6-2010

How robust is undominated Nash implementation?

Takashi KUNIMOTO
Singapore Management University, tkunimoto@smu.edu.sg

How Robust is Undominated Nash Implementation?

Article · January 2008

1 author:

Takashi Kunimoto
Singapore Management University

How Robust is Undominated Nash Implementation? *

Takashi Kunimoto†

First Version: February 2007
This Version: June 2010

Abstract. Palfrey and Srivastava (1991) show that almost any social choice correspondence (SCC) is implemented in undominated Nash equilibrium, a refinement of Nash equilibrium. By requiring solution concepts to have closed graph in the limit of complete information, Chung and Ely (2003) investigate the robustness of undominated Nash implementation. Their robustness test concludes that when preferences are strict (or more generally, hedonic), only monotonic SCCs can be implemented in the closure of the undominated Nash (equilibrium) correspondence.

This paper re-examines this robustness test. I show that almost any SCC is implemented in the closure of the undominated Nash correspondence, provided that the planner is certain that there is "approximate" common knowledge. I also show that only monotonic SCCs can be implemented in the closure of the undominated Nash correspondence, provided that the planner is only nearly certain that there is approximate common knowledge. Therefore, this robustness test, on the one hand, generates new restrictions imposed on the set of implementable SCCs, and on the other hand, clarifies the extent to which the permissive implementation results are sustained.

JEL classification: C72, D78, D82, D83

Keywords: Approximate common knowledge; Implementation; Monotonicity; Robustness; Undominated Nash equilibrium

*This paper is a substantial revision of "Robust Implementation under Approximate Common Knowledge." I am indebted to Herakles Polemarchakis and Roberto Serrano for generous support and advice. I am very thankful to an associate editor and anonymous referees for detailed comments and suggestions which significantly improved this paper. I am grateful to Max Amarante, Jeff Ely, Aviad Heifetz, Olivier Tercieux, Licun Xue, and Takehiko Yamato for valuable comments and discussions. I also benefit from the audiences of seminars and conferences in various places. All remaining errors are my own. Financial supports from NSF grant SES-0133113, the Stephen Robert Dissertation Fellowship, McGill startup grant, McGill Professorship Fund, FQRSC, and SSHRC of Canada are gratefully acknowledged.

†Department of Economics, McGill University, 855 Sherbrooke Street West, Montreal, Quebec, Canada, H3A 2T7, takashi.kunimoto@mcgill.ca
URL: http://people.mcgill.ca/takashi.kunimoto/

1 Introduction

Let us consider a society consisting of a finite set of individuals. The society has a social choice rule (or correspondence) which associates with each state of the world a subset of possible outcomes. A goal of implementation theory is to characterize the relationship between the structure of the institution (or mechanism) through which individuals interact and the outcome of that interaction, given a social choice rule and a domain of environments. This paper addresses *complete information* environments and asks the question of *full and exact* implementation: the search for mechanisms whose entire set of equilibrium outcomes "exactly" coincides with the given rule. In Section 6, I briefly discuss *virtual (approximate)* implementation, which only requires that the entire set of equilibrium outcomes "approximate" the rule.

The fundamental work on Nash implementation is initiated by Maskin (1999).[3] Maskin (1999) shows that a condition called *monotonicity* is necessary and almost sufficient for Nash implementation. It turns out that monotonicity is quite a demanding condition and the literature tried to obtain less restrictive characterizations using *refinements* of Nash equilibrium. In particular, Palfrey and Srivastava (1991) introduce a new solution concept, *undominated Nash equilibrium* and prove that almost any social choice correspondence (SCC) is implemented in undominated Nash equilibrium when there are at least three players. This is a striking result in that it says that for almost any social choice rule, it is possible to construct a decentralizing procedure whose resulting outcomes are precisely those prescribed by the rule. However, Jackson, Palfrey, and Srivastava (1994) argue that the power of this permissive result derives from the fact that we have not imposed any restrictions on the implementing mechanism. They propose a condition, "boundedness," which requires that if an action is weakly dominated, then it is weakly dominated by an undominated action. Jackson et al (1994) characterize a set of SCCs which are implemented in undominated Nash equilibrium by bounded mechanisms. While the restriction to bounded mechanisms eliminates some SCCs which are undominated Nash implemented with "unbounded" mechanisms, the remaining set of implementing SCCs is still larger than that for Nash implementation.

Chung and Ely (2003) elaborate on the power of the permissive result from a different perspective. Their argument is based on the consideration of "robustness" of implementability to the relaxation of the complete information assumption. The mechanisms used in the proofs of those classical results make conspicuous use of the assumption of complete information. Complete information entails common knowledge of payoffs, an assumption generally taken to be at best a simplifying assumption, but often a less innocuous one. Chung and Ely address the following question. Suppose that the planner acknowledges that complete information is an idealization and that in the true environment players may be uncertain about the state of the world. Then, one asks: what SCCs can be implemented by mechanisms that provide the desired outcome in *all* equilibria of environments that are "close" to complete information? Chung and Ely conclude that

[3] The original paper had been circulated since 1977.

this type of robustness test yields surprisingly strong restrictions on implementability. Without imposing any restrictions on the mechanism used, they show that only monotonic SCCs can be robustly implemented. This means that even when *all* the equilibrium outcomes under complete information yield the planner's desired outcomes, there may be environments close to complete information with equilibrium outcomes far from the desired set.

This paper revisits Chung and Ely (2003) and characterizes the class of nearby environments which is sufficient to obtain the monotonicity-as-a-necessary-condition result. First, I look at the class of nearby environments where the planner is *certain* that there is *approximate* common knowledge in the sense of Monderer and Samet (1989) (I will define this in Section 2.3). In this class of nearby environments, I show that if preferences are strict, the undominated Nash equilibrium correspondence has closed graph in the limit of complete information (Theorem 1). Furthermore, in the same class of nearby environments, I show that if preferences are strict, the undominated Nash equilibrium correspondence is lower hemi-continuous in the limit of complete information (Theorem 2). Combining Palfrey and Srivastva's (1991) mechanism with Theorems 1 and 2, I show that almost any SCC is implementable in the closure of the undominated Nash equilibrium correspondence (Corollary 3). Second, I look at a larger class of nearby environments where the planner is *nearly certain* that there is approximate common knowledge. In this larger class of nearby environments, it can be shown that the undominated Nash equilibrium correspondence exhibits a discontinuity in the limit of complete information in *any* mechanism that implements a non-monotonic SCC in undominated Nash equilibrium (Theorem 4). This is indeed Chung and Ely's theorem 1 and this paper's contribution lies in clarifying exactly "when" the monotonicity-as-a-necessary-condition result applies. Moreover, I extend the result to social choice "correspondences." Therefore, there is a big discrepancy between "certainty" and "near certainty" for the corresponding robustness requirements. This also explains how delicately the permissive implementation results of undominated Nash implementation rely on complete information. Thus, I propose a way of checking the robustness of the permissive implementation results and show that the robustness test, on the one hand, generates new restrictions imposed on the set of implementable SCCs, and on the other hand, clarifies the extent to which the permissive implementation results are sustained.

A final word is called for regarding some related but different concept of robust implementation. Following the approach of Bergemann and Morris (2005), Bergemann and Morris (2010) propose a different concept of robust full (exact) implementation. Their robustness is the requirement that implementation succeed in "all" type spaces coherent with a given payoff type space. Their approach is similar to this paper in that the payoff type space is embedded in a richer type space. However, I stress two differences: (1) their robustness test is *global* in the sense that they make *no* assumptions about players' higher order beliefs. That is, the planner takes into account that players can have any beliefs. On the other hand, this paper's robustness test substantially restrict players' admissible beliefs to be close to complete information. In this sense, this paper's robustness is much less demanding than theirs. (2) Bergemann and Morris assume that each player always knows his own payoff type[4], while this paper allows for the possibility that some player does not even know his own payoff type. In this sense, this paper's robustness is more demanding than theirs. Hence, logically speaking, this paper's robustness is neither stronger nor weaker than theirs.

The rest of the paper is organized as follows. In section 2, I formalize the setup and definitions. In section 3, I illustrate the main idea of the paper by example. In section 4, I show the main results: Theorem 1 identifies the conditions under which the undominated Nash correspondence has closed graph in the limit of complete information and Theorem 2 shows that under the same conditions identified by Theorem 1, the undominated Nash correspondence is also lower hemi-continuous in the limit of complete information. In section 5, I use the main results of the previous section and propose a way of checking the robustness of implementation results. Section 6 gives concluding remarks and the Appendix contains omitted proofs from the text.

2 The Basic Setup

2.1 The Environment

There is a finite set $N = \{1, \dots, n\}$ of players. Each player i has a bounded utility function $u_i : A \times \Theta \to \mathbb{R}$, where $\Theta = \{\theta_1, \dots, \theta_K\}$ denotes the finite set of *payoff* states, and A denote the set of *pure* outcomes.[5] I assume that for each player $i \in N$, any $\theta \in \Theta$, and any $a, a' \in A$, $u_i(a;\theta) = u_i(a';\theta)$ if and only if $a = a'$. This is the assumption of *strict preferences*.[6] A *social choice correspondence* (SCC) is a mapping F which associates a subset of A with each $\theta \in \Theta$. A single-valued social choice correspondence is a *social choice function* (SCF) denoted f. Hence, any selection of SCC F is a social choice function.

2.2 Type Space

While maintaining that the payoff space Θ is common knowledge, I want to allow for players to have richer beliefs about the payoff state than complete information prescribes. To do so, I embed the payoff space Θ into a state space Ω, where Ω is the finite set of states of the world.[7] Throughout the paper, I fix the state space Ω, which is perfectly general in the sense of the degree and nature of the information structures that it permits. Associated with each embedding is an onto mapping $\xi : \Omega \to \Theta$. Let Ξ be the set of all such mappings from $\Omega \to \Theta$. Let Ψ be the space of all partitions of Ω, the elements

[4]Artemov, Kunimoto, and Serrano (2010) call this the *coherence* assumption. In other words, this coherence assumption will be violated in this paper's analysis.

[5]All the results in the current paper do not depend upon a particular representation of $u_i(\cdot)$.

[6]Strict preferences can be extended to more general hedonic preferences, as Chung and Ely (2003) did. I find strict preferences easier to explain the results and opt for this formulation.

[7]When I only consider finite mechanisms, I can also handle the countably infinite Ω in the rest of the paper.

of which are in $2^{\Omega}\backslash\{\emptyset\}$. For each $\Pi_i \in \Psi$ and $\omega \in \Omega$, I denote by $\Pi_i(\omega)$ the element of Π_i which contains ω. I call $\Pi_i \in \Psi$ player i's *partition correspondence*. Denote $\Pi = (\Pi_i)_{i \in N} \in \Psi^n$. I call $(\Pi, \xi) \in \Psi^n \times \Xi$ a *type space* and fix it throughout.

2.3 Mechanisms

A *mechanism* $\Gamma = (M, g)$ consists of a *message space* M and an *outcome function* $g : M \to A$. Here $M \equiv \times_{i \in N} M_i$ and M_i is player i's message space. Throughout the paper, I will impose the following restriction on the class of mechanisms considered.

Assumption 1 (M1: Countability) *M_i is countable for each $i \in N$*

Remark: Since Ω is assumed to be finite, by $M1$, the pure strategy space is at most countable. Thus, I find M1 very useful to define the convergence in the pure strategy space. See also the remark after Proposition 1. In addition to its technical usefulness, M1 is also rich enough to encompass the integer games or devices alike when one constructs mechanisms. Note that the integer game is the one in which each player has to announce an integer and becomes the dictator when his integer is the largest one. This is an important inclusion because most of canonical mechanisms in the literature use devices like integer games.

Recall that a type space (Π, ξ) is fixed throughout. Fix a mechanism Γ satisfying M1. Player i's pure strategy σ_i is a mapping from Ω to M_i, which is Π_i-measurable. I will denote the space of player i's pure strategies by Σ_i. The space of pure strategy profiles is denoted $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ with generic element σ. Denote $\Sigma_{-i} \equiv \Sigma_1 \times \cdots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \cdots \times \Sigma_n$.[8] Throughout I assume that for each $i \in N$, Σ_i is a countable space with the discrete topology. Let Σ be endowed with the product topology.[9] I shall show below that if Σ is endowed with the discrete topology, the mechanism is *well-behaved*.

Let $\mathcal{P}$ be the space of all priors over Ω. Let $\text{supp}(\mu) \equiv \{\omega \in \Omega | \mu(\omega) > 0\}$. Throughout the paper, I assume that any prior $\mu \in \mathcal{P}$ has *full support* on Θ: for any $\theta \in \Theta$, there exists $\omega \in \text{supp}(\mu)$ with $\xi(\omega) = \theta$.

For any $\omega, \omega' \in \Omega$, define player i's belief over Ω conditional on $\Pi_i(\omega)$ as follows:

$$\mu(\omega'|\Pi_i(\omega)) = \begin{cases} \mu(\omega') \Big/ \sum_{\tilde{\omega} \in \Pi_i(\omega)} \mu(\tilde{\omega}) & \text{if } \omega' \in \Pi_i(\omega) \text{ and } \omega' \in \text{supp}(\mu) \\ 0 & \text{otherwise.} \end{cases}$$

When considering decision making under uncertainty, I assume that each player's preferences under uncertainty are represented by *expected utility*.[10] I define player i's *best*

[8] Similar notation will be used for products of other sets.

[9] Then, even with the discrete topology, M1 guarantees that the pure strategy space maintains nice topological properties, such as σ-compactness and separability. This is no longer true if Σ is uncountable.

[10] Most of the results can be extended from expected utility to much more general representations for preferences under uncertainty. Some additional assumptions I need for the results are very similar to Assumptions 1 and 2 of Chung and Ely (2003).

response correspondence as $b_i^\Gamma : \mathcal{P} \times \Sigma_{-i} \to \Sigma_i$: for any $\mu \in \mathcal{P}$ and $\sigma_{-i} \in \Sigma_{-i}$,

$$b_i^\Gamma(\mu, \sigma_{-i}) = \left\{ \sigma_i \in \Sigma_i \left| \begin{array}{c} \sum_{\tilde{\omega}\in\Omega} \mu(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i'(\tilde{\omega}), \sigma_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) \right] \geq 0 \\ \forall \omega \in \text{supp}(\mu),\ \forall \sigma_i' \in \Sigma_i \end{array} \right. \right\}.$$

The way I introduce players' beliefs over Ω into the analysis reveals that the common prior assumption (CPA) is maintained throughout. One important question is whether or not the CPA is necessary. Theorems 1 and 2, however, are intact as long as the *common support assumption*, which is implied by the CPA, is fulfilled.[11] Therefore, I conclude that the CPA is crucial to the extent that the common support assumption is indispensable for Theorems 1 and 2.[12] Next, I define the best response correspondence as $\psi_\Gamma^{BR} : \mathcal{P} \times \Sigma \to \Sigma$ as follows: for any $\mu \in \mathcal{P}$ and $\sigma \in \Sigma$,

$$\psi_\Gamma^{BR}(\mu, \sigma) = \{\tilde{\sigma} \in \Sigma |\ \tilde{\sigma}_i \in b_i^\Gamma(\mu, \sigma_{-i})\ \ \forall i \in N\}.$$

The graph of ψ_Γ^{BR} is just the subset of $\mathcal{P} \times \Sigma \times \Sigma$, defined by

$$\text{graph } \psi_\Gamma^{BR} \equiv \left\{ (\mu, \sigma, \hat{\sigma}) \in \mathcal{P} \times \Sigma \times \Sigma |\ \hat{\sigma} \in \psi_\Gamma^{BR}(\mu, \sigma) \right\}.$$

I say that $\{\mu^k\}_{k=1}^\infty$ is an *elaboration* of a prior μ if, for any $\omega \in \Omega$, $|\mu^k(\omega) - \mu(\omega)| \to 0$ as $k \to \infty$. This is the point-wise convergence of priors over Ω. The next proposition shows that if there exists $\{(\mu^k, \sigma^k)\}_{k=1}^\infty$ with $\sigma^k \to \sigma$ and $\mu^k \to \mu$ such that for each k, σ^k is a Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$, then σ is a Bayesian Nash equilibrium of the game $\Gamma(\mu)$. Note that under complete information, this σ is not necessarily a Nash equilibrium but rather induces a correlated equilibrium distribution because the players may correlate their actions via information in the elaborations. See the example of Section 3.3 for this correlation effect. In fact, if C1 holds (to be defined in Section 2.4), Theorem 1 of this paper will show that this σ reduces to a Nash equilibrium of the complete information game.

Proposition 1 (Γ is well-behaved) *Let Γ be a mechanism such that Σ is endowed with the discrete product topology. Then, Γ is **well-behaved** in the following sense: for any $(\mu, \sigma) \in \mathcal{P} \times \Sigma$, $(\mu, \sigma, \sigma) \in$ graph ψ_Γ^{BR} whenever there exists a sequence $\{(\mu^k, \sigma^k)\}_{k=1}^\infty$ such that (i) $(\mu^k, \sigma^k, \sigma^k) \in$ graph ψ_Γ^{BR} for each k and (ii) $(\mu^k, \sigma^k) \to (\mu, \sigma)$ as $k \to \infty$.*

Remark: The fact that Γ is well-behaved is consistent with the use of the integer games because there exists "no" Nash equilibrium in the integer game. Hence, requiring that Γ be well-behaved is far weaker than requiring that for any $\mu \in \mathcal{P}$, there always exist

[11] The common support assumption says that if some player assigns zero probability to a certain event, then all other players assign zero probability to the same event.

[12] For example, Theorem 1 of this paper cannot be true without the common support assumption, i.e., one can modify the example in Section 3.1 by dropping the common support assumption in which the undominated Nash equilibrium correspondence exhibits a discontinuity at a complete information even if all the sufficient conditions for Theorem 1 are satisfied. I am grateful to an anonymous referee for pointing this out.

a Bayesian Nash equilibrium in the game $\Gamma(\mu)$. Notice that M1 makes the use of the discrete topology reasonable when I consider the convergence in the pure strategy space Σ.

Proof of Proposition 1: Consider a sequence $\{(\mu^k, \sigma^k)\}_{k=1}^{\infty}$ such that $(\mu^k, \sigma^k, \sigma^k) \in$ graph ψ_Γ^{BR} for each k and $(\mu^k, \sigma^k) \to (\mu, \sigma)$ as $k \to \infty$. Since $\sigma^k \to \sigma$ in the discrete product topology, we obtain that $\sigma^k = \sigma$ for each k large enough. This implies that $(\mu^k, \sigma, \sigma) \in$ graph ψ_Γ^{BR} for each k large enough. What we want to show is that $(\mu, \sigma, \sigma) \in$ graph ψ_Γ^{BR}. Suppose, on the contrary, that $(\mu, \sigma, \sigma) \notin$ graph ψ_Γ^{BR}. Then, there exist $i \in N$, $\sigma_i^{'}$, and $\omega \in \text{supp}(\mu)$ such that

$$\sum_{\tilde{\omega} \in \Omega} \mu(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma_i^{'}(\tilde{\omega}), \sigma_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega})) \right] > 0.$$

By the continuity of expected utility, for each k large enough, we also obtain

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma_i^{'}(\tilde{\omega}), \sigma_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega})) \right] > 0.$$

This implies that $(\mu^k, \sigma^k, \sigma^k) \notin$ graph ψ_Γ^{BR} for each k large enough, which is a contradiction. ■

2.4 Embedding Complete Information into Incomplete Information Structures

In order to talk about complete and near-complete information, I define complete information priors and notions of closeness to complete information in terms of topology over the set of priors $\mathcal{P}$. To define those, I need some definitions. This paper uses Monderer and Samet's (1989) concept of "common p-belief" as an approximation to common knowledge, which is common 1-belief.

Fix a type space (Π, ξ) and a prior $\mu \in \mathcal{P}$. Let $B_i^q(E) \equiv \{\omega \in \Omega | \mu(E|\Pi_i(\omega)) \geq q\}$ denote the set of states in which player i assigns probability at least q to the event E. I call this player i's *q-belief operator*. In particular, when $q = 1$, I call B_i^1 player i's *1-belief operator* corresponding to player i's knowledge operator. An event E is said to be *q-evident* if $E \subset B_i^q(E)$ for all $i \in N$. This means that whenever E is true, everyone believes with probability at least q that E is true. An event E is said to be *common q-belief* at ω if there exists a q-evident event F such that $\omega \in F \subset \bigcap_{i \in N} B_i^q(E)$. For any event E, define $CB^q(E) \equiv \bigcap_{i \in N} B_i^q(E)$. I will loosely say that an event E is *approximate common knowledge* at ω if E is common q-belief at ω, for q close to 1. In particular, an event E is said to be *common knowledge* at ω if it is common 1-belief at ω, that is, $\omega \in E \subset CB^1(E)$.

For any $\theta \in \Theta$, define $\mathcal{G}^\theta \equiv \{\omega \in \Omega |\ \xi(\omega) = \theta\}$ as the set of states that corresponds to the payoff state θ. Next, I define complete information.

Definition 1 *$\mu \in \mathcal{P}$ is said to be a* ***complete information*** *prior if, for any $\omega \in \Omega$, $\mu(\omega) = 0$ whenever $\omega \notin \bigcup_{\theta\in\Theta} CB^1(\mathcal{G}^\theta)$.*

The above requirement says that it is always common knowledge which payoff state is being realized. Let $\Gamma(\theta)$ denote a complete information game in which θ is common knowledge. Fixing a type space (Π, ξ), this paper concerns an elaboration $\{\mu^k\}_{k=1}^{\infty}$ converging to a complete information prior μ.[13] Throughout the paper, I maintain the following restriction on the class of elaborations. I will mention it wherever I do not need this assumption.

Assumption 2 (C1: Consistency) *Every elaboration $\{\mu^k\}_{k=1}^{\infty}$ of a complete information prior μ is* ***consistent*** *in the following sense: for any $\theta \in \Theta$, there exists $\underline{q} \in (0, 1)$ such that for any $q \in [\underline{q}, 1]$ and any $i \in N$, there exists $\omega \in supp(\mu)$ with $\xi(\omega) = \theta$ for which $B_i^q(\mathcal{G}^\theta) = \Pi_i(\omega)$.*

Remark: C1 says that at all states where player i believes with high probability that some payoff state is realized, he can only choose one action (or message) uniformly over those states. Then, a violation of C1 may introduce extra correlations so that the focus on Nash equilibrium might be inadequate.[14] In Section 3.3, I will show by example that without C1, one can construct a Bayesian Nash equilibrium that induces a correlated equilibrium distribution of the complete information game whose support involve a *non-Nash* equilibrium. This implies that C1 cannot be completely dispensed with for the undominated Nash equilibrium correspondence to have closed graph in the limit of complete information.[15]

Define as $\mathcal{G}(\varepsilon)$ the set of all states in which there is a common $(1-\varepsilon)$-belief about which payoff state being realized as follows:

$$\mathcal{G}(\varepsilon) = \left\{\omega \in \Omega \middle| \ \exists\theta \in \Theta \text{ such that } \Gamma(\theta) \text{ is common } (1-\varepsilon)\text{-belief at } \omega\right\}.$$

Now, I am ready to introduce a special class of elaborations called d^*-elaborations below:

[13] While this looks a very similar argument in Fudenberg, Kreps, and Levine (1988), Dekel and Fudenberg (1990), and Kajii and Morris (1997a), I make a very distinct argument. In this paper, I fix the payoff space and perturb only the players' beliefs over the fixed payoff space so that the set of messages in the mechanism remain cheap talk and do not enter directly into the payoff functions. These authors are, on the other hand, concerned with the situation in which the set of payoff states is not common knowledge, i.e., there are "crazy" types.

[14] The reader is referred to Section 2.3 of Kajii and Morris (97a) for similar arguments.

[15] Let μ be a complete information prior and σ be a Bayesian Nash equilibrium of the game $\Gamma(\mu)$. Indeed, C1 guarantees that σ is a Bayesian Nash equilibrium of the game $\Gamma(\mu)$ if and only if $\sigma(\omega)$ is a Nash equilibrium of the complete information game $\Gamma(\xi(\omega))$ for any $\omega \in \text{supp}(\mu)$. I consider the type space of Chung and Ely (2003): Let $\Omega = \Omega_1 \times \cdots \times \Omega_n$, where $\Omega_i = \Theta$ for each $i \in N$ is considered i's private signal space. For each $\theta \in \Theta$, $\xi(\theta, \ldots, \theta) = \theta$. For each $i \in N$ and all $\omega, \omega' \in \Omega$, $\Pi_i(\omega) \neq \Pi_i(\omega')$ if and only if $\omega_i \neq \omega_i'$. Hence, C1 holds trivially in their setup.

Definition 2 $\{\mu^k\}_{k=1}^{\infty}$ *is said to be a* d^****-elaboration*** *of a complete information prior* μ *if it is an elaboration of* μ *and there exists the corresponding sequence* $\{\varepsilon^k\}_{k=1}^{\infty}$ *such that* $\varepsilon^k \geq 0$ *and* $\mu^k(\mathcal{G}(\varepsilon^k)) = 1$ *for each* k *and* $\varepsilon^k \to 0$ *as* $k \to \infty$.

d^*-elaborations describe a situation where the planner is "certain" that there is approximate common knowledge. In order to ask for a more demanding robustness test, I define a slightly coarser elaboration, which is called a d^{**}*-elaboration.*

Definition 3 $\{\mu^k\}_{k=1}^{\infty}$ *is said to be a* d^{**}***-elaboration*** *of a complete information prior* μ *if it is an elaboration of* μ *and there exists the corresponding sequence* $\{\varepsilon^k\}_{k=1}^{\infty}$ *such that* $\varepsilon^k \geq 0$ *and* $\mu^k(\mathcal{G}(\varepsilon^k)) \geq 1 - \varepsilon^k$ *for each* k *and* $\varepsilon^k \to 0$ *as* $k \to \infty$.

d^{**}-elaborations describe a situation where the planner is "nearly certain" that there is approximate common knowledge. In the Appendix, I show that d^*-elaborations and d^{**}-elaborations each generate a different topology, respectively.[16] Moreover, it can be easily seen that there is a one-to-one correspondence between the topologies for a set X and the convergence classes on it. Since any d^*-elaboration is a d^{**}-elaboration as well, but the converse is not true, the topology induced by d^{**}-elaborations is strictly coarser (weaker) than the one induced by d^*-elaborations.

2.5 Domination and Equilibrium

Recall that a type space (Π, ξ) is fixed throughout. Given a prior $\mu \in \mathcal{P}$, the mechanism Γ determines a Bayesian game $\Gamma(\mu)$ in which each player i's type is $\Pi_i(\omega)$ in state ω and after observing his type, player i selects a message from the set M_i. Player i's pure strategy in the game $\Gamma(\mu)$ is a function $\sigma_i : \Omega \to M_i$ which is Π_i-measurable. Let $\sigma \equiv (\sigma_1, \ldots, \sigma_n)$ be a strategy profile in the game $\Gamma(\mu)$. With these notations, I shall define Nash equilibrium (NE) and Bayesian Nash equilibrium (BNE), respectively.

Definition 4 *A strategy profile* σ *is a* ***Bayesian Nash equilibrium*** *(BNE) of* $\Gamma(\mu)$ *if for each* $i \in N$, *state* $\omega \in \Omega$, *and strategy* σ_i', *we have*

$$\sum_{\tilde{\omega} \in \Omega} \mu(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g((\sigma_i', \sigma_{-i})(\tilde{\omega})); \xi(\tilde{\omega})) \right] \geq 0.$$

By C1 (consistency), it is clear that under a complete information prior μ, a strategy profile σ is a Bayesian Nash equilibrium of $\Gamma(\mu)$ if and only if, for any $\omega \in \text{supp}(\mu)$, $\sigma(\omega)$ is a Nash equilibrium of the complete information game $\Gamma(\xi(\omega))$.[17] The following is a definition of *weak* dominance.

[16] Monderer and Samet (1989) use d^{**}-elaborations for their robust equilibrium analysis in complete information games. Kajii and Morris (1998) and Monderer and Samet (1996) extend d^{**}-elaborations into general incomplete information games. All three papers analyze the "lower" hemi-continuity of (ε-) Bayesian Nash equilibrium correspondence. Theorem 2 of this paper is along this line.

[17] In Section 3.3, I will show by example that this equivalence breaks down if C1 (consistency) is violated.

Definition 5 *Let $\Gamma(\mu)$ be an incomplete information game. A strategy σ_i is* ***dominated*** *for some $\omega \in supp\ (\mu)$ if there exists a strategy $\sigma_i^{'}$ such that for every strategy profile σ_{-i} for $j \neq i$,*

$$\sum_{\tilde{\omega}\in\Omega} \mu(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g((\sigma_i^{'}, \sigma_{-i})(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega})) \right] \geq 0.$$

with strict inequality for at least one σ_{-i}. A strategy σ_i is ***undominated*** *if it is not dominated for any $\omega \in supp(\mu)$.*

Finally I shall define undominated Bayesian Nash equilibrium, which, once again by C1 (consistency), is equivalent to undominated Nash equilibrium under any complete information prior μ.

Definition 6 *A strategy profile σ is an* ***undominated Bayesian Nash equilibrium*** *(U) of $\Gamma(\mu)$ if it is a Bayesian Nash equilibrium (BNE) of $\Gamma(\mu)$ for which σ_i is undominated for each $i \in N$.*

2.6 The Closure of the Solution Correspondences

Recall that a type space (Π, ξ) is fixed throughout. Given a mechanism Γ, I denote the undominated Bayesian Nash equilibrium correspondence by $\psi_\Gamma^U : \mathcal{P} \to \Sigma$ where each element σ of $\psi_\Gamma^U(\mu)$ is an undominated Bayesian Nash equilibrium of $\Gamma(\mu)$ (where U stands for undominated Bayesian Nash equilibrium). That is,

$$\psi_\Gamma^U(\mu) = \{\sigma \in \Sigma |\ \exists \text{ undominated Bayesian Nash equilibrium } \sigma \text{ of } \Gamma(\mu)\}.$$

And let the graph of ψ_Γ^U be the subset of $\mathcal{P} \times \Sigma$, defined by

$$\text{graph } \psi_\Gamma^U \equiv \left\{(\mu, \sigma) \in \mathcal{P} \times \Sigma |\ \sigma \in \psi_\Gamma^U(\mu)\right\}.$$

Note that Chung and Ely (2003) take the set of outcomes A as a Hausdorff topological space, define $\mathcal{A} \equiv A^\Omega$ as a product set, and endow $\mathcal{A}$ with the product topology. Then, they define $\psi_\Gamma^{CE} : \mathcal{P} \to \mathcal{A}$ as their undominated Bayesian Nash equilibrium correspondence. Consider a convergent sequence $\{(\mu^k, \sigma^k)\}_{k=1}^\infty$ such that $(\mu^k, g \circ \sigma^k) \in \text{graph } \psi_\Gamma^{CE}$ for each k and $(\mu^k, g \circ \sigma^k) \to (\mu, \alpha) \in \mathcal{P} \times \mathcal{A}$ as $k \to \infty$. While $g \circ \sigma^k$ does converge to α by hypothesis, this does not necessarily mean that σ^k converges to some σ as $k \to \infty$. Therefore, Chung and Ely implicitly consider a much coarser topology on Σ than this paper's discrete topology on Σ. As I argued, M1 (countability) makes the discrete topology a reasonable choice for Σ. I consider this difference as an advantage of this paper's approach.

This paper asks the following question: What SCCs can be implemented by mechanisms that provide the desired outcomes in "all" equilibria in "all" environments that are close to complete information? If the planner is concerned with this question, he must consider the possibility that there may be environments close to complete information where the set of equilibrium outcomes is undesirably large. At the same time, the planner wants to make sure that in "all" environments that are close to complete information, there always exists at least one equilibrium outcome that coincides with any selection of the SCC. Thus, as in Chung and Ely (2003), I consider the "closure" of the solution correspondence ψ_Γ^U. First, I define the closure of the graph ψ_Γ^U:

Definition 7 *Let Γ be a mechanism such that Σ is endowed with the discrete topology. Then, we say that $(\mu,\sigma) \in \left[\text{graph } \psi_\Gamma^U\right]^*$ if there exists a sequence $\{(\mu^k,\sigma^k)\}_{k=1}^\infty$ such that (i) $(\mu^k,\sigma^k) \in \text{graph } \psi_\Gamma^U$ for each k; (ii) $(\mu^k,\sigma^k) \to (\mu,\sigma)$; and (iii) $\{\mu^k\}_{k=1}^\infty$ is a d^*-elaboration. Similarly, we can define $(\mu,\sigma) \in \left[\text{graph } \psi_\Gamma^U\right]^{**}$ by instead requiring that $\{\mu^k\}_{k=1}^\infty$ be a d^{**}-elaboration.*

Next, with the above definition, define the following two different closures of the undominated Bayesian Nash equilibrium correspondence: for any $\mu \in \mathcal{P}$,

$$\begin{aligned}\left[\psi_\Gamma^U\right]^*(\mu) &= \left\{\sigma \in \Sigma \middle| \ (\mu,\sigma) \in \left[\text{graph } \psi_\Gamma^U\right]^*\right\}; \text{ and} \\ \left[\psi_\Gamma^U\right]^{**}(\mu) &= \left\{\sigma \in \Sigma \middle| \ (\mu,\sigma) \in \left[\text{graph } \psi_\Gamma^U\right]^{**}\right\}.\end{aligned}$$

In the next section, by way of example, I construct a mechanism Γ^* such that $[\psi_{\Gamma^*}^U]^{**}(\mu) \neq \psi_{\Gamma^*}^U(\mu)$ at any complete information prior μ. I will postpone a much more general result (Theorem 4) until Section 5.2. In Section 4, on the other hand, I will identify conditions under which for *any* mechanism Γ, $[\psi_\Gamma^U]^*(\mu) = \psi_\Gamma^U(\mu)$ at any complete information prior μ (Corollary 2).

3 Illustration

Let me illustrate the main idea of this paper through a series of examples. In Section 3.1, I will estabilish the discontinuity of the undominated Nash correspondence with respect to d^{**}-elaborations. In Section 3.2, I can explicitly construct a d^*-elaboration satisfying C1 (consistency). Finally, in Section 3.3, I will show that C1 cannot be completely dispensed with for the undominated Nash correspondence to have closed graph in the limit of complete information.

3.1 $[\psi_\Gamma^U]^{**}(\mu) \neq \psi_\Gamma^U(\mu)$ where μ is a complete information prior

There are two players, called Andy and Bob. There are three possible outcomes, a, b, and c. There are two possible payoff states, called θ and θ'. I assume that the players' preferences over the three outcomes are strict and state dependent as follows[18]:

Andy: $u_A(b;\theta) > u_A(a;\theta) > u_A(c;\theta)$ and $u_A(b;\theta') > u_A(a;\theta') > u_A(c;\theta')$
Bob: $u_B(a;\theta) > u_B(c;\theta) > u_B(b;\theta)$ and $u_B(a;\theta') > u_B(b;\theta') > u_B(c;\theta')$

[18] The example is adapted from Jackson and Srivastava (1996, Example 5), which is also used in Chung and Ely (2003).

$\Gamma^* = (M,g)$		Bob	
		m_B	m_B'
Andy	m_A	a	a
	m_A'	c	b

Table 1: A mechanism Γ^* that UNE implements f^*

		Bob's signal	
		0	1
Andy's signal	0	$1-p$	0
	1	$p\varepsilon$	$p(1-\varepsilon)$

Table 2: An Example of d^{**}-elaborations

The planner has an SCF $f^* : \{\theta, \theta'\} \to \{a, b, c\}$ with the property that $f(\theta) = a$ and $f(\theta') = b$. It is easy to see that f^* is not monotonic (See Section 5.2 for its definition).[19] This is an environment with complete information. The planner's objective is to devise a mechanism whose *unique* undominated Nash equilibrium outcome coincides with that of f^* in every payoff state. The mechanism $\Gamma^* = (M, g)$ (shown in Table 1) is given, in which Andy chooses the row and Bob chooses the column. Here $M = M_A \times M_B = \{m_A, m_A'\} \times \{m_B, m_B'\}$ refers to the set of message profiles. The outcome function $g : M \to \{a, b, c\}$ assigns to each message profile m an alternative $g(m) \in \{a, b, c\}$.

$\Gamma^*(\theta)$ denotes a complete information game in which the payoff state is θ. The profile (m_A, m_B), leading to outcome a, is the unique undominated Nash equilibrium of the game $\Gamma^*(\theta)$ and the only undominated Nash equilibrium of the game $\Gamma^*(\theta')$ is (m_A', m_B') leading to the outcome b. Note that (m_A, m_B) is a "dominated" Nash equilibrium of the game $\Gamma^*(\theta')$. Hence, the mechanism Γ^* implements f^* in undominated Nash equilibrium.

Suppose that players may be uncertain about the state of the world. To fix the idea, I shall construct the "nearby" environments in which Andy and Bob may be uncertain about the payoff state as shown in Table 2. There are three states of the world in the nearby environments parameterized by ε: $(0, 0), (1, 0)$, and $(1, 1)$. Here $(0, 0)$ stands for the state where the payoff state is θ'. $(1, 0)$ and $(1, 1)$ stand for the states where the payoff state is θ. With probability p, the state θ is realized and the state θ' is realized with probability $1 - p$. The row is Andy's signal and the column is Bob's signal. Each player only observes his own signal.[20] Note also that C1(consistency) holds for this elaboration.

[19] I shall prove here that f^* is not monotonic. Suppose, by way of contradiction, that f^* is monotonic. Since $f^*(\theta) \neq f^*(\theta')$, by monotonicity, there must exist player i and outcome y such that $u_i(y; \theta') > u_i(f^*(\theta); \theta')$ and $u_i(f^*(\theta); \theta) \geq u_i(y; \theta)$. Such an player must be Bob because Andy has the state uniform preference. Because $f^*(\theta) = a$ and a is the best outcome for Bob, there is no such better y, which is a contradiction.

[20] The complete information prior is embedded in this set of nearby environments when $\varepsilon = 0$.

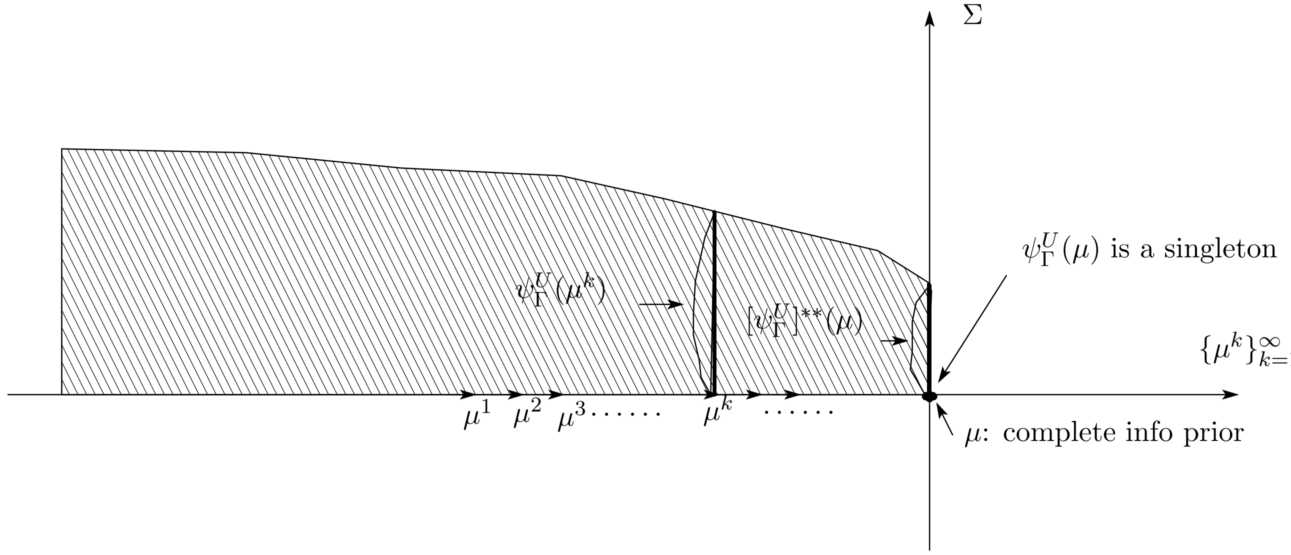

Figure 1: A schematic diagram: $\psi_\Gamma^U(\mu) \neq [\psi_\Gamma^U]^{**}(\mu)$

Suppose, under this noisy information structure, that Bob has the belief that: "Andy plays $m_A^{'}$ when he receives signal 1 and plays m_A when he receives signal 0."

Given Bob's belief specified above, playing m_B after receiving signal 0 gives the following lottery:

- $a = g(m_A, m_B)$ in $\Gamma^*(\theta')$ with probability $(1-p)/(1-p+p\varepsilon)$;
- $c = g(m_A^{'}, m_B)$ in $\Gamma^*(\theta)$ with probability $p\varepsilon/(1-p+p\varepsilon)$.

Given Bob's belief specified above, playing $m_B^{'}$ after receiving signal 0 gives the following lottery:

- $a = g(m_A, m_B^{'})$ in $\Gamma^*(\theta')$ with probability $(1-p)/(1-p+p\varepsilon)$;
- $b = g(m_A^{'}, m_B^{'})$ in $\Gamma^*(\theta)$ with probability $p\varepsilon/(1-p+p\varepsilon)$.

It turns out that m_B can be a strict best response for *any* $\varepsilon > 0$. Then, the planner is no longer confident that m_B is dominated in the game $\Gamma^*(\theta')$. Especially, at $(1, 0)$, Andy knows the game is $\Gamma^*(\theta)$ but Bob believes with high probability that the game is $\Gamma^*(\theta')$. This is a sense in which these nearby environments are far from complete information, no matter how small its likelihood is. It is not difficult to show that the nearby environments is a d^{**}-elaboration. See also Figure 1 for the summary of this subsection.

		Public signal	
		0	1
Nature's signal	0	$1-p$	0
	1	$p\varepsilon$	$p(1-\varepsilon)$

Table 3: An Example of d^*-elaborations

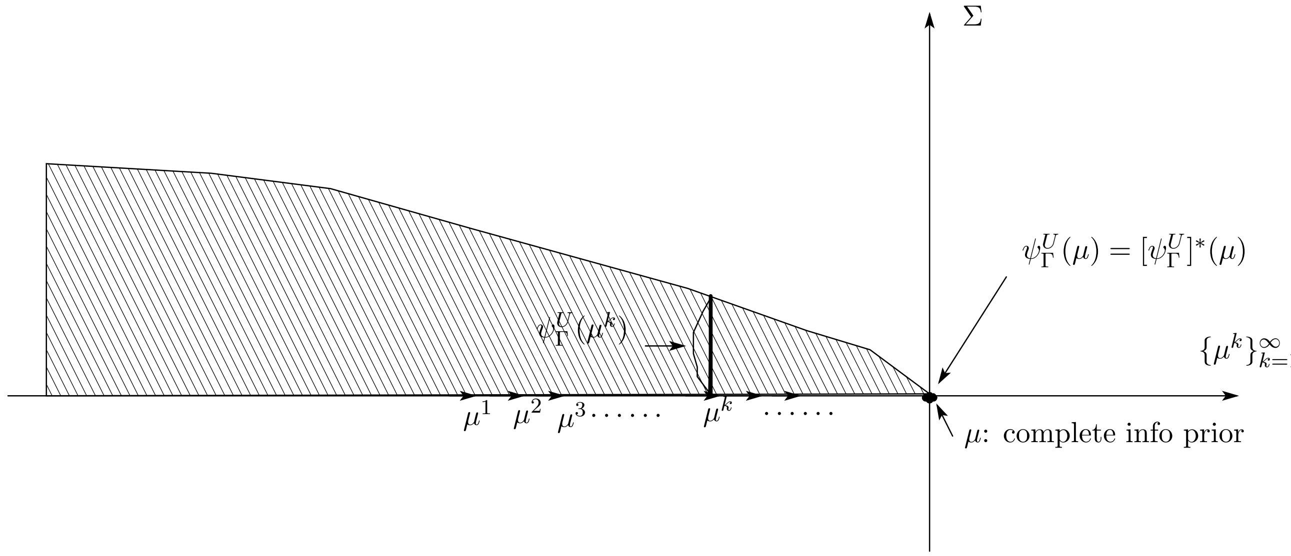

Figure 2: A schematic diagram: $\psi_\Gamma^U(\mu) = [\psi_\Gamma^U]^*(\mu)$

3.2 $[\psi_\Gamma^U]^*(\mu) = \psi_\Gamma^U(\mu)$ where μ is a complete information prior

Although I borrow most of the setups from Section 3.1, I consider alternative nearby environments: Like the previous elaboration, there are three states such that $(0,0)$ corresponds to the payoff state θ' and $(1,0)$ and $(1,1)$ correspond to the payoff state θ. Unlike the previous one, the row is Nature's signal and the column is both players' public signal. Each player never observes Nature's signal but both players commonly observe the public signal. Note that C1 holds for this elaboration as well. This information structure can be described in Table 3. It is easy to show that the nearby environments described by Table 3 is a d^*-elaboration. This example also shows that the existence of d^*-elaborations satisfying C1 is not a vacuous assumption. In Section 4, I will formally show that the "closedness" of the undominated Nash equilibrium correspondence can be restored by focusing on d^*-elaborations instead. See also Figure 2 for the summary of this subsection.

$\tilde{\Gamma}(\theta)$		Bob	
		m_B	$m_B^{'}$
Andy	m_A	$0,0$	$7,2$
	$m_A^{'}$	$2,7$	$6,6$

Table 4: Complete Information Game $\tilde{\Gamma}(\theta)$

		Bob's signal	
		0	1
Andy's signal	0	0	1/3
	1	1/3	1/3

Table 5: An Example of d^*-elaborations that do not satisfy C1

3.3 C1 cannot be completely dispensed with

I still have the same Andy and Bob from the previous two subsections but now have an otherwise different example: there are four outcomes $A = \{a, b, c, d\}$. Consider the following mechanism $\tilde{\Gamma} = (M, g)$ where $M_A = \{m_A, m_A^{'}\}$; $M_B = \{m_B, m_B^{'}\}$; $g(m_A, m_B) = a$; $g(m_A, m_B^{'}) = b$; $g(m_A^{'}, m_B) = c$; and $g(m_A^{'}, m_B^{'}) = d$. Assume that there is only one payoff state θ and each player has strict preferences at that state. The complete information $\tilde{\Gamma}(\theta)$ is given in Table 4. $\tilde{\Gamma}(\theta)$ is known to be a "game of chicken" and there are two pure strategy Nash equilibria: $(m_A, m_B^{'})$ and $(m_A^{'}, m_B)$. Now, I introduce three states of the world $(0,1), (1,0), (1,1)$ such that all the states correspond to payoff state θ. The fist component of the state is Andy's private signal and the second component of the state is Bob's private signal. Given this interpretation, I have the following partition Π_A and Π_B: $\Pi_A(0,1) = \{(0,1)\}$; $\Pi_A(1,0) = \Pi_A(1,1) = \{(1,0),(1,1)\}$; and $\Pi_B(1,0) = \{(0,1)\}$; $\Pi_B(1,0) = \Pi_B(1,1) = \{(0,1),(1,1)\}$, respectively.

I consider the nearby environments described by Table 5. Indeed, this nearby environments is a d^*-elaboration for a trivial reason. Consider the following strategy profile σ: $\sigma_A(0,1) = m_A; \sigma_A(1,0) = \sigma_A(1,1) = m_A^{'}; \sigma_B(1,0) = m_B$; and $\sigma_B(0,1) = \sigma_B(1,1) = m_B^{'}$. It is easy to show that σ is a Bayesian Nash equilibrium of the games. Note also that σ induceds a correlated equilibrium distribution. For any $q \in [0,1]$, I have that $B_A^q(\mathcal{G}^\theta) = \{(0,1)\} \cup \{(1,0),(1,1)\}$ and $B_B^q(\mathcal{G}^\theta) = \{(1,0)\} \cup \{(0,1),(1,1)\}$. Hence, C1 is violated in this elaboration. Since $(m_A^{'}, m_B^{'})$ is "not" a Nash equilibrium of the complete information game $\Gamma(\theta)$, the associated Bayesian Nash equilibrium correspondence does not have closed graph. Thus, C1 cannot be completely dispensed with if the undominated Nash correspondence is required to have closed graph in the limit of complete information.

4 The Main Results

Recall that a type space (Π, ξ) is fixed throughout. I am now ready to state one of the main results of this paper.

Theorem 1 *Suppose C1 holds. Let $\Gamma = (M, g)$ be a mechanism satisfying M1. Then, if preferences are strict, $[\psi_\Gamma^U]^*(\mu) \subseteq \psi_\Gamma^U(\mu)$ at any complete information prior μ.*

Proof of Theorem 1: Let $\{\mu^k\}_{k=1}^\infty$ be a d^*-elaboration of a complete information prior μ satisfying C1 such that there exists the corresponding sequence $\{q^k\}_{k=1}^\infty$ converging to 1 for which there is a common q^k-belief at any $\omega \in \text{supp}(\mu^k)$ about which payoff state being realized for each k.[21] Let $\{\sigma^k\}_{k=1}^\infty$ be the corresponding sequence of undominated Bayesian Nash equilibrium strategy profiles. By M1 (countability) and Proposition 1 (Γ is well-behaved), we can guarantee the existence of $\sigma \in \Sigma$ such that $\sigma^k = \sigma$ for each k big enough. Suppose, by way of contradiction, that there exist $\bar{\omega} \in \text{supp}(\mu)$ and $\theta \in \Theta$ with $\xi(\bar{\omega}) = \theta$ such that $\sigma(\bar{\omega}) = m^*$ is "not" an undominated Nash equilibrium of the game $\Gamma(\theta)$.

We must consider two cases: (1) m^* is a weakly dominated strategy profile of the game $\Gamma(\theta)$; and (2) m^* is not a Nash equilibrium of the game $\Gamma(\theta)$. In the rest of the proof, we will show that for each k big enough, there exists $\hat{\sigma}_i^k$ that either dominates σ_i^k or is a better reply to σ_{-i}^k than σ_i^k, which is a contradiction.

(1) m^* is a weakly dominated strategy profile of $\Gamma(\theta)$

By our hypothesis, there exists a nonempty subset of players $I \subseteq N$ such that for any $i \in I$, there exists $m_i^{'}$ with the following two properties:

- $u_i(g(m_i^{'}, \tilde{m}_{-i}); \theta) \geq u_i(g(m_i^*, \tilde{m}_{-i}); \theta)$ for any $\tilde{m}_{-i} \in M_{-i}$;
- $u_i(g(m_i^{'}, \hat{m}_{-i}); \theta) > u_i(g(m_i^*, \hat{m}_{-i}); \theta)$ for some $\hat{m}_{-i} \in M_{-i}$.

Choose one player $i \in I$. Without loss of generality, assume that $\mu^k(\omega) > 0$ for any $\omega \in \Pi_i(\bar{\omega})$ and any k. If this is not the case, we can simply take k big enough and ignore such states in the rest of the argument. We construct a sequence of player i's strategies $\{\hat{\sigma}_i^k\}_{k=1}^\infty$ with the following properties:

- $\hat{\sigma}_i^k(\bar{\omega}) = m_i^{'}$;
- $\hat{\sigma}_i^k(\omega) = \sigma_i^k(\omega)$ for any $\omega \notin \Pi_i(\bar{\omega})$;

By our hypothesis, $\sigma_i(\bar{\omega}) = m_i^*$ and σ_i is Π_i-measurable. So, $\sigma_i(\omega) = m_i^*$ for any $\omega \in \Pi_i(\bar{\omega})$. Since $\sigma_i^k(\bar{\omega}) = m_i^*$ for each k big enough, it can be shown that $\hat{\sigma}_i^k \neq \sigma_i^k$ for each k big enough. We shall show the following lemma.

[21] The existence of such an elaboration is shown to be non-vacuous in Section 3.2.

Lemma 1 *There exists $\bar{K} \in \mathbb{N}$ such that for any $k \geq \bar{K}$, $\hat{\sigma}_i^k$ dominates σ_i^k at any $\omega \in \Pi_i(\bar{\omega})$. That is, for any $\omega \in \Pi_i(\bar{\omega})$ and k big enough, the following two conditions hold:*

1. $\sum_{\tilde{\omega}\in\Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega}))\right] \geq 0$ *for any $\tilde{\sigma}_{-i}$;*

2. $\sum_{\tilde{\omega}\in\Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i^k(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega}))\right] > 0$ *for some $\hat{\sigma}_{-i}$.*

By Lemma 1 we aim at showing that σ^k is "not" an undominated Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$ for any $k \geq \bar{K}$, which is a contradiction.

Proof of Lemma 1: Since preferences are strict and m_i^* is dominated by $m_i^{'}$ in the game $\Gamma(\theta)$, we begin by noting the following important fact.

Fact 1 *For any $\tilde{m}_{-i} \in M_{-i}$, the following two must hold:*

1. $g(m_i^{'}, \tilde{m}_{-i}) \neq g(m_i^*, \tilde{m}_{-i}) \Longrightarrow u_i(g(m_i^{'}, \tilde{m}_{-i}); \theta) > u_i(g(m_i^*, \tilde{m}_{-i}); \theta)$;

2. $g(m_i^{'}, \tilde{m}_{-i}) = g(m_i^*, \tilde{m}_{-i}) \Longrightarrow u_i(g(m_i^{'}, \tilde{m}_{-i}); \tilde{\theta}) = u_i(g(m_i^*, \tilde{m}_{-i}); \tilde{\theta})$ *for any $\tilde{\theta}$.*

Recall that $\sigma_i^k(\omega) = m_i^*$ at any $\omega \in \Pi_i(\bar{\omega})$ for each k big enough. Since $\hat{\sigma}_i^k(\omega) = m_i^{'}$ for each $\omega \in \Pi_i(\bar{\omega})$ and each k, the previous fact further implies that, for any $\tilde{\sigma}_{-i} \in \Sigma_{-i}$, any $\tilde{\omega} \in \Pi_i(\bar{\omega})$, and any k big enough, the following two must be true:

$$\begin{aligned} g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})) \neq g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})) \quad &\Rightarrow \quad u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \theta) > u_i(g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \theta) \\ g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})) = g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})) \quad &\Rightarrow \quad u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \tilde{\theta}) = u_i(g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \tilde{\theta}) \ \ \forall \tilde{\theta} \end{aligned}$$

Since $\{\mu^k\}_{k=1}^{\infty}$ is a d^*-elaboration, for each k big enough, player i is certain that the payoff state θ is common q^k-belief at any state $\omega \in \Pi_i(\bar{\omega})$.[22] Recall that $q^k \to 1$ as $k \to \infty$. Moreover, by the continuity of expected utility, the boundedness of $u_i(\cdot)$, and the construction of $\hat{\sigma}_i^k$, we have the following: for any $\tilde{\sigma}_{-i} \in \Sigma_{-i}$, any $\omega \in \Pi_i(\bar{\omega})$, and any k big enough,

$$\sum_{\tilde{\omega}\in\Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega}); \xi(\tilde{\omega}))) - u_i(g(\sigma_i^k(\tilde{\omega}), \tilde{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega}))\right] \geq 0.$$

Hence, Condition 1 in Lemma 1 is satisfied at any state $\omega \in \Pi_i(\bar{\omega})$ for each k big enough. Let $\hat{\sigma}_{-i} \in \Sigma_{-i}$ be that $\hat{\sigma}_{-i}(\omega) = \hat{m}_{-i}$ for any $\omega \in \Omega$, as specified in our hypothesis. By the continuity of expected utility, we have that for any $\omega \in \Pi_i(\bar{\omega})$ and any k big enough,

$$\sum_{\tilde{\omega}\in\Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i^k(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega}))\right] > 0.$$

[22] This may not be true if $\{\mu^k\}_{k=1}^{\infty}$ is a d^{**}-elaboration. Set $\bar{\omega} = (0, 0)$, $\omega = (1, 0)$, and $i = B$ and $-i = A$ in the example in Section 3.1. Then, the payoff state is θ in state $\omega = (1, 0)$ and that the payoff state is θ' in state $\bar{\omega} = (0, 0)$. But player i cannot distinguish between $(0, 0)$ and $(1, 0)$.

Hence, Condition 2 in Lemma 1 is also satisfied for such $\hat{\sigma}_{-i}$. This implies $\hat{\sigma}_i^k$ dominates σ_i^k for each k big enough, which is a contradiction. This completes the proof of Lemma 1.[23] ■

(2) m^* is not a Nash equilibrium of $\Gamma(\theta)$

By our hypothesis, there exist an player i and a message $m_i^{'}$ with the following property:

$$u_i(g(m_i^{'}, m_{-i}^*); \theta) > u_i(g(m^*); \theta).$$

Recall that $\sigma^k(\bar{\omega}) = m^*$ for each k big enough. Once again, without loss of generality, assume $\mu^k(\omega) > 0$ for any $\omega \in \Pi_i(\bar{\omega})$ and any k. We construct a sequence of player i's strategies of $\{\hat{\sigma}_i^k\}_{k=1}^{\infty}$ with the following properties:

- $\hat{\sigma}_i^k(\bar{\omega}) = m_i^{'}$;
- $\hat{\sigma}_i^k(\omega) = \sigma_i^k(\omega)$ for any $\omega \notin \Pi_i(\bar{\omega})$;

As we already argued, $\hat{\sigma}_i^k \neq \sigma_i^k$ for each k big enough. We shall show the following lemma.

Lemma 2 *There exists $\bar{K} \in \mathbb{N}$ such that for any $k \geq \bar{K}$ and any $\omega \in \Pi_i(\bar{\omega})$, we have*

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \sigma_{-i}^k(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i^k(\tilde{\omega}), \sigma_{-i}^k(\tilde{\omega})); \xi(\tilde{\omega})) \right] > 0.$$

By Lemma 2 we aim to show that $\hat{\sigma}_i^k$ is a strictly better reply to σ_{-i}^k than σ_i^k for any $k \geq \bar{K}$ and any $\omega \in \Pi_i(\bar{\omega})$. Thus, this contradicts the hypothesis that σ^k is an undominated Bayesian Nash equilibrium of $\Gamma(\mu^k)$ for each k.

Proof of Lemma 2: Recall that $\sigma^k(\bar{\omega}) = m^*$ for each k big enough. Then, for each k big enough, we obtain

$$u_i(g(\hat{\sigma}_i^k(\bar{\omega}), \sigma_{-i}^k(\bar{\omega})); \theta) > u_i(g(\sigma^k(\bar{\omega})); \theta).$$

Since $\{\mu^k\}_{k=1}^{\infty}$ is a d^*-elaboration, player i is certain that the payoff state θ is common q^k-belief at any state $\omega \in \Pi_i(\bar{\omega})$. Recall that $q^k \to 1$ as $k \to \infty$. Thus, by C1 (consistency), for each k big enough, player i believes with probability q^k that $\sigma^k(\omega) = m^*$ for each $\omega \in \Pi_i(\bar{\omega})$. By the continuity of expected utility, the boundedness of $u_i(\cdot)$, and the construction of $\hat{\sigma}_i^k$, we obtain the following: for any $\omega \in \Pi_i(\bar{\omega})$ and any k big enough,

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\hat{\sigma}_i^k(\tilde{\omega}), \sigma_{-i}^k(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma^k(\tilde{\omega})); \xi(\tilde{\omega})) \right] > 0.$$

[23] It is important to note that C1 is not needed for Lemma 1.

We also claim that the above inequality holds even if $\{\mu^k\}_{k=1}^{\infty}$ is a d^{**}-elaboration. Suppose that there exists $\Omega^* \subset \Omega$ with $\mu^k(\Omega^*) \geq q^k$ for each k such that, by C1 (consistency), player i believes with probability q^k that $\sigma^k(\omega) = m^*$ for any $\omega \in \Pi_i(\bar{\omega}) \cap \Omega^*$. This is equivalent to saying that $\{\mu^k\}_{k=1}^{\infty}$ is a d^{**}-elaboration satisfying C1. By the continuity of expected utility, the boundedness of $u_i(\cdot)$, and the construction of $\hat{\sigma}_i^k$, the above inequality holds even in this case.

Hence, this shows that σ^k is "not" a Bayesian Nash equilibrium of $\Gamma(\mu^k)$ for each k big enough, which is a contradiction. This completes the proof of Lemma 2. ■

With Lemmas 1 and 2, we complete the proof of Theorem 1. ■

As I argued above, Lemma 2 is true even if $\{\mu^k\}_{k=1}^{\infty}$ is a d^{**}-elaboration. Moreover, Lemma 2 does not need the assumption of strict preferences or anything as such. Thus, I have the following corollary.

Corollary 1 *Suppose C1 holds. Let $\Gamma = (M, g)$ be a mechanism satisfying M1. Then, $[\psi_\Gamma^{BNE}]^{**}(\mu) \subseteq \psi_\Gamma^{BNE}(\mu)$ at any complete information prior μ. Here $\psi_\Gamma^{BNE} : \mathcal{P} \to \Sigma$ denotes the Bayesian Nash equilibrium correspondence.*

The next main result establishes the *lower* hemi-continuity of the undominated Bayesian Nash equilibrium correspondence with respect to the topology induced by d^{**}-elaborations, provided that preferences are strict and C1 (consistency) holds.

Theorem 2 *Assume that C1 holds and preferences are strict. Let $\Gamma = (M, g)$ be a mechanism. Then, $\psi_\Gamma^U(\mu) \subseteq [\psi_\Gamma^U]^{**}(\mu)$ at any complete information prior μ.*

Remark: Since the topology induced by d^*-elaborations is finer (stronger) than the one induced by d^{**}-elaborations, this theorem, a fortiori, shows that $\psi_\Gamma^U(\mu) \subseteq [\psi_\Gamma^U]^*(\mu)$ at any complete information prior μ as well. Note also that this result does not need any restrictions on the class of mechanisms.

Proof of Theorem 2: Let μ be a complete information prior and $\sigma \in \psi_\Gamma^U(\mu)$. By our hypothesis, we have that for any $\omega \in \text{supp}(\mu)$, $\sigma(\omega)$ is an undominated Nash equilibrium of the game $\Gamma(\xi(\omega))$. We shall claim that $\sigma \in [\psi_\Gamma^U]^{**}(\mu)$. Consider a d^{**}-elaboration $\{\mu^k\}_{k=1}^{\infty}$ such that there exist the corresponding sequence $\{q^k\}_{k=1}^{\infty}$ converging to 1 and $\Omega^* \subseteq \Omega$ with $\mu^k(\Omega^*) \geq q^k$ for which some payoff state is common q^k-belief at any $\omega \in \Omega^*$ for each k. The proof will be completed by the following Claims 1 and 2.

Claim 1: for each $i \in N$, σ_i is undominated in the game $\Gamma(\mu^k)$ for each k big enough.

Proof of Claim 1: Fix any player $i \in N$. Fix arbitrarily any $\theta \in \Theta$. Since μ has full support on Θ, there exists $\omega_\theta \in \text{supp}(\mu)$ such that $\xi(\omega_\theta) = \theta$. Without loss of generality, assume that $\mu^k(\omega) > 0$ for any $\omega \in \Pi_i(\omega_\theta)$ and any k. If this is not the case, we can simply take k big enough and ignore such states. Since σ_i is undominated in the complete information game $\Gamma(\mu)$, for any $\sigma_i' \in \Sigma_i$, there are only two mutually exclusive cases:

Case 1: there exists $\hat{\sigma}_{-i} \in \Sigma_{-i}$ such that

$$u_i(g(\sigma_i(\omega_\theta), \hat{\sigma}_{-i}(\omega_\theta)); \theta) > u_i(g(\sigma_i'(\omega_\theta), \hat{\sigma}_{-i}(\omega_\theta)); \theta).$$

Case 2: for any $\hat{\sigma}_{-i} \in \Sigma_{-i}$,

$$u_i(g(\sigma_i(\omega_\theta), \hat{\sigma}_{-i}(\omega_\theta)); \theta) = u_i(g(\sigma_i'(\omega_\theta), \hat{\sigma}_{-i}(\omega_\theta)); \theta).$$

First consider Case 1. Since $\{\mu^k\}_{k=1}^{\infty}$ is a d^{**}-elaboration, for each k big enough, player i is certain that the payoff state θ is common q^k-belief at any $\omega \in \Pi_i(\omega_\theta) \cap \Omega^*$. Recall that $\mu^k(\Omega^*) \geq q^k$ for each k and $q^k \to 1$ as $k \to \infty$. Therefore, by C1 (consistency), for each k big enough, player i believes with probability q^k that $\hat{\sigma}_{-i}(\omega) = \hat{\sigma}_{-i}(\omega_\theta)$ at each $\omega \in \Pi_i(\omega_\theta) \cap \Omega^*$. Due to the continuity of expected utility and the boundedness of $u_i(\cdot)$, we have the following: for any $\omega \in \Pi_i(\omega_\theta)$ and any k big enough,

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma_i(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i'(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega}))) \right] > 0.$$

Next consider Case 2. In this case, we can also have the following: for any $\hat{\sigma}_{-i}$ and any $\tilde{\omega} \in \Pi_i(\omega_\theta)$,

$$u_i(g(\sigma_i(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \theta) = u_i(g(\sigma_i'(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})); \theta).$$

Due to strict preferences, for any $\hat{\sigma}_{-i}$ and any $\tilde{\omega} \in \Pi_i(\omega_\theta)$, we have that $g(\sigma_i(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega})) = g(\sigma_i'(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega}))$. Thus, for any $\hat{\sigma}_{-i}$, any $\omega \in \Pi_i(\omega_\theta)$, and any k, we have

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma_i(\tilde{\omega}), \hat{\sigma}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma_i'(\tilde{\omega}), \hat{\sigma}_{-i}(\tilde{\omega}))) \right] = 0.$$

Since the previous argument goes through with any θ, Cases 1 and 2 together show that σ_i is undominated in the game $\Gamma(\mu^k)$ for each k big enough.■

Claim 2: σ is a Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$ for each k big enough.[24]

Proof of Claim 2: Suppose, by way of contradiction, that σ is "not" a Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$ for each k. That is, for each k big enough, there exist $i \in N$, σ_i', and $\omega_\theta \in \text{supp}(\mu^k) \cap \mathcal{G}^\theta$ with some $\theta \in \Theta$ such that for any $\omega \in \Pi_i(\omega_\theta) \cap \text{supp}(\mu^k)$,

$$\sum_{\tilde{\omega} \in \Omega} \mu^k(\tilde{\omega}|\Pi_i(\omega)) \left[u_i(g(\sigma_i'(\tilde{\omega}), \sigma_{-i}(\tilde{\omega})); \xi(\tilde{\omega})) - u_i(g(\sigma(\tilde{\omega})); \xi(\tilde{\omega}) \right] > 0,$$

where $\mathcal{G}^\theta$ denotes the set of states that correspond to payoff state θ. By C1 (consistency) and the full support of μ on Θ, we can guarantee without loss of generality that $\omega_\theta \in$

[24] Note that strict preferences are not needed for Claim 2.

$\text{supp}(\mu) \cap \mathcal{G}^\theta$.[25] Since $\{\mu^k\}_{k=1}^\infty$ is a d^{**}-elaboration, $u_i(\cdot)$ is bounded, and expected utility is continuous, we have

$$u_i(g(\sigma_i^{'}(\omega_\theta), \sigma_{-i}(\omega_\theta)); \theta) - u_i(g(\sigma(\omega_\theta)); \theta) > 0.$$

Define $\sigma_i^{''}$ as follows: for each $\omega \in \Omega$,

$$\sigma_i^{''}(\omega) = \begin{cases} \sigma_i^{'}(\omega_\theta) & \text{if } \omega \in \text{supp}(\mu) \cap \mathcal{G}^\theta \\ \sigma_i(\omega) & \text{otherwise.} \end{cases}$$

This implies that for any $\omega \in \text{supp}(\mu) \cap \mathcal{G}^\theta$,

$$u_i(g(\sigma_i^{''}(\omega), \sigma_{-i}(\omega)); \theta) > u_i(g(\sigma(\omega)); \theta).$$

Note that by construction, for any $\omega, \omega' \in \text{supp}(\mu) \cap \mathcal{G}^\theta$, we have $\sigma_i^{''}(\omega) = \sigma_i^{''}(\omega')$. This contradicts the hypothesis that σ is a Nash equilibrium of the complete information game $\Gamma(\theta)$. Therefore, σ is a Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$ for each k big enough.■

Claims 1 and 2 together show that σ is an undominated Bayesian Nash equilibrium of the game $\Gamma(\mu^k)$ for each k big enough. This completes the proof of Theorem 2.■

Theorems 1 and 2 together establish the next result. This will be used for the main result (Corollary 3) in the next section.

Corollary 2 *Suppose C1 holds. Let $\Gamma = (M, g)$ be a mechanism satisfying M1. Then, if preferences are strict, $[\psi_\Gamma^U]^*(\mu) = \psi_\Gamma^U(\mu)$ at any complete information prior μ.*

Proof of Corollary 2: The proof directly follows from Theorems 1 and 2. ■

5 Robust Implementation

Towards *robust implementation*, I want this paper's robustness criterion to satisfy the two requirements: First, it requires that there always exist one desirable equilibrium in "all" nearby environments. Second, it requires that no undesirable equilibria appear in "all" nearby environments. This reduces to proposing the closure of the equilibrium correspondence as a robustness test for implementation theory.

The following notation will be convenient. Fix a mechanism $\Gamma = (M, g)$. If $\mathcal{S}_\Gamma$ is a set of strategy profiles in the mechanism Γ such that for any selection f of F, there is $\sigma \in \mathcal{S}_\Gamma$ for which $g(\sigma(\omega)) = f(\xi(\omega))$ for any $\omega \in \text{supp}(\mu)$, then I will write $\mathcal{S}_\Gamma \sqsupset_\mu F$. Further, if $\mathcal{S}_\Gamma$ is a set of strategy profiles in the mechanism Γ such that $g(\sigma(\omega)) \in F(\xi(\omega))$ for each $\sigma \in \mathcal{S}_\Gamma$ and any $\omega \in \text{supp}(\mu)$, then we will write $\mathcal{S}_\Gamma \sqsubset_\mu F$. If $\mathcal{S}_\Gamma \sqsubset_\mu F$ and $\mathcal{S}_\Gamma \sqsupset_\mu F$, then we write $\mathcal{S}_\Gamma =_\mu F$.

[25] Here C1 is essential. Without C1, it is possible that $\Pi_i(\omega_\theta) \cap \text{supp}(\mu) = \emptyset$. In this case, we are unable to complete the rest of the argument. See also Kajii and Morris (1997b) in which they show by Example 4 that C1 cannot be dispensed with for the lower hemi-continuity of the Bayesian Nash equilibrium correspondence. However, in their Example 4, Kajii and Morris explicitly consider *mixed* strategies, which are ruled out in the current paper.

Definition 8 *A mechanism* Γ $[UNE]^*$*-implements an SCC* F *under a complete information prior* μ *if* $[\psi_\Gamma^U]^*(\mu) =_\mu F$*. Similarly, a mechanism* Γ $[UNE]^{**}$*-implements an SCC* F *under a complete information prior* μ *if* $[\psi_\Gamma^U]^{**}(\mu) =_\mu F$

An SCC F is said to be $[UNE]^*$-implementable if there exists a mechanism Γ such that $[\psi_\Gamma^U]^*(\mu) =_\mu F$ at any complete information prior μ. $[UNE]^{**}$-implementation can be also analogously defined.

Remark: $[UNE]^*$-implementation consists of two components: (i) $[\psi_\Gamma^U]^*(\mu) \sqsubset_\mu F$ and (ii) $[\psi_\Gamma^U]^*(\mu) \sqsubset_\mu F$. Requirement (i) says that for any selection f of F and "any" nearby environment, there always exists an equilibrium whose outcome coincides with f. Requirement (ii) says that in "all" nearby environments, every equilibrium outcome is consistent with F. The same argument can apply to $[UNE]^{**}$-implementation as well.

Palfrey and Srivastava (1991) show that almost any SCC is undominated Nash implementable. As a corollary combined with Theorems 1 and 2 of this paper, I will show that almost any SCC is $[UNE]^*$-implementable (Corollary 3). Thus, this paper's robustness test gives us a way of sustaining the permissive result of Palfrey and Srivastava (1991). By contrast, Chung and Ely (2003) show that only monotonic SCFs can be robustly undominated Nash implementable in their sense. Then, I show that Chung and Ely's robustness requirement is indeed $[UNE]^{**}$-implementability and clarify when Chung and Ely's monotonicity-as-a-necessary-condition result applies (Theorem 4). Moreover, I also extend this monotonicity-as-a-necessary-condition result to social choice *correspondences*.

5.1 How Robust is Undominated Nash Implementation?

An SCC F is said to be UNE-implementable if there exists a mechanism Γ such that $\psi_\Gamma^U(\mu) =_\mu F$ at any complete information prior μ. Nash-implementation can be analogously defined as well. To state Palfrey and Srivastava (1991)'s result on UNE implementation, I need one requirement called "no-veto-power."

Definition 9 *An SCC* F *satisfies* ***no-veto-power*** *if, for any* $\theta \in \Theta$*, whenever* $u_i(a;\theta) \geq u_i(b;\theta)$ *for any* $b \in A$ *and any* $i \in I \subseteq N$ *with* $|I| \geq n-1$ *players, we have* $a \in F(\theta)$*.*

Palfrey and Srivastava (1991) establish the following very permissive result for UNE implementation.

Theorem 3 (Palfrey and Srivastava (1991)) *Suppose that there are at least three players. Then, if an SCC* F *satisfies no-veto-power, it is UNE implementable.*

The result below is a robust version of Palfrey and Srivastava's (1991) permissive result.

Corollary 3 *Suppose that C1 holds and the set of outcomes A is a separable space. Assume further that there are at least three players and preferences are strict. Then, if an SCC F satisfies no-veto-power, it is $[UNE]^*$-implementable.*

Proof of Corollary 3: We build on a canonical mechanism proposed in the proof of Theorem 2 of Palfrey and Srivastava (1991). Since A is a separable space, we can take A^* as a countable dense subset of A and it suffices to focus on this A^* in the rest of the argument. In their canonical mechanism $\Gamma^{PS} = (M, g)$, Palfrey and Srivastava (1991) construct the following message space. For each $i \in N$, let

$$M_i = M_i^1 \times M_i^2 \times M_i^3 \times M_i^4 \times M_i^5 \text{ and } M = M_1 \times \cdots \times M_n$$

where

$$\begin{aligned} M_i^1 &= \{(a, \theta) \in A^* \times \Theta | \; a \in F(\theta)\}; \\ M_i^2 &= \Theta \\ M_i^3 &= \{-4, -3, -2, -1, 0, 1, \dots\} \\ M_i^4 &= \{0, 1, \dots\} \\ M_i^5 &= \Theta \end{aligned}$$

We omit the specification of the outcome function here. Since Θ is assumed to be finite, the message space M is at most countable. Hence, M1 holds for this mechanism. Palfrey and Srivastava (1991) show that $\emptyset \neq \psi^U_{\Gamma^{PS}}(\mu) =_\mu F$ for any complete information prior μ. By Corollary 2, we can also show that $\emptyset \neq \psi^U_{\Gamma^{PS}}(\mu) = [\psi^U_{\Gamma^{PS}}]^*(\mu) =_\mu F$ for any complete information prior μ. ■

This corollary clarifies the extent to which Palfrey and Srivastava's permissive implementation result can be sustained. I do not think that the permissive $[UNE]^*$-implementation result necessarily lend much support to the use of Palfrey and Srivastava's mechanism. Rather, this robustness test gives us a precise sense in which the Palfrey and Srivastava's mechanism is very fragile if we believe that $[UNE]^*$-implementation is a weak requirement. Indeed, I claim that the restrictiveness of d^*-elaborations is one of the reasons why undominated Nash implementation is very permissive. Because it explains how crucially Palfrey and Srivastava's result exploits the assumption of complete information.

Jackson, Palfrey, and Srivastava (1994) characterize "separable" environments within which they can design a mechanism that UNE implements *any* SCF, regardless of the number of players.[26] Hence, in separable environments with strict preferences, I conclude that any SCF is $[UNE]^*$-implementable irrespective of the number of players.

[26] See Jackson et al (1994) for the definition of separable environments.

5.2 When is Monotonicity Necessary?

The example in Section 3.1 exhibits the discontinuity of the undominated Bayesian Nash equilibrium correspondence. This result can be substantially generalized. More specifically, I will show that for *any* mechanism Γ that UNE implements a non-monotonic SCC, it follows that $[\psi_\Gamma^U]^{**}(\mu) \neq \psi_\Gamma^U(\mu)$ at any complete information prior μ. That is, monotonicity is a necessary condition for $[UNE]^{**}$-implementation. First, I give the formal definition of monotonicity.

Definition 10 *An SCC F is* ***monotonic*** *if for every pair of states θ and θ' and any $a \in F(\theta)$ such that*

$$(*)\ \ \forall i \in N,\ \forall b \in A,\ \ u_i(b;\theta') \geq u_i(a;\theta') \Longrightarrow u_i(b;\theta) \geq u_i(a;\theta),$$

we have $a \in F(\theta')$.

Maskin (1999) shows that monotonicity is a necessary and almost sufficient condition for Nash implementation. I shall show below that monotonicity is a necessary condition for $[UNE]^{**}$-implementation as well.

Theorem 4 *Suppose that preferences are strict. If an SCC F is $[UNE]^{**}$-implementable, it is necessarily monotonic.*[27]

Proof of Theorem 4: Suppose that F is a $[UNE]^{**}$-implementable SCC with the implementing mechanism $\Gamma = (M, g)$. Fix any $\theta, \theta' \in \Theta$ and any $a \in F(\theta)$. Suppose θ and θ' are two possible states satisfying $(*)$ in the condition of monotonicity (Definition 10). We will show that $a \in F(\theta')$.

Since the mechanism Γ UNE implements F by our hypothesis, there exists an undominated Nash equilibrium m^* of $\Gamma(\theta)$ such that $g(m^*) = a$. We claim that m^* is a Nash equilibrium of $\Gamma(\theta')$. If not, there must exist an player i and a message m_i such that $u_i(g(m_i, m_{-i}^*);\theta') > u_i(g(m^*);\theta')$. But by monotonic transformation $(*)$, this implies that $u_i(g(m_i, m_{-i}^*);\theta) > u_i(g(m^*);\theta)$, which is a contradicts the hypothesis that m^* is a Nash equilibrium of $\Gamma(\theta)$.

To avoid a trivial case where we can automatically conclude that $a \in F(\theta')$, we must assume that m^* is dominated in $\Gamma(\theta')$. Then let $I \subset N$ be the nonempty set of players for whom m_i^* is dominated in $\Gamma(\theta')$ for each $i \in I$. With abuse of notations, we use the expression that $|I| = I \geq 1$. We decompose the set of players of N into the following:

$$N = I \cup J = \{i_1, \ldots, i_I, j_1, \ldots, j_J\},$$

where $I = \{i_1, \ldots, i_I\}, J = \{j_1, \ldots, j_J\}$ and $I \cap J = \emptyset$. Again with abuse of notation, we use the expression that $|J| = J$.

[27] This result does not impose any restrictions on the class of mechanisms.

The proof builds on Chung and Ely's (2003) Theorem 1.[28] For $\varepsilon > 0$ sufficiently small, we construct a type space (Π, ξ) with the following properties: for $2 + I$ possible states, let $\tilde{\Omega} = \{\omega_0, \omega_1, \omega_2, \ldots, \omega_{1+I}\} \subset \Omega$; $\xi : \Omega \to \Theta$ satisfies the following: $\xi(\omega_0) = \theta'$ and $\xi(\omega_k) = \theta$ for each $k = 2, \ldots, I + 1$. For any $\tilde{\theta} \neq \theta, \theta'$, there exists $\omega_{\tilde{\theta}}$ such that $\mathcal{G}^{\tilde{\theta}} = \{\omega \in \Omega | \xi(\omega) = \tilde{\theta}\} = \{\omega_{\tilde{\theta}}\}$. That is, for any $\omega \notin \tilde{\Omega}$, the payoff state is common knowledge.

We can define the partition correspondence for each player $i \in I$ as follows:

$$\begin{aligned} \Pi_i(\omega_0) &= \Pi_i(\omega_{1+i}) = \{\omega_0, \omega_{1+i}\} \\ \Pi_i(\omega_1) &= \Pi_i(\omega_{1+k}) \text{ for all } k \in I \backslash \{i\} \\ &= \{\omega_1, \ldots, \omega_{1+I}\} \backslash \{\omega_{1+i}\} \end{aligned}$$

We can also define the partition correspondence for each player $j \in J$ as follows:

$$\begin{aligned} \Pi_j(\omega_0) &= \{\omega_0\} \\ \Pi_j(\omega_1) &= \Pi_j(\omega_2) = \cdots = \Pi_j(\omega_{1+I}) \\ &= \{\omega_1, \ldots, \omega_{1+I}\} \end{aligned}$$

Fix the above type space (Π, ξ) throughout. Let μ be a complete information prior such that, for any $\omega \in \tilde{\Omega}$, $\mu(\omega) > 0$ if and only if $\omega = \omega_0, \omega_1$. Let $\{\mu^\varepsilon\}_\varepsilon$ be an elaboration of a complete information prior μ with the following properties: $\mu^\varepsilon(\omega_0) = \mu(\omega_0); \mu^\varepsilon(\omega_1) = (1 - \varepsilon)\mu(\omega_1)$; and $\mu^\varepsilon(\omega_k) = \varepsilon\mu(\omega_1)/I$ for each $k = 2, \ldots, I$, where $\varepsilon > 0$. Note that C1 holds for this elaboration $\{\mu^\varepsilon\}$. This simply strengthens the result.

Define $\{\varepsilon^k\}_{k=1}^\infty$ such that $\varepsilon^k > 0$ for each k and $\varepsilon^k \to 0$ as $k \to \infty$. Let $\mu^k \equiv \mu^{\varepsilon^k}$ for each k. Then, $\{\mu^k\}_{k=1}^\infty$ is an elaboration of the complete information prior μ, which is the same as the one constructed in the proof of Theorem 1 of Chung and Ely (2003).[29] Next, we will show that $\{\mu^k\}_{k=1}^\infty$ is a d^{**}-elaboration of the complete information prior μ. The claim below will show that with probability $\varepsilon\mu(\omega_1)$, there is *no* approximate common knowledge at $\omega_2, \ldots, \omega_{1+I}$ about what payoff state being realized.

Lemma 3 *There is* ***no*** *approximate common knowledge at $\omega_2, \ldots, \omega_{1+I}$ about what payoff state is to be realized.*

The proof of Lemma 3 is left to the reader. Then, Lemma 3 shows that $\{\mu^k\}_{k=1}^\infty$ is a d^{**}-elaboration of the complete information prior μ. The rest of the proof is taken care of by Chung and Ely' Theorem 1. ■

Theorem 4 is a characterization of Chung and Ely's (2003) perturbation in their Theorem 1. One implication of this result is that $[UNE]^{**}$-implementation essentially

[28] The reason why I rewrite their proof is that it allows us to characterize the class of elaborations needed for their argument. Besides, the current proof can handle the case of social choice "correspondences" as well as functions.

[29] When I set $I = \{\text{Andy}\}$, and $J = \{\text{Bob}\}$ and $\mu(\omega_0) = 1 - p$ and $\mu(\omega_1) = p$, this is exactly the same elaboration constructed in Section 3.1.

implies Nash implementation. This paper and Chung and Ely (2003) do not impose any restrictions on the mechanism but rather require that the mechanism be robust to incomplete information. Jackson, Palfrey, and Srivastava (1994), on the contrary, argue that some of the power of the permissive undominated Nash implementation results derives from the fact that we have not imposed any restrictions on the implementing mechanism. They require that the implementing mechanism satisfy some "nice" properties. Specifically, they require that the implementing mechanisms be "bounded" and satisfy "the best response property." A mechanism is said to be bounded if whenever a message is weakly dominated, then it is weakly dominated by an undominated message. A mechanism is said to satisfy the best response property if each player has a best response to any message profiles of the other players in any state.[30] With these qualifications on the mechanisms, Jackson et al (1994) show that the *chained* condition, which is a non-trivial restriction but still weaker than monotonicity, is necessary for undominated Nash implementation. It is interesting to know whether or not there is an intermediate elaboration between d^* and d^{**} relative to which the chained condition is necessary for implementation in the closure of the undominated Nash equilibrium correspondence. This remains an open question.

5.3 $[NE]^{**}$-implementation

Given what I have obtained for $[UNE]^*$- and $[UNE]^{**}$-implementation, it is natural to investigate the question of $[NE]^{**}$-implementation. The current paper can offer some partial answer to this question. An SCC F is $[NE]^{**}$-implementable if there exists a mechanism Γ such that $[\psi_\Gamma^{BNE}]^{**}(\mu) =_\mu F$ at any complete information prior μ. The next result establishes a sufficiency result for $[NE]^{**}$-implementation.

Corollary 4 *Suppose that C1 holds and the set of outcomes A is a complete separable space. Assume further that there are at least three players and preferences are strict. Then, if an SCC F satisfies monotonicity and no-veto-power, it is $[NE]^{**}$-implementable.*

Remark: Corollary 4 shows that under some conditions, there is no difference between $[NE]^{**}$-implementation and the standard Nash implementation. Nevertheless, it still remains an open question whether or not one can obtain more general results for $[NE]^{**}$-implementation.

Proof of Corollary 4: We build on a canonical mechanism proposed in the proof of Theorem 3 of Maskin (1999). Since A is a separable space, we can take A^* as a countable dense subset of A and it suffices to focus on this A^* in the rest of the argument. In his canonical mechanism $\Gamma^M = (M, g)$, Maskin constructs the following message space. For each $i \in N$, let

$$M_i = M_i^1 \times M_i^2 \times M_i^3 \text{ and } M = M_1 \times \cdots \times M_n$$

[30]See their paper for the exact definitions

where $M_i^1 = \Theta$; $M_i^2 = A^*$; and $M_i^3 = \{1, 2, \dots\}$.

We omit the specification of the outcome function here. Since Θ is assumed to be finite, the message space M is at most countable. Hence, this mechanism satisfies M1 (countability). Maskin (1999) shows that if $n \geq 3$, and F satisfies monotonicity and no-veto-power, it follows that $\emptyset \neq \psi^{NE}_{\Gamma^M}(\mu) =_\mu F$ for any complete information prior μ. By Corollary 1 and Theorem 2, we can also show that $\emptyset \neq \psi^{NE}_{\Gamma^M}(\mu) = [\psi^{BNE}_{\Gamma^M}]^{**}(\mu) =_\mu F$ for any complete information prior μ. This completes the proof.■

6 Concluding Remarks

This paper proposes a robustness test for implementation results near complete information. I show that this robustness test, on the one hand, generates new restrictions imposed on the set of implementable SCCs, and on the other hand, clarifies the extent to which the permissive implementation results are sustained. I obtain two main results: (1) almost any SCC is $[UNE]^*$-implementable (Corollary 2); and (2) Only monotonic SCCs are $[UNE]^{**}$-implementable (Theorem 4 and Chung and Ely's (2003) Theorem 1). I will conclude this paper with two related topics: subgame perfect implementation and virtual implementation.

Aghion, Fudenberg, Holden, Kunimoto and Tercieux (2010) show that when the planner is only nearly certain that there is approximate common knowledge (i.e., he is subject to d^{**}-elaborations), only monotonic SCCs can be subgame perfect implemented in the closure of the sequential equilibrium correspondence. Note that this result does not need the assumption of strict preferences. Therefore, many implementation results using refinements of Nash equilibria are not robust to incomplete information. On the other hand, if one is only concerned with virtual or approximate (as opposed to exact) implementation, Kunimoto (2010) shows that, in quasilinear environments (or some generalization of it) where there are at least three players, any SCF is virtually (or approximately) implementable in iteratively undominated strategies even under almost complete information. Recall that the current paper focuses on "exact" (as opposed to virtual) implementation. Therefore, once the robustness to almost complete information is taken into account, there is a big difference between exact and virtual implementation.

7 Appendix

7.1 Topology Induced by d^*-Elaborations

Fix a type space (Π, ξ) throughout. Recall that Ω is finite. I consider the notion of the closeness of priors. Define d_0 by the rule

$$d_0(\mu, \mu') = \max_{\omega \in \Omega} |\mu(\omega) - \mu'(\omega)|.$$

Note that $d_0(\mu, \mu') = 0$ if and only if $\mu = \mu'$. Fix two priors μ and μ'. I will require extra conditions on conditional probabilities. Recall that $\mathcal{G}(\varepsilon)$ is the set of all states in which there is a common $(1-\varepsilon)$-belief about what game being played. Let

$$\begin{aligned} d_1(\mu) &= \min\{\varepsilon \mid \mu(\mathcal{G}(\varepsilon)) = 1\}, \text{ and} \\ d^*(\mu, \mu^{'}) &= \max\{d_0(\mu, \mu^{'}), d_1(\mu), d_1(\mu^{'})\}. \end{aligned}$$

Note that $d_1(\mu) = 0$ when μ is a complete information prior. By construction, d^* is non-negative and symmetric. However, it might be the case that $d^*(\mu, \mu) > 0$ unless μ is a complete information prior. This implies that d^* is not even a pseudo-metric. Therefore, I find it convenient to define a topology by specifying what nets converge to which points. I use Theorem 9 of Kelley (1955) (in p74) which shows that every convergence class is actually derived from a topology. Then, it remains to prove that any convergent net according to d^*-elaborations belongs to some convergence class.

Definition 11 *Let $\mathscr{C}$ be a class consisting of pairs (S, s), where S is a net in X and s a point. $\mathscr{C}$ is a convergence class for X if it satisfies the conditions listed below. We say that S converges $(\mathscr{C})$ to s or that $\lim_k S_k \equiv s$ $(\mathscr{C})$ if and only if $(S, s) \in \mathscr{C}$.*

1. *If S is a net such that $S_n = s$ for each n, then S converges $(\mathscr{C})$ to s.*
2. *If S converges $(\mathscr{C})$ to s, then so does each subnet of S.*
3. *If S does not converge $(\mathscr{C})$ to s, then there is a subnet of S, no subnet of which converges $(\mathscr{C})$ to s.*
4. *Let D be a directed set, let E_m be a directed set for each $m \in D$, let F be the product $D \times \prod_{m \in D} E_m$ and for $(m, f) \in F$, let $R(m, f) = (m, f(m))$. If $\lim_m \lim_n S(m, n) \equiv s$ $(\mathscr{C})$, then $S \circ R$ converges $(\mathscr{C})$ to s. Here, $S(m, n)$ is a member of a topological space for each m in D and each n in E_m.*

Let $\mathscr{C}^*$ be a class consisting of all d^*-elaborations $\{\mu^k\}_{k=1}^{\infty}$ of some complete information prior μ such that $d^*(\mu^k, \mu) \to 0$ as $k \to \infty$.

Proposition 2 *Let $\mathscr{C}^*$ be a class given above. Then, $\mathscr{C}^*$ is a convergence class for $\mathcal{P}$, where $\mathcal{P}$ is the space of all priors over Ω.*

Proof of Proposition 2: Recall that we fix a type space (Π, ξ) throughout. We must check four properties for the convergence class. Let μ be a complete information prior. Set $\{\mu^k\}_{k=1}^{\infty}$ as $\mu^k = \mu$ for each k. Then, we have that $d^*(\mu^k, \mu^*) = 0$ for each k, therefore, $(\{\mu^k\}, \mu) \in \mathscr{C}^*$. Thus, $\mathscr{C}^*$ satisfies property 1. Let $d^*(\mu^k, \mu) \to 0$ as $k \to \infty$ for some complete information prior μ, that is, $\mu^k \to \mu$ $(\mathscr{C}^*)$. It is straightforward to see that any subnet of $\{\mu^k\}$ also converges to μ $(\mathscr{C}^*)$. Hence, property 2 is satisfied for $\mathscr{C}^*$. Suppose that μ^k does not converge to μ as $k \to \infty$ according to $\mathscr{C}^*$. Then, there exists $\delta > 0$ for which there exists $\bar{k}$ such that $d^*(\mu^k, \mu) \geq \delta$ for each $k \geq \bar{k}$. Consider a subnet $\{\mu^l\}_{l=1}^{\infty} \equiv \{\mu^k\}_{k=\bar{k}}^{\infty}$. By construction, there is $\delta > 0$ such that $d^*(\mu^l, \mu) \geq \delta$ for each l. Then, it is straightforward to see that no subnet of $\{\mu^l\}_{l=1}^{\infty}$ converges to

μ according to $\mathscr{C}^*$. Thus, property 3 is satisfied for $\mathscr{C}^*$. Let a double indexed net $S(k,l) \equiv \{\{\mu^{k_l}\}_{l=1}^{\infty}\}_{k=1}^{\infty}$. Now we know by our hypothesis that for any $\varepsilon > 0$, there exist $\bar{k}$ and $\bar{l}$ such that $d^*(\mu^{k_l}, \mu) < \varepsilon$ for any $k \geq \bar{k}$ and any $l \geq \bar{l}$. Then, in order to check if property 4 is satisfied, it remains to show that, for any ε, we are able to find a member $(k, f) \in F$ such that, if $(n, g) \geq (k, f)$, then $d^*(\mu^{n_{g(n)}}, \mu) < \varepsilon$. By our hypothesis we can choose $k \in D$ so that $d^*(\lim_l \mu^{n_l}, \mu) < \varepsilon$ for any $n \geq k$. For each such n, choose a member $f(n) \in E_n$ such that $d^*(\mu^{n_l}, \mu) < \varepsilon$ for any $l \geq f(n)$. If n is a member of D which does not follow k, let $f(n)$ be an arbitrary member of E_n. If $(n, g) \geq (k, f)$, then $n \geq k$, hence $d^*(\mu^{n_l}, \mu) < \varepsilon$, and since $g(n) \geq f(n)$, we have that $d^*(\mu^{n_{g(n)}}, \mu) < \varepsilon$. Thus, property 4 is satisfied.■

Theorem 9 of Kelley (1955): Let $\mathscr{C}$ be a convergence class for a set X, and for each subset A of X, let A^c be the set of all points s such that, for some net S in A, S converges ($\mathscr{C}$) to s. Then ^{c} is a closure operator, and $(S, s) \in \mathscr{C}$ if and only if S converges to s relative to the topology associated with ^{c}.

7.2 Topology Induced by d^{**}-Elaborations

I consider a slightly coarser topology than that induced by d^*-elaborations. Fix two priors μ and $\mu^{'}$. Let

$$\tilde{d}_1(\mu) = \min \left\{ \varepsilon \middle| \mu\left(\mathcal{G}(\varepsilon)\right) \geq 1 - \varepsilon \right\}.$$

Define $d^{**}(\mu, \mu^{'})$ as follows:

$$d^{**}(\mu, \mu^{'}) = \max \left\{ d_0(\mu, \mu^{'}), \tilde{d}_1(\mu), \tilde{d}_1(\mu^{'}) \right\}.$$

Note that d^{**} is non-negative and symmetric. However, it might be the case that $d^{**}(\mu, \mu) > 0$ unless μ is a complete information prior. As I argued for d^*, this implies that d^{**} is not even a pseudo-metric. Once again, I find it convenient to define a topology via convergence class. Clearly, any convergent net according to d^*-elaborations is always a convergent net according to d^{**}. But the converse is not generally true. Let $\mathscr{C}^{**}$ be a class consisting of all d^{**}-elaborations $\{\mu^k\}_{k=1}^{\infty}$ of some complete information prior μ such that $d^{**}(\mu^k, \mu) \to 0$ as $k \to \infty$. In the previous section, I establish the equivalence between the topology induced by d^*-elaborations and the corresponding convergence class $\mathscr{C}^*$. Analogously, I can show that $\mathscr{C}^{**}$ is a convergence class for the set of information priors, $\mathcal{P}$. Applying Theorem 9 of Kelley (1955), I conclude that the convergence class $\mathscr{C}^{**}$ indeed generates a topology over $\mathcal{P}$. Furthermore, it can seen from the proof of Kelley's Theorem 9 that there is a one-to-one correspondence between the topologies for a set $\mathscr{P}^*(\mathscr{P}^{**})$ and the convergence classes on it. Therefore, the topology induced by d^{**}-elaborations is strictly coarser than the one induced by d^*-elaborations.

References

[1] P. Aghion, D. Fudenberg, R. Holden, T. Kunimoto, and O. Tercieux: "Subgame Perfect Implementation Under Value Perturbations, *mimeo*, (2010).

[2] G. Artemov, T. Kunimoto, and R. Serrano: "Robust Virtual Implementation with Incomplete Information: Towards a Reinterpretation of the Wilson Doctrine," *mimeo*, (2010).

[3] D. Bergemann and S. Morris: "Robust Mechanism Design," *Econometrica*, 73, (2005), 1771-1813.

[4] D. Bergemann and S. Morris: "Robust Implementation in General Mechanims," *mimeo*, (2010), Yale University and Princeton University.

[5] K-S. Chung and J. Ely: "Implementation with Near-Complete Information," *Econometrica*, 71, (2003), 857-871.

[6] E. Dekel and D. Fudenberg: "Rational Behavior with Payoff Uncertainty," *Journal of Economic Theory*, 52, (1990), 243-267.

[7] D. Fudenberg, D.M. Kreps, and D.K. Levine: "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44, (1988), 354-380.

[8] M.O. Jackson, T.R. Palfrey, and S. Srivastava: "Undominated Nash Implementation in Bounded Mechanisms," *Games and Economic Behavior*, 6, (1994), 474-501.

[9] M.O. Jackson and S. Srivastava: "Characterizations of Game Theoretic Solution Concepts which Lead to Impossibility Theorems," *Review of Economic Studies*, 63, (1996), 23-38.

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

[11] A. Kajii and S. Morris: "Refinements and Higher Order Beliefs: A Unified Survery," *mimeo*, (1997b).

[12] A. Kajii and S. Morris: "Payoff Continuity in Incomplete Information Games," *Journal of Economic Theory*, 82, (1998), 267-276.

[13] J.L. Kelley: "*General Topology*," (1955), D. Van Nostrand Company.

[14] T. Kunimoto: "Robust Implementation under Approximate Common Knowledge," *mimeo*, November 2005.

[15] T. Kunimoto: "Robust Virtual Implementation with Almost Complete Information, *mimeo*, (2010).

[16] E. Maskin: "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, (1999), 23-38.

[17] D. Monderer and D. Samet: "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, (1989), 170-90.

[18] D. Monderer and D. Samet: "Proximity of Information in Games with Incomplete Information," *Mathematics of Operations Research*, 21, (1996), 707-725.

[19] T. R. Palfrey and S. Srivastava: "Nash Implementation Using Undominated Strategies," *Econometrica*, 59, (1991), 479-501.