APPROXIMATE ROBUSTNESS OF EQUILIBRIUM TO INCOMPLETE INFORMATION

Ori Haimanko and Atsushi Kajii

Discussion Paper No. 12-09

September 2012

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

Approximate Robustness of Equilibrium to Incomplete Information*

Ori Haimanko
Department of Economics, Ben-Gurion University

Atsushi Kajii
KIER, Kyoto University

August 31, 2012

Abstract

We relax the Kajii and Morris (1997a) notion of equilibrium robustness by allowing approximate equilibria in close incomplete information games. The new notion is termed "approximate robustness". The approximately robust equilibrium correspondence turns out to be upper hemicontinuous, unlike the (exactly) robust equilibrium correspondence. As a corollary of the upper hemicontinuity, it is shown that approximately robust equilibria exist in all two-player zero-sum games and all two-player two-strategy games, whereas (exactly) robust equilibria may fail to exist for some games in these categories.

JEL Classification Number: C72.

Keywords: incomplete information, robustness, Bayesian Nash equilibrium, ε-equilibrium, upper hemicontinuity, zero-sum games.

1 Introduction

Kajii and Morris (1997a) – henceforth KM – proposed a refinement of Nash equilibrium, based on the idea that an equilibrium should not change much if the information in a game becomes incomplete to a certain degree. More

*An earlier version of this paper was entitled "On Continuity of Robust Equilibria". Kajii acknowledges financial support from JSPS Grant-in-Aid for Scientific Research No.(S)20223001.

precisely, given a complete information game g, a game with incomplete information is considered to be "close" to g if the sets of players and actions are the same as in g, and, with high probability, each player knows that his payoffs are given by g (though there need not be common or approximate common knowledge of payoffs). A Nash equilibrium of g is said to be *robust to incomplete information* if every incomplete information game sufficiently close to g possesses a Bayesian-Nash equilibrium such that both equilibria induce similar distributions over actions.

KM motivated their concept of robustness by pointing out that an analyst who wishes to model some strategic environment as a complete information game (that describes the environment correctly with high probability) is unlikely to be aware of the fine details of the true information structure, which the players know and take into account in their strategic decisions. "If it is guaranteed that the analyst's prediction based on the complete information game is not qualitatively different from some equilibrium of the real incomplete information game being played, then the analyst will be justified in ignoring subtle informational complications." (KM, p. 1283)

We echo this motivation, and shall keep the above notion of "closeness" to g. We shall however relax the assumption that equilibrium behavior in close incomplete information games is exact, by allowing approximate Bayesian Nash equilibria. Given $\varepsilon \geq 0$, we say that a Nash equilibrium of g is *ε-robust to incomplete information* if every incomplete information game sufficiently close to g possesses a Bayesian-Nash (interim) *ε-equilibrium* such that both equilibria induce similar action distributions.

The concept of robustness of KM is obviously identical to our notion when we take $\varepsilon = 0$, i.e., our 0-robustness is just the KM-robustness. For $\varepsilon > 0$, the notion of ε-robustness is less demanding. However, ε-robustness may also imply implausible behavior in nearby incomplete information games, where players may consistently ε-deviate from their best responses, no matter how close the incomplete information games are to g. This is what the following definition is set to rule out. We say that a Nash equilibrium is *approximately robust to incomplete information* if it is ε-robust for *any* $\varepsilon > 0$.[1]

Our notion of approximate robustness constitutes a mild and natural extension of KM-robustness. Unlike KM, we do allow players to make small mistakes – slight deviations from their best responses – in incomplete information games that are close to g, in approximating the behavior in a

[1] The notion of robustness based on approximate best responses has been proposed by Chen and Xiong (2011) in the context of selecting rationalizable actions in incomplete information games via perturbing higher order beliefs. However, their notion acts as a *refinement* of the Weinsten and Yildiz (2007) selection procedure in which the best responses are exact, while in our case approximate robustness *extends* KM-robustness.

Nash equilibrium of g. But, to keep the spirit of exactness set forth in KM, the definition of approximate robustness requires that these mistakes become vanishingly small as the incomplete information games "converge" to g. Thus, the analyst in the KM story will still do well by choosing an approximately robust equilibrium (henceforth, ARE) in the complete information game g, as this prediction is quite justifiable – in the real incomplete information game close to g players do not need to depart from rationality beyond some practically negligible bound, if at all, to arrive at the predicted action distribution.

Part of the conceptual appeal of approximate robustness lies in the fact that the set of ARE is well behaved. The main result of this paper, Theorem 2, shows that the correspondence which maps each complete information game to the (possibly empty) set of its ARE is upper hemicontinuous. Since the robustness embodies the continuity with respect to information, it goes without saying that the additional, *implied*, aspect of continuity, with respect to the base complete information game, provides an important support for the notion of approximate robustness. In contrast, the KM-robust equilibrium correspondence is not upper hemicontinuous, as we will show in Section 5.

The upper hemicontinuity of the ARE correspondence has an immediate application. It implies that, if the set of ARE is non-empty for a class of games, then the set is non-empty for the closure of that class. We use this fact to show that every two-player *zero-sum* game possesses an ARE (see Corollary 3), as does every two-player two-action game (see Corollary 4). These claims do not hold in general for KM-robust equilibria (henceforth, KM-RE), however. We will show in Section 5 that a zero-sum game may not have a KM-RE unless there is a unique saddle point (which is then a KM-RE by Proposition 3.2 in KM).

Since there are games that possess an ARE but not a KM-RE, a fortiori approximate robustness is strictly weaker than KM-robustness, despite an a priori similarity of the two notions. This weakness stresses another useful aspect of approximate robustness – adopting it as an alternative to the KM notion has the effect of strictly extending the domain of games in which a robust equilibrium exists.[2]

Our paper is organized as follows. The basic notations pertaining to games of complete and incomplete information are presented in Section 2.

[2]Such an extension is a much needed step, as there are games without a KM-RE, and there are only limited KM-RE existence results (see, e.g., KM, Ui (2001), Morris and Ui (2005)). The general scope of this extension remains an open question, however, and some games of interest are obviously left out. For instance, Oyama and Takahashi (2011) describe an open set of two-player supermodular games without a KM-RE, but non-existence of ARE in these games can also be established in a similar way.

Section 3 introduces our notions of ε-robustness and approximate robustness. Section 4 contains our main result on upper hemicontinuity of the ARE correspondence (Theorem 2), supplemented by Corollaries 3 and 4 that establish existence of ARE in two subclasses of two-player games. Finally, Section 5 considers an example of a 4×4 zero-sum game, which shows simultaneously that the existence of an ARE does not guarantee the existence of a KM-RE, and that the KM-RE correspondence is not upper hemicontinuous.

2 Preliminaries

2.1 Complete Information Games

We follow the notation of KM as close as possible for ease of comparison. Throughout the analysis we fix a finite set of *players* $\mathcal{I} = \{1, 2, ..., I\}$ and a finite set A_i of *actions* for each player $i \in \mathcal{I}$. Denote by $A := \times_{i\in\mathcal{I}} A_i$ the set of players' action profiles. We shall denote $\times_{j\neq i} A_j$ by A_{-i} and a generic element of A_{-i} by a_{-i}. Similar conventions will be used whenever they are clear from the context. A *complete information game* is given by an I-tuple $g = (g_i)_{i\in\mathcal{I}}$, where $g_i : A \to \mathbb{R}$ is the *payoff function* of player i for each $i \in \mathcal{I}$.

For a given finite set B, denote by $\Delta(B)$ the simplex of probability vectors on B, i.e.,

$$\Delta(B) \equiv \left\{ (s(b))_{b\in\mathcal{B}} \in \mathbb{R}_+^{\mathcal{B}} \mid \sum_{b\in B} s(b) = 1 \right\}.$$

An element of $\Delta(A_i)$ is referred to as a *mixed action* for player i and that of $\Delta(A)$ as an *action distribution*. The distance between two action distributions is measured by the sup norm: thus, for any $\mu, \mu' \in \Delta(A)$ we write

$$\|\mu - \mu'\| \equiv \max_{a\in\mathcal{A}} |\mu(a) - \mu'(a)|. \tag{1}$$

An *action distribution*, $\mu \in \Delta(A)$, is a *correlated equilibrium* of a game g if, for all $i \in \mathcal{I}$ and $a_i, a_i' \in A_i$,

$$\sum_{a_{-i}\in A_{-i}} g_i(a_i, a_{-i}) \mu(a_i, a_{-i}) \geq \sum_{a_{-i}\in A_{-i}} g_i(a_i', a_{-i}) \mu(a_i, a_{-i}).$$

An action distribution μ is a *Nash equilibrium* of g if it is a correlated equilibrium, and is a product distribution induced by a mixed action profile, i.e., for all $a \in A$,

$$\mu(a) = \underset{i\in\mathcal{I}}{\times} \mu_i(a_i), \tag{2}$$

where $\mu_i \in \Delta(A_i)$ is the marginal distribution of μ on A_i. Whenever convenient, a Nash equilibrium μ will be represented by the mixed action profile $(\mu_1, ..., \mu_I)$. Denote by $NE(g)$ the set of Nash equilibria of g.

2.2 Incomplete Information Games

In line with KM, we now extend the definition of a game to allow uncertainty and incomplete information.

The underlying uncertainty in an *incomplete information game* is described by a probability space (Ω, P), where Ω is a countable[3] *set of states of nature*, and P is a countably additive probability measure on Ω which is the *common prior belief* of the players about the actual state of nature. The *information* of player i is given by a (possibly infinite) partition Q_i of Ω. The payoffs to player i are determined by a state dependent payoff function, $u_i : A \times \Omega \to \mathbb{R}$. The incomplete information game with the above attributes will be denoted by $\mathcal{U} = \left\{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u_i\}_{i\in\mathcal{I}}\right\}$.

Given $\omega \in \Omega$, denote by $Q_i(\omega)$ the unique element of Q_i that contains ω; if ω is the actual state of nature, player i only knows that the realized state belongs to $Q_i(\omega)$. We will henceforth assume[4] that every information set of every player is possible, i.e., that $P(Q_i(\omega)) > 0$ for all $i \in \mathcal{I}$ and $\omega \in \Omega$. Under this assumption the conditional probability of state ω given information set $Q_i(\omega)$, written $P(\omega|Q_i(\omega))$, is well-defined by the rule $P(\omega|Q_i(\omega)) = \frac{P(\omega)}{P[Q_i(\omega)]}$.

A *(behavioral) strategy* of player i is a Q_i-measurable function $\sigma_i : \Omega \to \Delta(A_i)$; $\sigma_i(a_i|\omega)$ will denote the probability that player i chooses action a_i given ω. A *strategy profile* is an I-tuple $\sigma = (\sigma_i)_{i\in\mathcal{I}}$ where σ_i is a strategy of player i. We denote by $\sigma(a|\omega)$ the probability that action profile $a = (\cdots, a_i, \cdots)$ is chosen given ω under σ; i.e., $\sigma(a|\omega) = \prod_{i\in\mathcal{I}} \sigma_i(a_i|\omega)$. Write $\mathcal{X}_i$ for the set of strategies for player i, and $\mathcal{X}$ for the set of all the strategy profiles. Also denote by $\mathcal{X}_{-i}$ the set of strategy profiles of players other than player i, and write σ_{-i} for $(\sigma_j)_{j\neq i}$.

Abusing notation, we extend the domain of each u_i to mixed strategies and thus write $u_i(\sigma(\omega), \omega)$ for $\sum_{a\in A} u_i(a, \omega)\sigma(a|\omega)$. When $\omega \in \Omega$ occurs, the *interim payoff* of strategy profile σ to player i is given by the conditional expectation

$$U_i(\sigma|\omega) \equiv \sum_{\omega\in Q_i(\omega)} \sum_{a\in A} u_i(a, \omega)\sigma(a|\omega) P[\omega|Q_i(\omega)]. \tag{3}$$

[3]The countability assumption is made to avoid measure theoretic complications, just as in KM.

[4]This simplifying assumption is also made in KM.

Thus far our setup has been identical to that of KM. We now extend the scope of KM by considering approximate, and not just exact, equilibria in incomplete information games. For $\varepsilon \geq 0$ a strategy profile $\hat{\sigma} \in \mathcal{X}$ is an *(interim) Bayesian ε-Nash equilibrium* of $\mathcal{U}$ (henceforth, ε-*BE* for short) if, for every player i, for all $\sigma_i \in \mathcal{X}_i$ and for all $\omega \in \Omega$,

$$U_i(\hat{\sigma}|\omega) \geq U_i(\sigma_i, \hat{\sigma}_{-i}|\omega) - \varepsilon \tag{4}$$

Denote by $BE_\varepsilon(\mathcal{U})$ the set of all ε-BE of $\mathcal{U}$.

Notice that the slack ε is chosen uniformly across the states of nature, and so the notion of ε-*BE* is much stronger than what one might regard as an "ex ante" ε-equilibrium. By the principle of dynamic optimization, a 0-*BE* is just the standard Bayesian Nash equilibrium of $\mathcal{U}$ (*BE* for short).

An action distribution, $\hat{\mu} \in \Delta(A)$, is an *ε-BE equilibrium action distribution* of $\mathcal{U}$ if there exists a $\hat{\sigma} \in BE_\varepsilon(\mathcal{U})$ which induces $\hat{\mu}$; that is, $\hat{\mu}(a) = \sum_{\omega\in\Omega} \hat{\sigma}(a|\omega) P(\omega)$ for every $\alpha \in A$.

3 Approximate Robustness

Following KM, an incomplete information game $\mathcal{U}$ is deemed close to a complete information game g if the payoff structure under $\mathcal{U}$ is equal to g with high probability. Formally, for a given incomplete information game $\mathcal{U}$ we define for every $i \in \mathcal{I}$:

$$\Omega_i(\mathcal{U}, g) \equiv \{\omega : u_i(a, \omega') = g_i(a) \text{ for all } a \in A,\ \omega' \in Q_i(\omega)\}, \tag{5}$$

and set $\Omega(\mathcal{U}, g) \equiv \cap_i \Omega_i(\mathcal{U}, g)$. An incomplete information game $\mathcal{U}$ is said to be a *δ-elaboration* of a complete information game g if $P(\Omega(\mathcal{U}, g)) \geq 1 - \delta$.[5]

The following definition extends the KM notion of informational robustness in that it does not require the BE in elaborations to be exact:[6]

Definition 1 *Given a complete information game g and $\varepsilon \geq 0$, an action distribution μ is* ε-robust *to incomplete information in g (ε-*RE *for short),*

[5] Our δ-elaboration is a slight extention of the original notion of KM, who require that $P(\Omega(\mathcal{U}, g)) = 1 - \delta$. It is easy to see that both definitions of δ-elaboration would give rise to identical concepts of robustness (introduced in Definition 1 below). However, our notion streamlines several mathematical arguments in the sequel, as any δ-elaboration is also a δ'-elaboration in our sense, for any $\delta' > \delta$.

[6] One could define a weaker concept by restricting elaborations to *canonical* elaborations, as in Kajii and Morris (1997b) and Ui (2001). It will become clear that all the results and comments we report in this paper remain valid for the weaker notion. It is however an open question whether this is a *strictly* weaker notion.

if for any $\tau > 0$, there exists $\overline{\delta} > 0$ with the following property: any δ-elaboration $\mathcal{U}$ of g with $0 \leq \delta \leq \overline{\delta}$ possesses an ε-BE action distribution ν such that with $\|\mu - \nu\| \leq \tau$. An action distribution μ is approximately robust *to incomplete information if it is ε-robust for any $\varepsilon > 0$.*

Notice that if μ is ε-robust to incomplete information, it must be an ε-Nash equilibrium[7] of g (as follows from Definition 1 by considering the degenerate 0-elaboration of g with $|\Omega| = 1$). Thus an approximately robust action distribution is necessarily a Nash equilibrium of g, and we will refer to it as an *approximately robust equilibrium* (*ARE* for short) from now on. The set of ARE in the game g is denoted by $ARE\,(g)$.

As was said, our notions of ε- and approximate robustness extend the definition of robustness introduced in KM, that considers only exact (0-)BE equilibria in elaborations. Thus the Nash equilibria which are KM-robust are precisely the 0-robust action distributions. It follows from the definition that any KM-robust action distribution is approximately robust, and thus the set of KM-robust equilibria (KM-RE for short) is a subset of $ARE\,(g)$.

Given the conceptual closeness of requirements that the notions of KM-robustness and approximate robustness impose on equilibria, one might conjecture that the sets of KM-RE and ARE coincide. It turns out, however, that approximate robustness is a strictly weaker notion. It will be shown in Section 5, where we construct a game g in which there is no KM-RE, but $ARE\,(g) \neq \phi$.

4 Results

4.1 Upper Hemicontinuity of the ARE Correspondence

We shall show that the approximate robustness exhibits a desirable continuity property: the correspondence which maps each complete information game into the set of approximate robust equilibria is upper hemicontinuous. Interestingly enough, the analogous correspondence which maps a game into the set of its KM-robust equilibria is not upper hemicontinuous, as we elaborate in Section 5.

Formally, endow the set of all complete information games Γ with the metric d_Γ, given by

$$d_\Gamma\,(g, g') \equiv \max_{i \in \mathcal{I}} \max_{a \in \mathcal{A}} |g_i\,(a) - g'_i\,(a)|$$

[7] For $\varepsilon \geq 0$, a product distribution $\mu \in \Delta\,(A)$ is an ε-*Nash equilibrium* of g if for each $i \in \mathcal{I}$, μ_i is an ε-best response of i to μ_{-i}.

for every $g, g' \in \Gamma$. Note that the Nash equilibrium correspondence, $g \mapsto NE(g)$ is upper hemicontinuous with this metric. The next result shows that its sub-correspondence, the ARE correspondence $g \mapsto ARE(g)$, is also upper hemicontinuous on Γ:

Theorem 2 *Let $\{g^k\}_{k=1}^{\infty} \subset \Gamma$ and assume that, for each k, there exists $\mu^k \in ARE(g^k)$. If the limits $g \equiv \lim_{k\to\infty} g^k$ and $\mu \equiv \lim_{k\to\infty} \mu^k$ exist, then $\mu \in ARE(g)$.*

Proof. According to Definition 1, we must establish ε-robustness of μ for an arbitrarily chosen $\varepsilon > 0$. To this end, fix any $\tau > 0$. Since the games and the action distributions are convergent, there exists $k \geq 1$ such that the complete information game $\overline{g} \equiv g^k$ and its ARE $\overline{\mu} \equiv \mu^k$ satisfy

$$\|\bar{\mu} - \mu\| < \frac{\tau}{2},\ d_\Gamma(\overline{g}, g) < \frac{\varepsilon}{4}. \tag{6}$$

For any $0 \leq \delta$ and any δ-elaboration $\mathcal{U} = \{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u_i\}_{i\in\mathcal{I}}\}$ of g, denote by $\bar{\mathcal{U}} = \{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{\bar{u}_i\}_{i\in\mathcal{I}}\}$ the incomplete information game where, for every $a \in A$, $\omega \in \Omega$ and $i \in \mathcal{I}$,

$$\overline{u}_i(\omega, a) \equiv \begin{cases} \overline{g_i}(a), & \text{if } \omega \in \Omega_i(\mathcal{U}, g); \\ u_i(\omega, a), & \text{otherwise.} \end{cases} \tag{7}$$

That is, $\bar{\mathcal{U}}$ is obtained by replacing g_i with $\bar{g}_i$ whenever player i knows (in $\mathcal{U}$) that his payoff is given by g_i. Clearly, $\bar{\mathcal{U}}$ is a δ-elaboration of the game $\bar{g}$, and we shall call it a δ-elaboration of $\overline{g}$ *induced* by $\mathcal{U}$ for later reference.

Note that the second inequality in (6) implies via (7) and (3) that for any strategy profile σ,

$$\left|U_i(\sigma|\omega) - \bar{U}_i(\sigma|\omega)\right| < \frac{\varepsilon}{4} \tag{8}$$

for every $i \in \mathcal{I}$, at every $\omega \in \Omega$, where U_i and $\bar{U}_i$ are the interim payoffs of i in $\mathcal{U}$ and $\bar{\mathcal{U}}$, respectively, defined as in (3). Combining (8) with (4) in the definition of ε-BE, it is readily confirmed that every $\frac{\varepsilon}{2}$-BE strategy profile $\hat{\sigma}$ of $\bar{\mathcal{U}}$ is also an ε-BE of $\mathcal{U}$; that is,

$$BE_{\frac{\varepsilon}{2}}(\bar{\mathcal{U}}) \subset BE_\varepsilon(\mathcal{U}). \tag{9}$$

Recall that $\bar{\mu}$ is approximately robust in $\bar{g}$ by assumption, and in particular it is $\frac{\varepsilon}{2}$-robust in $\bar{g}$. So there exists $0 < \overline{\delta}$ such that for any $0 \leq \delta \leq \overline{\delta}$, any induced δ-elaboration $\bar{\mathcal{U}}$ of $\overline{g}$ possesses some $\widehat{\sigma}_{\bar{\mathcal{U}}} \in BE_{\frac{\varepsilon}{2}}(\bar{\mathcal{U}})$ which induces an action distribution $\widehat{\mu}_{\bar{\mathcal{U}}}$ such that

$$\|\widehat{\mu}_{\bar{\mathcal{U}}} - \overline{\mu}\| \leq \frac{\tau}{2}. \tag{10}$$

Now, for any δ-elaboration $\mathcal{U}$ of g with $0 \leq \delta \leq \bar{\delta}$, consider the induced elaboration $\bar{\mathcal{U}}$, and $\hat{\sigma}_{\bar{\mathcal{U}}} \in BE_{\frac{\varepsilon}{2}}(\bar{\mathcal{U}})$, $\widehat{\mu}_{\bar{\mathcal{U}}} \in \Delta(A)$ as above. By (9) we have $\widehat{\sigma}_{\bar{\mathcal{U}}} \in BE_{\varepsilon}(\mathcal{U})$, and by the first inequality in (6), and (10), also

$$\|\widehat{\mu}_{\bar{\mathcal{U}}} - \mu\| \leq \tau. \tag{11}$$

Since $\bar{\delta}$ as above can be found for any $\tau > 0$, we conclude that μ is ε-robust in the game g. And, since $\varepsilon > 0$ was chosen arbitrarily, μ is in fact approximately robust, as we claimed. ■

4.2 Existence of ARE in Subclasses of Two-Player Games

An immediate consequence of the upper hemicontinuity property established in Theorem 2 is that if the set of ARE is non-empty for a class of games, then the set is non-empty for the closure of the class. Here we apply this observation to show the existence of ARE in all two-player zero-sum games, and all two-player two-action games.

Corollary 3 *If g is a two-player zero-sum game, i.e., $I = 2$ and $g_2 = -g_1$, then $ARE(g) \neq \phi$.*

Proof. By Bohnenblust et al (1950), a complete information two-player zero-sum game has a unique pair of optimal mixed actions (and thus a unique Nash equilibrium) for a *generic*[8] payoff matrix of player 1. By Proposition 3.2 in KM and the discussion following it, this unique Nash equilibrium is a KM-RE.

Thus, given a zero-sum game g, there exists a sequence $\{g^k\}_{k=1}^{\infty}$ of zero-sum games such that $\lim_{k\to\infty} g^k = g$, and, for each $k \geq 1$, the game g^k has a KM-RE (which is in particular an ARE). A limit point of these ARE must belong to $ARE(g)$ by Theorem 2, and thus we have $ARE(g) \neq \phi$. ■

Corollary 4 *If g is a two-player two-action game, i.e., $I = |A_1| = |A_2| = 2$, then $ARE(g) \neq \phi$.*

Proof. The sufficient conditions of Propositions 3.2 and 5.3, and Corollary 5.6 in KM imply the existence of KM-RE in a *generic* 2×2 game (KM, p. 1301). As in the proof of Corollary 3, the application of Theorem 2 shows that $ARE(g) \neq \phi$ in *any* 2×2 game g. ■

[8] I.e., the claim holds for all payoff matrices in some dense and open (w.r.t. the Euclidean topolgy) subset of the space of all real-valued $|A^1| \times |A^2|$ matrices.

5 Approximate Robustness is weaker than KM-robustness

We now present an example of a two-player zero-sum game g that does not possess a KM-RE. Since ARE exist in all two-player zero-sum games by Corollary 3, this example demonstrates that the notion of approximate robustness is strictly weaker than the notion of KM-robustness. Moreover, as at least some Nash equilibria in g are limit points of KM-RE in nearby games (by the proof of Corollary 3), the non-existence of KM-RE also implies that the KM-RE correspondence is not upper hemicontinuous. This stands in contrast to the upper hemicontinuity of the ARE correspondence $g \mapsto ARE\left(g\right)$ that was established in Theorem 2.

Consider a zero-sum two-player game g, in which both players have four actions, and the payoffs of player 1 are given by the following matrix (where an action of player 1 (resp., 2) is represented by a choice of row (resp., column)):

$$\begin{array}{ccccc} & \mathbf{c}^1 & \mathbf{c}^2 & \mathbf{c}^3 & \mathbf{c}^4 \\ \mathbf{r}^1 & 1 & -1 & 0 & 0 \\ \mathbf{r}^2 & -1 & 1 & 0 & 0 \\ \mathbf{r}^3 & 0 & 0 & 1 & -1 \\ \mathbf{r}^4 & 0 & 0 & -1 & 1 \end{array}. \tag{12}$$

Note that, if players' choices are confined to either the first two rows/columns, or the last two rows/columns, then they play the matching pennies game. Any strategy which selects the first two rows/columns with equal probability, and the last two rows/columns with equal probability is a mixed equilibrium strategy.

We shall verify below that this game has no KM-RE. We do so by constructing two elaboration sequences, both of which approach the game g, such that each elaboration has a *unique* BE (we will establish this by employing the standard contagion argument).[9] But it will then become clear that the (uniquely determined) action distributions induced by the BE converge to different limits in the two sequences, which implies that no action distribution can be a KM-RE.

For the first sequence, fix $0 < \delta < 1$. In what follows we describe a 2δ-elaboration $\mathcal{U}_\delta = \left\{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u_i\}_{i\in\mathcal{I}}\right\}$ of game g. Let

$$\Omega = \{(k,k) \mid k \in \mathbb{Z}_+\} \cup \{(k+1,k) \mid k \in \mathbb{Z}_+\},$$

[9]By construction, our elaborations of g will also be *canonical* in the sense of Kajii and Morris (1997b) and Ui (2001), as every player i will have a strictly dominant pure action outside the set $\Omega_i\left(\mathcal{U}, g\right)$.

and assume that each player i can discern only the i^{th} coordinate in each state $(t_1, t_2) \in \Omega$, i.e., that

$$Q_1\left((t_1, t_2)\right) = \{t_1\} \times \{\max(t_1 - 1, 0), t_1\}$$

and

$$Q_2\left((t_1, t_2)\right) = \{t_2, t_2 + 1\} \times \{t_2\}.$$

Furthermore, let

$$P\left(\{(k, k)\}\right) = \delta\left(1 - \delta\right)^{2k} \text{ and } P\left(\{(k + 1, k)\}\right) = \delta\left(1 - \delta\right)^{2k+1}$$

for all $k \geq 0$. The following table illustrates the information structure and the prior.

	$t_2 = 0$	$t_2 = 1$	$t_2 = 2$	...	$t_2 = k - 1$	$t_2 = k$	...
$t_1 = 0$	δ						
$t_1 = 1$	$\delta\left(1 - \delta\right)$	$\delta\left(1 - \delta\right)^2$					
$t_1 = 2$		$\delta\left(1 - \delta\right)^3$	$\delta\left(1 - \delta\right)^4$				
$\vdots$			$\vdots$	$\ddots$			.
$t_1 = k$					$\delta\left(1 - \delta\right)^{2k-1}$	$\delta\left(1 - \delta\right)^{2k}$	
$t_1 = k + 1$						$\delta\left(1 - \delta\right)^{2k+1}$	...
$\vdots$							$\ddots$

The state-dependent payoff functions $(u_i)_{i=1,2}$ are determined by the following rule: (i) if $t_1 > 0$, then the pure action payoffs are given by the game g in (12); (ii) if $t_1 = 0$, then the pure action payoffs are given by the zero-sum game $\tilde{g}$ represented by the following payoff matrix:[10]

$$\begin{array}{ccccc} & \mathbf{c}^1 & \mathbf{c}^2 & \mathbf{c}^3 & \mathbf{c}^4 \\ \mathbf{r}^1 & 0 & 9 & 9 & 9 \\ \mathbf{r}^2 & 1 & 10 & 10 & 10 \\ \mathbf{r}^3 & 0 & 9 & 9 & 9 \\ \mathbf{r}^4 & 0 & 9 & 9 & 9 \end{array}. \tag{13}$$

[10]Note that with our specification of payoffs, the incomplete information game $\mathcal{U}_\delta$ is zero-sum in *every* state of nature. Thus, our method will show that the game g has no KM-robust equilibrium even under the restriction that all nearby incomplete information elaborations must be of zero-sum nature. However, the same proof would work if we assumed that, for *both* $i = 1, 2$, player i's payoffs are given by his payoffs in $\tilde{g}$ when $t_i = 0$, and by his payoffs in g when $t_i > 0$. Thus, the elaborations we use can be given another special feature (in addition to being canonical in the sense of Kajii and Morris (1997b) and Ui (2001), as every player i is "committed" to his strictly dominant action given $t_i = 0$) – it can be assumed that each player i is "commited" to some strategy precisely when he know that his payoffs are not given by g, i.e., when $\Omega_i\left(\mathcal{U}_\delta, g\right)$ does not occur.

Claim 5 *For every* $0 < \delta < 1$, *elaboration* $\mathcal{U}_\delta$ *has a* unique *BE* $(\widehat{\sigma}_1, \widehat{\sigma}_2)$, *where*

$$\widehat{\sigma}_1\left(\mathbf{r}^2 |\, (k, \max(k-1,0))\right) = \widehat{\sigma}_1\left(\mathbf{r}^2 |\, (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^1 |\, (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^1 |\, (k+1,k)\right) = 1 \tag{14}$$

if $k \geq 0$ *is even, and*

$$\widehat{\sigma}_1\left(\mathbf{r}^1 |\, (k, k-1)\right) = \widehat{\sigma}_1\left(\mathbf{r}^1 |\, (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^2 |\, (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^2 |\, (k+1,k)\right) = 1 \tag{15}$$

if $k \geq 0$ *is odd.*

Proof. Consider a BE $(\widehat{\sigma}_1, \widehat{\sigma}_2)$ of $\mathcal{U}_\delta$. Note that pure action $\mathbf{r}^2$ is strictly dominant for player 1 conditional on $t_1 = 0$, and also that $\mathbf{c}^1$ is strictly dominant for player 2 conditional on $t_2 = 0$ (since the conditional probability that player 2 assigns at $t_2 = 0$ to the event that $\tilde{g}$ in (13) is played, equals $\frac{\delta}{\delta+\delta(1-\delta)} = \frac{1}{2-\delta} > \frac{1}{2}$). Thus

$$\widehat{\sigma}_2\left(\mathbf{c}^1 |\, (0,0)\right) = \widehat{\sigma}_2\left(\mathbf{c}^1 |\, (1,0)\right) = 1, \tag{16}$$

and

$$\widehat{\sigma}_1\left(\mathbf{r}^2 |(0,0)\right) = 1. \tag{17}$$

The relations (16) and (17) thus establish (14) for $k = 0$.

Next, conditional on $t_1 = 1$, player 1 knows that the game is given by g. He believes that $t_2 = 0$ with probability $\frac{\delta(1-\delta)}{\delta(1-\delta)+\delta(1-\delta)^2} = \frac{1}{2-\delta}$ and that $t_2 = 1$ with the complementary probability $\frac{1-\delta}{2-\delta}$. Taking into account that $\mathbf{c}^1$ is played by player 2 at $t_2 = 0$ (as shown in (16)), the conditional expected payoff of player 1 at $t_1 = 1$ is given by the following matrix, where the rows correspond to the possible actions of 1 given $t_1 = 1$, and the columns correspond to the possible actions of 2 given $t_2 = 1$:

$$\begin{array}{ccccc} & \mathbf{c}^1 & \mathbf{c}^2 & \mathbf{c}^3 & \mathbf{c}^4 \\ \mathbf{r}^1 & 1 & \frac{\delta}{2-\delta} & \frac{1}{2-\delta} & \frac{1}{2-\delta} \\ \mathbf{r}^2 & -1 & -\frac{\delta}{2-\delta} & -\frac{1}{2-\delta} & -\frac{1}{2-\delta} \\ \mathbf{r}^3 & 0 & 0 & \frac{1-\delta}{2-\delta} & -\frac{1-\delta}{2-\delta} \\ \mathbf{r}^4 & 0 & 0 & -\frac{1-\delta}{2-\delta} & \frac{1-\delta}{2-\delta} \end{array}. \tag{18}$$

So no matter what player 2 plays at $t_2 = 1$, action $\mathbf{r}^1$ is strictly dominant for player 1 given $t_1 = 1$, and hence it must be played in any BE. We have thus shown that

$$\widehat{\sigma}_1\left(\mathbf{r}^1 |\, (1,0)\right) = \widehat{\sigma}_1\left(\mathbf{r}^1 |\, (1,1)\right)) = 1. \tag{19}$$

Similarly at $t_2 = 1$, using the fact that player 1's BE action at $t_1 = 1$ is $\mathbf{r}^1$ as shown in (19), and that player 2 attributes to $t_1 = 1$ probability $\frac{\delta(1-\delta)^2}{\delta(1-\delta)^2+\delta(1-\delta)^3} = \frac{1}{2-\delta}$, it can be shown that

$$\widehat{\sigma}_2\left(\mathbf{c}^1 \middle| (1,1)\right) = \widehat{\sigma}_2\left(\mathbf{c}^1 \middle| (2,1)\right) = 1. \tag{20}$$

The relations (19) and (20) therefore establish (15) for $k = 1$.

The argument can be done iteratively to obtain (14) and (15) for all $k > 1$.

■

Next, consider another 2δ-elaboration $\mathcal{U}'_\delta = \left\{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u'_i\}_{i\in\mathcal{I}}\right\}$, which is identical to $\mathcal{U}_\delta$ except for the payoff functions $\{u'_i\}_{i\in\{1,2\}}$ given as follows: (i) if $t_1 > 0$, pure action payoffs are given by the game g in (12); (ii) if $t_1 = 0$, pure action payoffs are given by the zero-sum game g' represented by the following payoff matrix:

	$\mathbf{c}^1$	$\mathbf{c}^2$	$\mathbf{c}^3$	$\mathbf{c}^4$
$\mathbf{r}^1$	9	9	9	0
$\mathbf{r}^2$	9	9	9	0
$\mathbf{r}^3$	10	10	10	1
$\mathbf{r}^4$	9	9	9	0

Then the following result can be established using arguments symmetric to those in the proof of Claim 5:

Claim 6 *For every* $0 < \delta < 1$, *elaboration* $\mathcal{U}'_\delta$ *has a* unique *BE* $(\widehat{\sigma}_1, \widehat{\sigma}_2)$, *where*

$$\widehat{\sigma}_1\left(\mathbf{r}^3 \middle| (k, \max(k-1)), 0)\right) = \widehat{\sigma}_1\left(\mathbf{r}^3 \middle| (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^4 \middle| (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^4 \middle| (k+1,k)\right) = 1$$

if $k \geq 0$ *is even, and*

$$\widehat{\sigma}_1\left(\mathbf{r}^4 \middle| (k,k-1)\right) = \widehat{\sigma}_1\left(\mathbf{r}^4 \middle| (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^3 \middle| (k,k)\right) = \widehat{\sigma}_2\left(\mathbf{c}^3 \middle| (k+1,k)\right) = 1$$

if $k \geq 0$ *is odd.*

It follows from the description of BE in Claim 5 that, when $\delta \to 0$, the (uniquely determined) BE action distribution in $\mathcal{U}_\delta$ converges to $\mu \in \Delta(A)$, which is the uniform distribution on the set $\{\mathbf{r}^1, \mathbf{r}^2\} \times \{\mathbf{c}^1, \mathbf{c}^2\} \subset A$. Similarly, from Claim 6, the unique BE action distribution in $\mathcal{U}'_\delta$ converges to $\mu' \in \Delta(A)$, which is the uniform distribution on the set $\{\mathbf{r}^3, \mathbf{r}^4\} \times \{\mathbf{c}^3, \mathbf{c}^4\} \subset A$. The limits are therefore distinct (in fact, supported on disjoint subsets of A), as we have asserted, confirming that there is no KM-RE in the game g.

To complete the discussion, we demonstrate that the game g does possess ARE, as guaranteed by Corollary 3. In fact, the game has multiple ARE; for instance, the following three equilibria of g are ARE:

$$(\mu_1^*, \mu_2^*) = \left((\frac{1}{2}, \frac{1}{2}, 0, 0), (\frac{1}{2}, \frac{1}{2}, 0, 0) \right), \tag{21}$$

$$(\mu_1^{**}, \mu_2^{**}) = \left((0, 0, \frac{1}{2}, \frac{1}{2}), (0, 0, \frac{1}{2}, \frac{1}{2}) \right), \tag{22}$$

$$(\mu_1^{***}, \mu_2^{***}) = \left((\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}), (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}) \right). \tag{23}$$

To see that (μ_1^*, μ_2^*) is an ARE, notice that, for every $\varepsilon > 0$, it is the *unique* Nash equilibrium in the zero-sum game $g^*(\varepsilon)$ where the payoffs of player 1 are given by the matrix

	$\mathbf{c}^1$	$\mathbf{c}^2$	$\mathbf{c}^3$	$\mathbf{c}^4$
$\mathbf{r}^1$	1	-1	ε	ε
$\mathbf{r}^2$	-1	1	ε	ε
$\mathbf{r}^3$	$-\varepsilon$	$-\varepsilon$	1	-1
$\mathbf{r}^4$	$-\varepsilon$	$-\varepsilon$	-1	1

Thus (μ_1^*, μ_2^*) is in fact a KM-RE of $g^*(\varepsilon)$ by Proposition 3.2 in KM, and so it is an ARE a fortiori. Since $\lim_{\varepsilon \to 0} g^*(\varepsilon) = g$, Theorem 2 implies that (μ_1^*, μ_2^*) is an ARE of g. By a symmetric argument, it can be readily seen that (μ_1^{**}, μ_2^{**}) is another ARE of g. Finally, to show that $(\mu_1^{***}, \mu_2^{***})$ is an ARE, it suffices to point out that it is the unique Nash equilibrium of the zero-sum game $g^{***}(\varepsilon)$ with the payoff matrix of player 1 given by

	$\mathbf{c}^1$	$\mathbf{c}^2$	$\mathbf{c}^3$	$\mathbf{c}^4$
$\mathbf{r}^1$	$1+\varepsilon$	-1	0	0
$\mathbf{r}^2$	-1	$1+\varepsilon$	0	0
$\mathbf{r}^3$	0	0	$1+\varepsilon$	-1
$\mathbf{r}^4$	0	0	-1	$1+\varepsilon$

It has been already pointed out that, by Claim 5 , for a specific family of elaborations $(\mathcal{U}_\delta)_{0<\delta<1}$ of the game g, the limit action distribution of the unique (exact) BE is uniform on $\{\mathbf{r}^1, \mathbf{r}^2\} \times \{\mathbf{c}^1, \mathbf{c}^2\}$, i.e., it corresponds to (μ_1^*, μ_2^*). However, the proof of Claim 5 cannot be adapted to show uniqueness of an ε-BE distribution in $\mathcal{U}_\delta$ (for $\varepsilon > 0$). In fact, $\mathcal{U}_\delta$ must possess multiple ε-BE distributions for all small enough δ, for otherwise neither (μ_1^{**}, μ_2^{**}) nor $(\mu_1^{***}, \mu_2^{***})$ would be ARE, contrary to what we have just shown. Technically, what drives the uniqueness result is the existence of a strictly dominant action for player 1 in the game with payoff matrix (18). But the dominant action $\mathbf{r}^1$ beats $\mathbf{r}^2$ only by a tiny margin, which becomes vanishingly small as $\delta \to 0$. Thus in an ε-BE, where players can make small but non-negligible (compared to δ) mistakes in their best responses, the choice of $\mathbf{r}^1$ by player 1 is by no means necessary when $t_1 = 1$. This is what allows players in an ε-BE to shift the weight from the actions prescribed by the unique (exact) BE $(\widehat{\sigma}_1, \widehat{\sigma}_2)$ to other actions, thereby generating action distributions corresponding to (μ_1^{**}, μ_2^{**}) or $(\mu_1^{***}, \mu_2^{***})$ in the limit as $\delta \to 0$. But our argument that establishes the approximate robustness of (μ_1^{*}, μ_2^{*}), (μ_1^{**}, μ_2^{**}), and $(\mu_1^{***}, \mu_2^{***})$ works *indirectly*, without the need to explicitly construct an ε-BE in every elaboration close to g, having the property that the corresponding action distributions approximate the chosen equilibrium.

Remark 7 *Similar but simpler arguments can be used to show that KM-RE may not exist in a two-player two-action game, unlike the ARE (whose existence in such games is guaranteed by Corollary 4). Indeed, consider a two-player coordination game h, whose payoffs given by the following matrix:*

	$\mathbf{c}^1$	$\mathbf{c}^2$
$\mathbf{r}^1$	$(1,1)$	$(0,0)$.
$\mathbf{r}^2$	$(0,0)$	$(1,1)$

Amend the construction of elaboration $\mathcal{U}_\delta = \left\{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u_i\}_{i\in\mathcal{I}}\right\}$ above, by replacing the game g with h, and the game $\tilde{g}$ with $\widetilde{h}$ whose payoffs are described by:

	$\mathbf{c}^1$	$\mathbf{c}^2$
$\mathbf{r}^1$	$(10,10)$	$(10,0)$.
$\mathbf{r}^2$	$(0,10)$	$(0,0)$

The unique BE of $\mathcal{U}_\delta$ consists of playing the first *pure action in every state of nature, by both players. In a similarly amended elaboration $\mathcal{U}'_\delta = \left\{\Omega, P, \{Q_i\}_{i\in\mathcal{I}}, \{u'_i\}_{i\in\mathcal{I}}\right\}$, the unique BE prescribes playing the* second *pure action in every state of nature, by both players. As the corresponding BE action distributions are supported on distinct action profiles, $(\mathbf{r}^1, \mathbf{c}^1)$ and $(\mathbf{r}^2, \mathbf{c}^2)$, it is clear that the game h has no KM-RE.*

6 Concluding Remarks

The notion of KM-robustness stipulates that a robust NE of the game g must be imitable by a certain BE action distribution in any incomplete information elaboration of g, provided that elaboration is sufficiently "close" to g. The notion of closeness that underlies KM-robustness is rather loose, as it does not take into account information details that are not directly related to each player's knowledge of his own payoffs, and hence KM-robustness is a very strong property[11]. If one strengthen the notion of closeness, however, a weaker concept of robustness is obtained, and Kajii and Morris (1997b) take this route. Obviously, analogous exercises can be carried out for the approximate robustness. We discuss one particular case of importance below.

Recall that in an incomplete information elaboration of g, there is a high probability that each player knows that his payoffs are given by g but it is *not* required that, with high probability, it is common knowledge that g is being played. Since the game being common knowledge is often deemed as a prerequisite for an NE being played, in the context of a robustness excercise it is of interest to check the implication of an additional requirement that, for players in an elaboration, it is "close to common knowledge" that g is being played. Kajii and Morris (1997b) consider the following mode of convergence of elaborations to g, which they call *limit common knowledge (LCK)*: the probability of the event where it is common p-belief that g is played, must approach 1, for every $p < 1$.[12] They show that any semi-strict NE (which is just a slight generalization of a strict NE) is KM-robust w.r.t. LCK elaborations. But not all NE are KM-robust w.r.t. LCK elaborations.

In our context, call an NE ε-robust w.r.t. LCK elaborations if for any sequence of elaborations LCK-converging to g, there is a sequence of ε-BE action distributions in these elaborations which converges to the NE. An NE is then LCK-approximately robust (LCK-ARE) if it is ε-robust w.r.t. LCK elaborations for any $\varepsilon > 0$.

It can then be shown that *any* NE is an LCK-ARE, by applying Theorem B of Monderer and Samet (1989, p. 184). The argument works roughly as follows. Choose an NE, and fix any sequence of elaborations LCK-converging to g. To establish ε-robustness of that NE, construct a strategy profile in each elaboration using the following steps: (i) fix a high $p < 1$; (ii) on those information sets of a player where it is common p-belief that g is played, assign him the (mixed) action prescribed by the given NE; (iii) consider the game in which players only take action on information sets *not* included in step (ii) – and play actions fixed in step (ii) elsewhere – and take some equilibrium of this game. This fully specifies all players' strategies in each elaboration. Note that, by construction, players best-respond on information sets where their actions have been determined in step (iii), and on information sets where their actions have been set as the given NE actions in step (ii),

[11] See Footnote 2.

[12] See Kajii and Morris (1997b) for a precise definition.

players ε-best respond if the elaboration is close enough to g, since they know that the opponents play the NE actions with high probability.

Viewed in a wider context, this result shows that LCK-approximate robustness completely characterizes the NE of any complete information game g, conforming with the analysis of Monderer and Samet (1989); to put it differently, we can accommodate the analysis of Monderer and Samet (1989) in the study of informational robustness of equilibria. Although no new mathematical technique is involved in this observation, it is worthy of note, since it casts approximate robustness in another role, as one of the conceptual uderpinnings of Nash equilibrium.

We conclude with an open question. As has been pointed out in the proofs of Corollaries 3 and 4, both KM-RE and ARE exist in *generic* two-player zero-sum, and two-player two-action, games, and it can also be shown that these equilibria are generically unique[13]. Thus, in generic games of these two categories, the notions of KM-robustness and approximate robustness are equivalent, though we know (KM, Oyama and Takahashi (2011), Section 5 of this paper) that the equivalence does not hold in general. It is not known to us, at present, whether the notions differ on an open set of complete information games, or coincide (or, at least, are equivalent in terms of existence) for a generic choice of game.

[13]For instance, a generic two-player zero-sum game has a unique correlated equilibrium by the proof of Corollary 3, and as such it is the unique ARE (since all limit points of ε-BE action distributions in incomplete information elaborations converging to g must be ε-correlated equilibria, and any limit point of the latter, when $\varepsilon \to 0$, is an (exact) correlated equilibrium).

References

1. Bohnenblust, H.F., S. Karlin, and L. S. Shapley (1950). "Solutions of Discrete Two-Person Games," Contributions to the theory of games (ed. by H. W. Kuhn and A. W. Tucker), Princeton, Princeton University Press.

2. Chen, Y.-C. and S. Xiong (2011). "Robust Selection of Rationalizability," mimeo.

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

4. Kajii, A., and S. Morris (1997b). "Refinements and Higher Order Beliefs: A Unified Survey," Northwestern Math Center Discussion Paper No. 1197

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

6. Morris, S., and T. Ui (2005). "Generalized Potentials and Robust Sets of Equilibria," *Journal of Economic Theory* **124**, pp. 45-78.

7. Oyama, D., and S. Takahashi (2011). "On the Relationship between Robustness to Incomplete Information and Noise-Independent Selection in Global Games," *Journal of Mathematical Economics* **47 (6),** pp. 683-688.

8. Ui, T. (2001). "Robust Equilibria of Potential Games," *Econometrica* **69**, pp. 1373-1380.

9. Weinstein, J. and M. Yildiz (2007), "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinement," *Econometrica* **75**, pp. 365-400.