_j(\mathbf{n}) \\
 &\quad + r \sum_{i \neq j} \sum_{J \neq \emptyset: i \notin J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \\
 &\leq \sum_{i \neq j} \sum_{J \neq \emptyset: j \notin J, i \in J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_j(\mathbf{n}) \\
 &\quad + r \sum_{i \neq j} \sum_{J \neq \emptyset: i \notin J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \\
 &\leq r \sum_{i \neq j} \sum_{J \neq \emptyset: j \notin J, i \in J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \\
 &\quad + r \sum_{i \neq j} \sum_{J \neq \emptyset: i \notin J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \\
 &= r \sum_{i \neq j} \left(\sum_{J \neq \emptyset: j \notin J, i \in J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \right)
 \end{aligned}$$

$$\begin{aligned}
 & + \sum_{J \neq \emptyset: i \notin J} \sum_{\mathbf{n}: \min(\mathbf{n})=m, \arg \min(\mathbf{n})=J} \pi_i(\mathbf{n}) \Big) \\
 & \leq r \sum_{i \neq j} \sum_{\mathbf{n}: \min(\mathbf{n})=m} \pi_i(\mathbf{n}),
 \end{aligned}$$

as claimed. \square

In the following, we consider the case where $p_i = p$ for all $i \in \mathcal{I}$.

Let $x_j(k) = \sum_{i \in \mathcal{I}} x_j(i, k)$ and

$$\mathbf{x}(k) = (x_1(k), x_2(k), \dots, x_I(k))'.$$

Lemma A.2.3. *If $p_i = p$ for all $i \in \mathcal{I}$, then*

$$\mathbf{x}(k) \leq (r\mathbf{R})\mathbf{x}(k-1) + \varepsilon(r\mathbf{R})\mathbf{1}'.$$

Proof. By Lemma A.2.1,

$$\begin{aligned}
 x_j(k) &= x_j(j, k) + \sum_{i \neq j} x_j(i, k) \\
 &\leq x_j(j, k) + \sum_{i \neq j} r x_i(i, k) \\
 &\leq \frac{p_j}{1-p_j} \sum_{h \neq j} x_j(h, k-1) + \frac{p_j}{1-p_j} y_j \\
 &\quad + \sum_{i \neq j} \left(\frac{r p_i}{1-p_i} \sum_{h \neq i} x_i(h, k-1) + \frac{r p_i}{1-p_i} y_i \right) \\
 &\leq \frac{r p_j}{1-p_j} \sum_{h \neq j} x_h(h, k-1) + \frac{p_j}{1-p_j} y_j \\
 &\quad + \sum_{i \neq j} \left(\frac{r p_i}{1-p_i} \sum_{h \neq i} x_i(h, k-1) + \frac{r p_i}{1-p_i} y_i \right) \\
 &= \frac{r p}{1-p} \sum_{i \neq j} \sum_{h \in \mathcal{I}} x_i(h, k-1) + \frac{p}{1-p} \left(y_j + r \sum_{i \neq j} y_i \right) \\
 &= \frac{r p}{1-p} \sum_{i \neq j} x_i(k-1) + \frac{p}{1-p} \left(y_j + r \sum_{i \neq j} y_i \right),
 \end{aligned}$$

where by Lemma A.2.2,

$$\begin{aligned}
 y_j + r \sum_{i \neq j} y_i &= \sum_{\mathbf{n}: n_j > 0, \min(\mathbf{n})=0} \pi_j(\mathbf{n}) + r \sum_{i \neq j} \sum_{\mathbf{n}: n_i > 0, \min(\mathbf{n})=0} \pi_i(\mathbf{n}) \\
 &\leq r \sum_{i \neq j} \sum_{\mathbf{n}: \min(\mathbf{n})=0} \pi_i(\mathbf{n}) \\
 &\leq (I-1)r\varepsilon.
 \end{aligned}$$

Hence, we have

$$x_j(k) \leq \frac{rp}{1-p} \sum_{i \neq j} x_i(k-1) + (I-1) \frac{rp}{1-p} \varepsilon$$

for all $j \in \mathcal{I}$, or

$$\mathbf{x}(k) \leq (r\mathbf{R})\mathbf{x}(k-1) + \varepsilon(r\mathbf{R})\mathbf{1}'$$

as claimed. \square

Proposition A.2.4. *Let $P_i(E^c) \leq \varepsilon$. If $p_i = p$ for all $i \in \mathcal{I}$, then*

$$1 - P_j([B_*^p]^K(E)) \leq \varepsilon[(\mathbf{I} + r\mathbf{R} + \dots + (r\mathbf{R})^K)\mathbf{1}']_j$$

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

Proof. By Lemma A.2.3,

$$\begin{aligned} \mathbf{x}(K) &\leq (r\mathbf{R})\mathbf{x}(K-1) + \varepsilon(r\mathbf{R})\mathbf{1}' \\ &\leq (r\mathbf{R})((r\mathbf{R})\mathbf{x}(K-2) + \varepsilon(r\mathbf{R})\mathbf{1}') + \varepsilon(r\mathbf{R})\mathbf{1}' \\ &= (r\mathbf{R})^2\mathbf{x}(K-2) + \varepsilon(r\mathbf{R} + (r\mathbf{R})^2)\mathbf{1}' \\ &\vdots \\ &\leq (r\mathbf{R})^K\mathbf{x}(0) + \varepsilon(r\mathbf{R} + (r\mathbf{R})^2 + \dots + (r\mathbf{R})^K)\mathbf{1}' \\ &= \varepsilon(r\mathbf{R} + (r\mathbf{R})^2 + \dots + (r\mathbf{R})^K)\mathbf{1}'. \end{aligned}$$

Hence, we have

$$\begin{aligned} 1 - P_j([B_*^p]^K(E)) &= \sum_{\mathbf{n}: \min(\mathbf{n})=0} \pi_j(\mathbf{n}) + \sum_{\mathbf{n}: 0 < \min(\mathbf{n}) \leq K} \pi_j(\mathbf{n}) \\ &\leq \varepsilon + x_j(K) \\ &\leq \varepsilon + \varepsilon[(r\mathbf{R} + (r\mathbf{R})^2 + \dots + (r\mathbf{R})^K)\mathbf{1}']_j \\ &= \varepsilon[(\mathbf{I} + r\mathbf{R} + \dots + (r\mathbf{R})^K)\mathbf{1}']_j, \end{aligned}$$

as claimed. \square

Theorem A.2.5. *Let $P_i(E^c) \leq \varepsilon$. If $\{1 + r(I-1)\}p < 1$, then*

$$1 - P_j([B_*^p]^\infty(E)) \leq \frac{1-p}{1 - \{1 + r(I-1)\}p} \varepsilon$$

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

Proof. Observe first that

$$(r\mathbf{R})^k\mathbf{1}' = \left\{ (I-1) \frac{rp}{1-p} \right\}^k \mathbf{1}'.$$

If $\{1 + r(I-1)\}p < 1$, then

$$[(\mathbf{I} + r\mathbf{R} + \dots + (r\mathbf{R})^K)\mathbf{1}']_j = \sum_{k=0}^K \left\{ (I-1) \frac{rp}{1-p} \right\}^k$$

converges as $K \rightarrow \infty$ to

$$\frac{1}{1 - (I - 1)\frac{rp}{1-p}} = \frac{1 - p}{1 - \{1 + r(I - 1)\}p}.$$

Hence, by Proposition A.2.4 we have the desired inequality. \square

References

- [1] R.J. Aumann, Subjectivity and correlation in randomized strategies, *J. Math. Econ.* 1 (1974) 67–96.
- [2] R.J. Aumann, Correlated equilibrium as an expression of Bayesian rationality, *Econometrica* 55 (1987) 1–18.
- [3] P. Battigalli, M. Siniscalchi, Rationalization and incomplete information, *Adv. Theoretical Econ.* 3 (3) (2003) 1–44.
- [4] A. Brandenburger, E. Dekel, Rationalizability and correlated equilibria, *Econometrica* 55 (1987) 1391–1402.
- [5] A. Brandenburger, E. Dekel, Hierarchies of beliefs and common knowledge, *J. Econ. Theory* 59 (1993) 189–198.
- [6] E. Dekel, D. Fudenberg, D. Levine, Learning to play Bayesian games, *Games Econ. Behav.* 46 (2004) 282–303.
- [7] E. Dekel, D. Fudenberg, S. Morris, Interim correlated rationalizability, *Theoretical Econ.* 2 (2007) 15–40.
- [8] D. Fudenberg, D. Kreps, D. Levine, On the robustness of equilibrium refinements, *J. Econ. Theory* 44 (1988) 354–380.
- [9] J.C. Harsanyi, Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points, *Int. J. Game Theory* 2 (1973) 1–23.
- [10] A. Kajii, S. Morris, The robustness of equilibria to incomplete information, *Econometrica* 65 (1997) 1283–1309.
- [11] E. Kohlberg, J.-F. Mertens, On the strategic stability of equilibria, *Econometrica* 54 (1986) 1003–1037.
- [12] B.L. Lipman, Finite order implications of common priors, *Econometrica* 71 (2003) 1255–1267.
- [13] B.L. Lipman, Finite order implications of common priors in infinite models, *J. Math. Econ.* (2009), forthcoming.
- [14] J.-F. Mertens, S. Zamir, Formulation of Bayesian analysis for games with incomplete information, *Int. J. Game Theory* 14 (1985) 1–29.
- [15] D. Monderer, D. Samet, Approximating common knowledge with common beliefs, *Games Econ. Behav.* 1 (1989) 170–190.
- [16] S. Morris, R. Rob, H.S. Shin, p -Dominance and belief potential, *Econometrica* 63 (1995) 145–157.
- [17] S. Morris, T. Ui, Generalized potentials and robust sets of equilibria, *J. Econ. Theory* 124 (2005) 45–78.
- [18] D. Oyama, O. Tercieux, On the strategic impact of an event under non-common priors, mimeo, 2005.
- [19] D. Oyama, O. Tercieux, Iterated potential and robustness of equilibria, *J. Econ. Theory* 144 (2009) 1726–1769.
- [20] A. Rubinstein, The electronic mail game: Strategic behavior under ‘almost common knowledge’, *Amer. Econ. Rev.* 79 (1989) 385–391.
- [21] R.S. Simon, Games of incomplete information, ergodic theory, and the measurability of equilibria, *Israel J. Math.* 138 (2003) 73–92.
- [22] T. Ui, Robust equilibria of potential games, *Econometrica* 69 (2001) 1373–1380.
- [23] J. Weinstein, M. Yildiz, Finite-order implications of any equilibrium, MIT Department of Economics, Working Paper 04-06, 2004.
- [24] J. Weinstein, M. Yildiz, A structure theorem for rationalizability with application to robust predictions of refinements, *Econometrica* 75 (2007) 365–400.