Continuous Implementation*

Marion Oury
HEC Paris

Olivier Tercieux
Paris School of Economics and CNRS

Abstract

It is well-known that mechanism design literature makes many simplifying informational assumptions in particular in terms of common knowledge of the environment among players. In this paper, we introduce a notion of continuous implementation and characterize when a social choice function is continuously implementable. More specifically, we say that a social choice function is continuously (partially) implementable if it is (partially) implementable for types in the model under study and it continues to be (partially) implementable for types "close" to this initial model. We first show that if the model is of complete information a social choice function is continuously (partially) implementable *only if* it satisfies Maskin's monotonicity. We then extend this result to general incomplete information settings and show that a social choice function is continuously (partially) implementable *only if* it is fully implementable in iterative dominance. For finite mechanisms, this condition is also sufficient. We also discuss implications of this characterization for the virtual implementation approach.

Journal of Economic Literature Classification Numbers: C79, D82

Keywords: High order beliefs, Robust implementation

*First version: February 15, 2007. This version: May 11, 2009. We wish to thank Dirk Bergemann, Philippe Jehiel, Atsushi Kajii, Eric Maskin, Nicolas Vieille and seminar participants at Paris School of Economics, the EEA meeting in Budapest, Princeton University, the Institute for Advanced Study in Princeton, Rutgers University, Mc Guil University, the world congress of game theory in Northwestern, the workshop on new topics in mechanism design in Madrid, GRIPS workshop on Global Games in Tokyo and the PSE-Northwestern joint workshop. We are especially grateful to Stephen Morris for helpful remarks and discussions at various stages of this research. Olivier Tercieux thanks the Institute for Advanced Study at Princeton for financial support through the Deutsche Bank membership. He is also grateful to the Institute for its hospitality.

1 Introduction

The notion of partial – as opposed to full – implementation consists in designing games under which *some* equilibrium – but not necessarily *all* – yields the outcome desired by the social planner. Despite the fact that undesirable equilibria may potentially exist, partial implementation is widely used both in theoretical and applied works. One of the main reasons for its success is the celebrated revelation principle: if the desired outcomes can arise as an equilibrium in some mechanism, then it will arise in a truth-telling equilibrium of the direct mechanism. It is then informally argued that truth-telling is a focal point and that agents should coordinate on this equilibrium when it exists. If the social planner is not perfectly sure that the information structure of the model he has in mind corresponds exactly to the true situation, the very notion of truth-telling strategy becomes problematic. Nevertheless, the focal point argument used to defend the partial approach extends to this (more realistic) context: if an agent's type t is very "close" to a type $\bar{t}$ of the initial model, then reporting the message that corresponds to type $\bar{t}$ may reasonably be seen as a focal point for type t.

Following this line of thought and taking into account the doubts a social planer may have about his model, we characterize when a social choice function can be partially *continuously* implemented. More specifically, we require that in any perturbation of the initial model, there exists an equilibrium that yields the desired outcome, not only at all types of the initial model but also at all types "close" to initial types. Our main results state that this continuity requirement leads to necessary (and sufficient) conditions that are tightly linked to full implementation. Otherwise stated, this paper shows that the partial implementation paradigm is very fragile when slight modifications of the information structures are allowed.

In a first step, we focus on the simple case in which the initial model is of complete information. In this specific setting, widely used in mechanism design, the approach of partial implementation is very permissive. For instance, when there are at least three agents, any social choice function can be partially implemented[1]. Let us be more specific and describe our continuity requirement in this setting. A profile of complete information types may be seen as a profile of degenerate hierarchies of beliefs where every player knows the realized state of nature, every player knows that everyone knows and so on... Put in another way, the modeler studies a set of (degenerate) hierarchies of beliefs where

[1]The mechanism allowing such a permissive result is direct. Just assume that whenever at least $n-1$ players out of n send the same state of nature, the mechanism assigns the outcome desired by the planer at this state. In any other case, the mechanism assigns some arbitrary outcome.

some given state of nature is commonly known. In our setting, a profile of incomplete information types t is considered to be close to a profile t_θ of complete information types (where it is common knowledge that the real state of nature is θ) if, t induces hierarchies of beliefs where each player believes with a high probability that payoffs are given by θ, each player believes with a high probability that each player believes with a high probability that payoffs are given by θ, etc... up to high but finite order. In our first result, we show that a social choice function is continuously implementable *only if* it satisfies Maskin's monotonicity. Many social choice functions are not monotonic and hence not continuously implementable. Since Maskin's monotonicity is necessary and (almost) sufficient for full Nash implementation in complete information settings (Maskin (1999)), this result builds a first bridge between partial and full implementation. In other words, a lack of full implementation can be problematic even if the social planer is only willing to partially implement the social choice function.

Our continuity requirement naturally extends to the case where the initial model is of incomplete information. To formalize this, we use the method introduced by Harsanyi (1967) and developed in Mertens and Zamir (1985). Each type in the initial model is mapped into a hierarchy of beliefs. Then, following the interim approach due to Weinstein and Yildiz (2007), we define a notion of "nearby" types. As already underlined in the complete information setting, this notion, formally described by the product topology in the universal type space, captures the restrictions on the modeler's ability to observe the players' (high order) beliefs. In this general setting, we provide our main result: if a social choice function is continuously implementable, then it must also be fully implementable in rationalizable messages. More precisely, we show that if some mechanism continuously implements a social choice function f, then we can extract from the initial mechanism a "smaller" mechanism that fully implements f in rationalizable messages. For finite mechanisms, this condition is also sufficient[2]. Borgers (1995) shows that full implementation in rationalizable messages is a demanding notion when considering large preference domains. However, under complete information, Bergemann and Morris (2009a) establishes a tight connection between this notion and full implementation in Nash equilibrium. Bergemann and Morris (2009b,c) provide necessary and sufficient conditions for full implementation in rationalizable messages while Bergemann and Morris (2007) studies an application to ascending auctions.

Virtual implementation corresponds to the requirement that the outcomes specified by the social choice function arise with probability arbitrarily close to –but not necessarily

[2] As will be discussed further, for infinite mechanisms, the existence of an equilibrium is not ensured and so this condition need not be sufficient.

equal to– one. Moving to virtual implementation may be seen as natural when considering continuity requirements. A corollary of our main result states that a social choice function is continuously virtually (partially) implementable with finite mechanisms if and only if it is virtually (fully) implementable in rationalizable messages with finite mechanisms. While apparently similar to our main result, this characterization is actually much less demanding. Indeed, under virtual implementation, very permissive sufficient conditions have been established by Abreu and Matsushima (1992a,b) for the solution concept of rationalizability. More precisely, in complete information settings, Abreu and Matsushima (1992a) shows that under very weak domain restrictions, if there are more than three players, any social choice function is virtually implementable in rationalizable messages. Abreu and Matsushima (1992b) extends this result to incomplete information: they show that Bayesian Incentive Compatibility and a measurability condition, which seems weak and generically satisfied[3], are both necessary and sufficient. In other terms, under the virtual approach, the gap between partial implementation (which is equivalent to Bayesian Incentive Compatibility) and full implementation in rationalizable messages is quite small. Since mechanisms used in this literature are finite, this means that when moving to virtual implementation, our continuity requirement leads to much less severe restrictions than for exact implementation. We interpret this result as a new argument in favor of the virtual approach: even if the social planer is only interested in partial implementation, considering Abreu and Matsushima's mechanisms makes sense.

Since the seminal paper by Rubinstein (1989) on the e-mail game, several approaches have been followed to analyze the connection between high order beliefs and strategic behavior; the so-called notion of *robustness* due to Kajii and Morris (1997), the *global games* argument due to Carlsson and Van Damme (1993) and the *interim approach* due to Weinstein and Yildiz (2007). These works share the common assumption that in the perturbed models, some types may have preferences that are radically different from those of types in the initial model[4]. Indeed, the behavior of these specific types is used as a starting point for contagion processes that drive results in these analyses. Note that the meaning of such an assumption in the mechanism design context (where the social planer fixes the game form) would be problematic. However, in the present paper, we show that the logic of implementation makes this assumption unnecessary. Indeed, in mechanism design, several different states of natures are ex ante possible for the social planer. Our

[3] For instance, as noted in Abreu and Matsushima (1992b) or Bergemann and Morris (2009d), a simple sufficient condition for all social choice functions to satisfy the measurability condition, is type diversity: every type has distinct preferences over lotteries unconditional on others' types.

[4] This corresponds to the notion of "crazy types" in the robustness approach and to that of "dominance regions" in global games or the interim approach.

argument in the proof uses this multiplicity and shows that this setting is then rich enough: partial implementation in the initial model is used as an (endogenous) starting point for the contagion process at equilibrium. It is then enough to assume that sending a message may involve an (arbitrarily) small cost. Since in many real economic situations sending a message is costly[5], this technical assumption[6] is in the spirit of our local requirements: mechanisms that are not robust to an arbitrarily small departure from the assumption of costless messages are rather undesirable.

Our results also contribute to the literature on the so-called "Wilson doctrine"[7]. Bergemann and Morris (2005) is one of the first attempts to relax the implicit common knowledge assumptions made in the mechanism design literature.[8] In their setting, the modeler when choosing a mechanism has no information on the real situation that will finally prevail among the agents[9]. Consequently, their notion of robust implementation follows a "global approach": a social choice function[10] is robustly (partially) implementable if it is (partially) implementable *on all possible models.* They show that a social choice function is robustly implementable if and only if it is ex-post implementable[11]. On the contrary, we assume that the planer has some specific model in mind and is quite confident about it. As a consequence, our requirement is only local: the social choice function must be implemented only at types “close” to types in the initial model. This is the reason why

[5]For instance, sending a message in real life situations may consist in filling in a questionnaire which is time-consuming and hence costly. It may sometimes require to present costly physical proofs such as observable characteristics of products, endowments... See Bull and Watson (2007) or Kartik and Tercieux (2009) for details.

[6]In case a player has several best responses against some belief, this assumption allows us to build a small perturbation of the environment where this player has a unique best response.

[7]Wilson (1987) writes "I foresee the progress of game theory as depending on successive reductions in the base of common knowledge required to conduct useful analyses of practical problems. Only by repeated weakening of common knowledge assumptions will the theory approximate reality".

[8]Another related paper is Chung and Ely (2001). They study full implementation in undominated Nash equilibrium. They show that (under hedonic preferences), while almost all social choice functions are fully implementable in undominated Nash equilibrium, only monotonic social choice functions can be fully implemented in undominated Nash if we also require that no discontinuity occurs at complete information information. There are two main differences with our work: first we focus on partial implementation, second the topology behind their continuity requirement is different from ours. See Kunimoto (2008) for additional details on their underlying topology.

[9]Alternatively, Artemov, Kunimoto and Serrano (2007) consider that the planer knows the (finite set of) first-order beliefs of the agents.

[10]Bergemann and Morris (2005) also consider social choice correspondences.

[11]Ex-post implementation requires that each agent’s strategy be optimal for every possible realization of the types of other agents. The possibility of ex-post implementation has been recently studied (see Jehiel *et al* (2006) and Bikhchandani (2006)).

ex-post implementation is not necessary in our setting.[12]

This paper is organized as follows. Section 2 analyses the complete information case. In Section 3, we extend our notion of continuous implementation to incomplete information and give our main result. We conclude with a discussion of further issues in Section 4.

2 Complete Information Case

We first introduce the complete information setting and the notion of implementation under complete information. Then, we define our main notion of continuous implementation.

2.1 Complete Information Implementation

We consider a finite set of players $\mathcal{I} = \{1, ..., I\}$. Each agent i has a bounded utility function $u_i : A \times \Theta \to \mathbf{R}$ where Θ is the finite[13] set of states of nature and A is the set of outcomes endowed with an arbitrary topology. A social choice function is a mapping $f : \Theta \to A$.[14] If the true state of nature is θ, the planner would like the outcome to be $f(\theta)$. A mechanism specifies a message set M_i for each agent i and a mapping g from message profiles to outcomes. More precisely, we write M as an abbreviation for $\prod_{i\in\mathcal{I}} M_i$ and for each player i, M_{-i} for $\prod_{j\neq i} M_j$.[15] A mechanism $\mathcal{M}$ is a pair (M, g) where the outcome function $g : M \to A$ assigns to each message profile m an alternative $g(m) \in A$. In what follows, we assume that message spaces are countable.[16] By a slight abuse of notations, we will sometimes note m for the degenerate distribution in $\Delta(M)$ assigning probability 1 to m.

For each $\theta \in \Theta$, a mechanism $\mathcal{M} = (M, g)$ induces a complete information game $\Gamma(\mathcal{M}, \theta) = [\mathcal{I}, \{M_i\}_{i\in\mathcal{I}}, \{u_i(g(.), \theta)\}_{i\in\mathcal{I}}]$ where each agent i's payoff when message profile m is sent is $u_i(g(m), \theta)$. We also denote the set of pure Nash equilibria in $\Gamma(\mathcal{M}, \theta)$ by

[12] It is not sufficient either. This is due to the fact that Bergemann and Morris (2005) use the so-called "known own payoff type" universal type space. To be more specific, they define the set of states of nature Θ by $\Theta = \Pi_{i\in\mathcal{I}}\Theta_i$ where Θ_i is the set of player i's payoff types. Then, they assume that there is common knowledge that each player i knows his payoff type θ_i.

[13] The finiteness assumption is used to prove our main result (Theorem 2) but is not needed to establish the necessity of Maskin monotonicity (Theorem 1).

[14] In the paper, we restrict our attention to social choice function for simplicity. Extensions to social choice correspondences will be discussed further.

[15] Similar abbreviations will be used throughout the paper for analogous objects.

[16] As will become clear, this assumption will allow us to prove our necessary conditions for continuous implementation using only models with a countable set of types. When moving to sufficient conditions, having models with countable set of types will be useful to apply standard existence theorems; see footnote 16.

$NE(\mathcal{M},\theta) = \{m \in M :$ for each i, $u_i(g(m_i,m_{-i}),\theta) \geq u_i(g(m'_i,m_{-i}),\theta)$ for all $m'_i \in M_i\}$. Implementation literature (in particular partial implementation) often focuses on the equilibrium concept of pure (and not mixed) Nash. We recall the definitions of partial and full implementation under complete information.

Definition 1 *A social choice function $f : \Theta \to A$ is partially implementable if there exists a mechanism $\mathcal{M}$ such that for each θ, there exists $m^* \in NE(\mathcal{M},\theta)$ such that $g(m^*) = f(\theta)$.*

Definition 2 *A social choice function $f : \Theta \to A$ is fully implementable if there exists a mechanism $\mathcal{M}$ such that for each θ, $NE(\mathcal{M},\theta) \neq \emptyset$ and for any $m^* \in NE(\mathcal{M},\theta)$ we have: $g(m^*) = f(\theta)$.*

2.2 Continuous Implementation

To define our notion of continuous implementation,we embed the complete information setting in a richer setting that allows to perturb high order beliefs. We also relax the assumption that sending a message is perfectly costless.

Small costs of messages

We assume that sending a message may be slightly costly. Indeed, sending a message usually requires to fill in a questionnaire, to write a letter and sometimes to present costly physical proofs such as observable characteristics of goods, endowments... A recent literature in implementation takes into account costs of messages[17]. We believe that a mechanism implementing a social choice function should still implement it when we allow for slight departures from the assumption of costless messages. In order to formalize this idea, we proceed as follows.

Given a mechanism $\mathcal{M} = (M,g)$, for each player i, we define a cost function $c_i : M_i \times \tilde{\Theta} \to \mathbb{R}_+$ where $\tilde{\Theta}$ is the space of states of nature associated with costs of messages. We assume that the state space $\tilde{\Theta}$ is rich enough. More precisely, it is defined by

$$\tilde{\Theta} = \bigcup_{i \in \mathcal{I}} \bigcup_{m_i \in M_i} \{\tilde{\theta}^{m_i}\} \cup \{\tilde{\theta}^0\}$$

where for each player i and each message m_i, we have $c_i(m_i, \tilde{\theta}^0) = 0$, $c_i(m_i, \tilde{\theta}^{m_i}) = 0$ and $c_i(m'_i, \tilde{\theta}^{m_i}) = \eta$ for all $m'_i \neq m_i$, where η is a strictly positive parameter that can be chosen arbitrarily close to 0. When no confusion arises, we will omit the dependence with respect to η. Note that since M has been assumed to be countable, $\tilde{\Theta}$ is also countable. Next, we write $\Theta^* = \Theta \times \tilde{\Theta}$ for the extended set of states of nature. For a given state of

[17] See for instance Bull and Watson (2007), Deneckere and Severinov (2008), Matsushima (2008) and Kartik and Tercieux (2009).

nature $\theta^* = (\theta, \tilde{\theta}) \in \Theta^*$, the utility function of player i for a given message m_i and a given outcome a is $u_i(a, \theta) - c_i(m_i, \tilde{\theta})$. For notational convenience, we will sometimes identify θ and $(\theta, \tilde{\theta}^0)$.

Technically, the above construction is used to break ties. More precisely, if a type is indifferent between several messages, we can slightly perturb his information so that this type has a unique best reply. Our basic point is that we want to allow a rich enough uncertainty such that for each message there exists a state of nature where the cost of this message is smaller than the other messages' costs. Note that this assumption is reminiscent of the richness assumption assumed in Weinstein and Yildiz (2007). However, it is much weaker. Indeed, the richness assumption states that for any player i and any message m_i, there exists a state nature where m_i is strictly dominant for player i.

Models

There are two main classes of situations with incomplete information. The first one consists in situations with an ex ante stage during which each player observes a private signal about the payoffs, and the joint distribution of signals and payoffs is commonly known. These situations are naturally modelled using a standard type space. The second class, on which we focus in this paper, consists in genuine situations of incomplete information, i.e. situations with no ex ante stage: each player begins with a hierarchy of beliefs. We follow the standard Harsanyi (1967)'s approach and model these hierarchies of beliefs by introducing a hypothetical ex ante stage leading to a standard type space. This allows us to study strategic behavior of players at types that are considered to be close to a given original model.

A model $\mathcal{T}$ is a pair (T, κ) where $T = T_1 \times ... \times T_I$ is a countable[18] type space and $\kappa_{t_i} \in \Delta(\Theta^* \times T_{-i})$ denotes the associated beliefs for each $t_i \in T_i$. Given a mechanism $\mathcal{M}$ and a model $\mathcal{T}$, we write $U(\mathcal{M}, \mathcal{T})$ for the induced incomplete information game. In this game, a (behavioral) strategy of a player i is any measurable function $\sigma_i : T_i \to \Delta(M_i)$. We will note $\sigma_i(m_i \mid t_i)$ for the probability that strategy σ_i assigns to message m_i when player i is of type t_i. For each $i \in \mathcal{I}$ and for each belief $\pi_i \in \Delta(\Theta^* \times M_{-i})$, set

$$BR_i(\pi_i \mid \mathcal{M}) = \arg\max_{m_i \in M_i} \sum_{(\theta, \tilde{\theta}, m_{-i}) \in \Theta^* \times M_{-i}} \pi_i(\theta, \tilde{\theta}, m_{-i}) \left[u_i(g(m_i, m_{-i}), \theta) - c_i(m_i, \tilde{\theta}) \right].$$

Given any type t_i and any strategy profile σ_{-i}, we write $\pi_i(\cdot \mid t_i, \sigma_{-i}) \in \Delta(\Theta^* \times M_{-i})$ for the joint distribution on the underlying uncertainty and the other players' messages induced by t_i and σ_{-i}.

[18] This assumption is just made to ensure existence of Bayes Nash equilibrium in finite games which will turn out to be useful when we deal with sufficient conditions for continuous implementation.

Definition 3 *A profile of strategies* $\sigma = (\sigma_1, ..., \sigma_I)$ *is a Bayes Nash equilibrium in* $U(\mathcal{M}, \mathcal{T})$ *if for each* $i \in \mathcal{I}$ *and for each* $t_i \in T_i$,

$$m_i \in Supp(\sigma_i(t_i)) \Rightarrow m_i \in BR_i(\pi_i(\cdot \mid t_i, \sigma_{-i}) \mid \mathcal{M}).$$

As specified earlier, types close to complete information are types where it is mutually believed (with arbitrarily large probability) up to an arbitrarily large order that a given complete information situation occurred. More formally, given a model (T, κ) and any type t_i in type space T_i, we can compute the belief of t_i on Θ^* (i.e. his *first-order belief*) by

$$h_i^1(t_i) = \text{marg}_{\Theta^*}\kappa_{t_i}$$

We can compute the second-order belief of t_i, i.e. his beliefs about $(\theta^*, h_1^1(t_1), ..., h_I^1(t_I))$, by setting

$$h_i^2(t_i) = \kappa_{t_i}(\{(\theta^*, t_{-i}) \mid (\theta^*, h_1^1(t_1), ..., h_I^1(t_I)) \in F\})$$

for each measurable $F \subseteq \Theta^* \times \Delta(\Theta^*)^I$. We can compute an entire hierarchy of beliefs by proceeding in this way. Hence, a type of a player i induces an infinite hierarchy of beliefs $(h_i^1(t_i), h_i^2(t_i), ..., h_i^k(t_i), ...)$ where $h_i^1(t_i) \in \Delta(\Theta^*)$ is a probability distribution on Θ^*, representing the beliefs of i about θ^*, $h_i^2(t_i) \in \Delta(\Theta^* \times \Delta(\Theta^*)^I)$ is a probability distribution representing the beliefs of i about θ^* and the other first order beliefs. Let us write $h_i(t_i)$ for the resulting hierarchy and $h_i^k(t_i)$ for the kth-order beliefs of type t_i.

The set of all belief hierarchies for which it is common knowledge that the beliefs are coherent (i.e., each player knows his beliefs and his beliefs at different orders are consistent with each other) is the universal type space (see Mertens and Zamir (1985) and Brandenburger and Dekel (1993)). We denote by $\mathcal{T}_i^*$ the set of player i's hierarchies of belief in this space and write $\mathcal{T}^* = \prod_{i \in \mathcal{I}} \mathcal{T}_i^*$.

In our formulation, two types t and $\bar{t}$ are "close" if there exists a sufficiently "large" k such that for each $l \leq k$, the lth-order beliefs $h^l(t)$ and $h^l(\bar{t})$ are close in the topology of convergence of measures. To be more precise, each $\mathcal{T}_i^*$ is endowed with the product topology, so that a sequence of types $\{t_i[n]\}_{n=0}^{\infty}$ converges to a type t_i, if, for each $k : h_i^k(t_i[n]) \to h_i^k(t_i)$ (i.e. $h_i^k(t_i[n])$ converges toward $h_i^k(t_i)$ in the topology of weak convergence of measures[19]). In such a case, we write $t_i[n] \to_P t_i$. We will sometimes use the metric $d^k(.,.)$ on the kth level beliefs[20] that metrizes the topology of weak convergence of measures.

[19] Recall that $h_i^k(t_i[n]) \in \Delta(X_{k-1})$ where $X_0 = \Theta^*$ and $X_k = [\Delta(X_{k-1})]^I \times X_{k-1}$.

[20] I.e. on $\Delta(X_{k-1})$ – see the previous footnote. One such metric is the Prokhorov metric; see Section 4.2.

We now introduce our main notion of continuous implementation. We first define the complete information model $\mathcal{T}^{CI} = (T^{CI}, \kappa)$ as follows. For each player i, $T_i^{CI} = \bigcup_{\theta \in \Theta} \{t_{i,\theta}\}$ and $\kappa_{t_{i,\theta}} = \delta_{(\theta, \tilde{\theta}^0, t_{-i,\theta})}$ where δ_x denotes the probability distribution that puts probability 1 on $\{x\}$. It is easily checked that $h_i(t_{i,\theta})$ is the hierarchy of player i's beliefs corresponding to common knowledge[21] of $(\theta, \tilde{\theta}^0)$. We will henceforth call the type $t_{i,\theta}$ a complete information type.

The modeler is interested in strategic behavior of types where players mutually believe (with high probability) up to a high (but finite) level that the state of nature is some $\theta \in \Theta$. Our continuity requirement will ensure that the social choice function is implemented not only at complete information types but also at types that are so close to the complete information situation that they cannot be ruled out by the modeler.

In what follows, for two models $\mathcal{T} = (T, \kappa)$ and $\mathcal{T}' = (T', \kappa')$ we will note $\mathcal{T} \supset \mathcal{T}'$ if $T \supset T'$ and for all $i, t_i \in T_i' : \kappa_{t_i} = \kappa'_{t_i}$.

Definition 4 *Fix a mechanism $\mathcal{M}$ and a model $\mathcal{T} \supset \mathcal{T}^{CI}$. We say that an equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ continuously implements f if for each $t_\theta \in T^{CI}$, (i) $\sigma(t_\theta)$ is pure and (ii) for any sequence $t[n] \to_P t_\theta$ where for each $n : t[n] \in T$, we have $g \circ \sigma(t[n]) \to f(\theta)$.*

Notice that point (i) maintains the requirement of pure strategy behavior usually assumed in implementation theory; this will allow for simple comparisons with existing results[22]. We now state a formal definition of continuous implementation.

Definition 5 *A social choice function $f : \Theta \to A$ is continuously implementable if there exists a mechanism $\mathcal{M}$, such that for any model $\mathcal{T} \supset \mathcal{T}^{CI}$, there is a Bayes Nash equilibrium σ in the induced game $U(\mathcal{M}, \mathcal{T})$ which continuously implements f.*

2.3 Monotonicity as a Necessary Condition

In this section, we show that any social choice function that is continuously implementable satisfies the well-known monotonicity condition as defined in Maskin (1999). This result, which is a first step toward our main result, reduces the gap between partial and full implementation since – as proved by Maskin – this monotonicity condition is necessary and "almost" sufficient for full implementation[23].

Let us first recall the definition of monotonicity for social choice functions.

[21] In this paper, we do not distinguish common knowledge and common belief.

[22] This assumption is dispensible for our main result (Theorem 2). In addition, provided that the definition of monotonicity is extended to lotteries, Theorem 1 also extends.

[23] Maskin (1999) showed that with more than three players together with the assumption that f satisfies the weak condition of No Veto Power, monotonicity actually implies full implementation.

Definition 6 *A social choice function f is monotonic if for every pair of states θ and θ' such that for each player i and for each $a \in A$,*

$$u_i(a,\theta) \leq u_i(f(\theta),\theta) \Rightarrow u_i(a,\theta') \leq u_i(f(\theta),\theta'), \qquad (\star)$$

we have $f(\theta) = f(\theta')$.

We now state the main theorem of this section.

Theorem 1 *A social choice function is continuously implementable only if it is monotonic.*

Proof of Theorem 1. Assume that there exists a mechanism $\mathcal{M} = (M, g)$ that continuously implements f. Pick $\theta, \theta' \in \Theta$ such that for each player i and for each $a \in A$, the relation $(\star)$ is satisfied. We want to show that $f(\theta) = f(\theta')$.

We show that there exists a model $\mathcal{T} = (T, \kappa)$ such that for any equilibrium σ that continuously implements f, there is a sequence of types $\{t[n]\}_{n=1}^{\infty}$ in T such that $t[n] \to_P t_{\theta'}$ and $g \circ \sigma(t[n]) \to f(\theta)$. By point (ii) of Definition 4: $g \circ \sigma(t[n]) \to f(\theta')$, which implies $f(\theta) = f(\theta')$.

For this purpose, we build the desired model (T, κ) in which for each player i, each set T_i satisfies:

$$T_i = T_i^{CI} \cup \left(\bigcup_{k=1}^{\infty} \bigcup_{m \in M} t_i(k,m) \right)$$

where $t_i(k,m)$ and κ are defined recursively as follows. For each $m = (m_1, ..., m_I) \in M$: $t_i(1,m)$ is such that

$$\text{marg}_{T_{-i}} \kappa_{t_i(1,m)}(t_{-i,\theta}) = 1,$$

$$\text{marg}_{\Theta} \kappa_{t_i(1,m)}(\theta') = 1,$$

and

$$\text{marg}_{\tilde{\Theta}} \kappa_{t_i(1,m)}(\tilde{\theta}^{m_i}) = 1.$$

In addition, for each $k \geq 2$, $t_i(k,m)$ is defined by

$$\text{marg}_{T_{-i}} \kappa_{t_i(k,m)}(t_{-i}(k-1,m)) = 1,$$

$$\text{marg}_{\Theta} \kappa_{t_i(k,m)}(\theta') = 1,$$

$$\text{marg}_{\tilde{\Theta}} \kappa_{t_i(k,m)}(\tilde{\theta}^{0}) = 1 - \frac{1}{k}, \text{ and, } \text{marg}_{\tilde{\Theta}} \kappa_{t_i(k,m)}(\tilde{\theta}^{m}) = \frac{1}{k}.$$

Observe that since M has been assumed to be countable, each T_i is countable, and so is T.

Now pick any equilibrium σ of the induced game $U(\mathcal{M}, \mathcal{T})$ that continuously implements f. By point (i) in Definition 4, $\sigma(t_\theta)$ is a pure Nash equilibrium in the complete information game $\Gamma(\mathcal{M}, \theta)$. In addition, point (ii) in Definition 4 implies that $g \circ \sigma(t_\theta) = f(\theta)$. In the sequel, we note $m_i^* := \sigma_i(t_{i,\theta})$ and $m^* := \sigma(t_\theta)$. We have for any player i and $m_i' \in M_i$:

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

and so,

$$u_i(f(\theta), \theta) \geq u_i(g(m_i', m_{-i}^*), \theta).$$

By $(\star)$, this implies that

$$u_i(f(\theta), \theta') \geq u_i(g(m_i', m_{-i}^*), \theta'),$$

which in turn implies

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

i.e. m^* is a pure Nash equilibrium in $\Gamma(\mathcal{M}, \theta')$. Otherwise stated, for each player i : $m_i^* \in BR_i(\delta_{(\theta', \tilde{\theta}^0, m_{-i}^*)} \mid \mathcal{M})$; in addition, it is easily checked that $\pi_i(\cdot \mid t_i(1, m^*), \sigma_{-i}) = \delta_{(\theta', \tilde{\theta}^{m_i^*}, m_{-i}^*)}$. Consequently, $\{m_i^*\} = BR_i(\delta_{(\theta', \tilde{\theta}^{m^*}, m_{-i}^*)} \mid \mathcal{M})$; and since σ is an equilibrium:

$$\sigma_i(t_i(1, m^*)) = m_i^*.$$

Using a similar reasoning, it is easy to show inductively that for all $k \geq 2$

$$\sigma_i(t_i(k, m^*)) = m_i^*.$$

This means that for each $k \geq 1$, $g \circ \sigma(t(k, m^*)) = g(m^*) = f(\theta)$ and so obviously, $g \circ \sigma(t(k, m^*)) \to f(\theta)$ (as $k \to \infty$) which completes the proof since $t(k, m^*) \to_P t_{\theta'}$ (as $k \to \infty$). ■

One may think that our strong result stems from lack of a common prior in the model we build. However, it is possible to slightly perturb the conditional beliefs of types in our model so that using an argument due to Lipman (2003, 2005), these types could be picked from models where players share a common prior.

Our paper focuses on social choice functions; for social choice correspondences, two definitions of partial implementation coexist. To be more specific, in the first definition, which is "weak", a social choice correspondence $F : \Theta \rightrightarrows A$ is partially implementable if there exists a mechanism $\mathcal{M}$ and a selection f of F such that the mechanism $\mathcal{M}$ partially implements f. In the second definition, which is "strong", a social choice correspondence $F : \Theta \rightrightarrows A$ is partially implementable if there exists a mechanism $\mathcal{M}$ that partially implements *each* selection f of F. Maskin (1999) gives a definition of Maskin monotonicity for social choice correspondences[24] and shows that under the same conditions as for social choice functions, this notion implies "strong" full implementation, i.e. that there exists a mechanism $\mathcal{M}$ such that for each $\theta : F(\theta) = g(NE(\mathcal{M}, \theta))$; note that the mechanism $\mathcal{M}$ partially implements each selection of F. Using a "strong" definition of continuous partial implementation, Theorem 1 can easily be extended to social choice correspondences.

It is clear that if a mechanism continuously implements a social choice function, then it must partially implement in NE this social choice function. If we add the requirement that it must partially implement in strict NE, then we may dispense with the assumption on cost of messages. In this case, a necessary condition would be the strict monotonicity condition (i.e. where the inequalities in the definition of monotonicity are replaced by strict inequalities).

3 Incomplete Information

So far, we focused our attention on situations where the planer has a complete information setting in mind. We now relax this assumption and consider the general case where the initial model of the planer is (potentially) an incomplete information one. As before, the modeler wants to see how strategic behavior is affected under his mechanism when the assumption that his model is common knowledge is relaxed. We now move to our main result which establishes a tight connection between our notion of continuous implementation and full implementation in rationalizable messages.

3.1 Definitions

We first extend the definition of continuous implementation to an incomplete information setting. In the sequel, we fix a finite[25] model $\bar{\mathcal{T}} = (\bar{T}, \bar{\kappa})$ which is the model the planer has in mind.

Definition 7 *Fix a mechanism $\mathcal{M}$ and a model $\mathcal{T} \supset \bar{\mathcal{T}}$, we say that an equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ continuously implements f if for each $\bar{t} \in \bar{T}$, (i) $\sigma(\bar{t})$ is pure and (ii) for any sequence $t[n] \to_P \bar{t}$ where for each $n : t[n] \in T$, we have $g \circ \sigma(t[n]) \to f(\bar{t})$.*

[24] A social choice correspondence is monotonic if for all θ, $b \in F(\theta)$ and any θ' such that for each player i and $a \in A$

$$u_i(a, \theta) \leq u_i(b, \theta) \Rightarrow u_i(a, \theta') \leq u_i(b, \theta')$$

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

[25] The finiteness assumption allows to prove Theorem 2 using only models that are countable which again will be useful when moving to sufficiency results.

Definition 8 *A social choice function $f : \bar{T} \to A$ is continuously implementable if there exists a mechanism $\mathcal{M}$, such that for each model $\mathcal{T} \supset \bar{\mathcal{T}}$, there is a Bayes Nash equilibrium σ in the induced game $U(\mathcal{M}, \mathcal{T})$ which continuously implements f.*

3.2 Necessary Condition

Our characterization result relies on the notion of full implementation in rationalizable messages. First, let us recall the definition of (interim correlated) rationalizability given in Dekel, Fudenberg and Morris (2006, 2007). Pick any profile of types t drawn from some arbitrary model $\mathcal{T} = (T, \kappa)$. For each i and t_i, set $R_i^0(t_i \mid \mathcal{M}, \mathcal{T}) = M_i$, and define for any integer $k > 0$ the sets $R_i^k(t_i \mid \mathcal{M}, \mathcal{T})$ iteratively, by

$$R_i^k\left[t_i \mid \mathcal{M}, \mathcal{T}\right] = \left\{ m_i \in M_i \;\middle|\; \begin{array}{c} m_i \in BR_i(\mathrm{marg}_{\Theta^* \times M_{-i}} \pi_i \mid \mathcal{M}) \text{ for some } \pi_i \in \Delta(\Theta^* \times T_{-i} \times M_{-i}) \\ \text{where } \mathrm{marg}_{\Theta^* \times T_{-i}} \pi_i = \kappa_{t_i} \text{ and} \\ \pi_i(\theta^*, t_{-i}, m_{-i}) > 0 \Longrightarrow m_{-i} \in R_{-i}^{k-1}(t_{-i} \mid \mathcal{M}, \mathcal{T}) \end{array} \right\}.$$

where $R_{-i}^{k-1}(t_{-i} \mid \mathcal{M}, \mathcal{T})$ stands for $\prod_{j \neq i} R_j^{k-1}(t_j \mid \mathcal{M}, \mathcal{T})$. The set of all rationalizable messages for player i (with type t_i) is

$$R_i^\infty(t_i \mid \mathcal{M}, \mathcal{T}) = \bigcap_{k=0}^{\infty} R_i^k(t_i \mid \mathcal{M}, \mathcal{T}) \text{ and } R^\infty(t \mid \mathcal{M}, \mathcal{T}) = \prod_{i=1}^{I} R_i^\infty(t_i \mid \mathcal{M}, \mathcal{T}).$$

Remark 1 *Lipman (1994) gives an alternative definition of rationalizability for the case of countable action sets. While his definition is consistent with common knowledge of rationality, the one we use in this paper is a coarser solution concept. Using a coarser solution concept strengthens our necessary condition; this condition will remain valid under any finer notion of rationalizability. Our sufficiency results will be proved for finite mechanisms where both concepts coincide.*

We say that a social choice function is *fully implementable in rationalizable messages*, or simply *fully rationalizable implementable*, if there is a mechanism $\mathcal{M}$ so that at each profile of types $\bar{t} \in \bar{T} : m \in R^\infty(\bar{t} \mid \mathcal{M}, \bar{\mathcal{T}}) \Rightarrow g(m) = f(\bar{t})$. In the sequel, for two mechanisms $\mathcal{M} = (M, g)$ and $\mathcal{M}' = (M', g')$, we write $\mathcal{M}' \subset \mathcal{M}$ if $M' \subset M$ and $g_{|M'} = g'$ where $g_{|M'}$ denotes the restriction of g to M'.

Our main theorem is stated as follows.

Theorem 2 *A social choice function $f : \bar{T} \to A$ is continuously implementable with a mechanism $\mathcal{M}$ only if it is fully rationalizable implementable by some mechanism $\mathcal{M}' \subset \mathcal{M}$.*

The above result states a necessary condition for continuous implementation. When moving to sufficiency, existence of equilibria becomes an issue. Given that the set of messages may be infinite, there may exist models where no equilibrium exists even if f is fully rationalizable implementable (in the model $\bar{\mathcal{T}}$). However, this condition will turn to be sufficient when considering finite mechanisms.

Note that if f is continuously implementable, then it is partially implementable. Hence, the previous result shows that full implementation in (Bayes) Nash equilibrium is a necessary condition for continuous implementation. Jackson (1991) has extended Maskin's monotonicity to incomplete information settings. He defines Bayesian monotonicity and shows that this notion is a necessary condition for full implementation in Nash equilibria in incomplete information settings. Hence, as a corollary of the above result, we get that Bayesian monotonicity is also necessary for continuous implementation which generalizes our Theorem 1.[26]

Finally, using the "weak" definition of partial implementation for social choice correspondences given in Section 2.3, it is possible to extend Theorem 2 and to establish that a necessary condition for F to be (weakly) continuously implementable is that F must be (weakly) fully implementable in rationalizable messages (i.e. there is a mechanism $\mathcal{M}$ and a selection f of F such that $\mathcal{M}$ full implements f in rationalizable messages.)[27]

Let us move now to the proof of Theorem 2. Since f is continuously implementable, there exists a mechanism $\mathcal{M} = (M, g)$, such that for any model $\mathcal{T} = (T, \kappa)$, there is a Bayes Nash equilibrium σ in the induced game $U(\mathcal{M}, \mathcal{T})$ where for each $\bar{t} \in \bar{T}$, (i) $\sigma(\bar{t})$ is pure and (ii) for any sequence $t[n] \to_P \bar{t}$ where for each $n : t[n] \in T$, we have $g \circ \sigma(t[n]) \to f(\bar{t})$. We let $\bar{\Sigma}$ be the set of pure Bayesian Nash equilibria of $U(\mathcal{M}, \bar{\mathcal{T}})$. Note that because $\bar{T}$ is finite and M is countable, $\bar{\Sigma}$ is countable. For each $\bar{\sigma} \in \bar{\Sigma}$, we build the set of message profiles $M(\bar{\sigma})$ in the following way.

For each player i and each positive integer k, we define inductively $M_i^k(\bar{\sigma})$. First, we set $M_i^0(\bar{\sigma}) = \bar{\sigma}_i(\bar{T}_i)$. Then, for each $k \geq 1:$

$$M_i^{k+1}(\bar{\sigma}) = BR_i(\Delta(\Theta \times \{\tilde{\theta}^0\} \times M_{-i}^k(\bar{\sigma})) \mid \mathcal{M}).$$

Recall that in the model $\bar{\mathcal{T}} = (\bar{T}, \bar{\kappa})$, $\text{marg}_{\tilde{\Theta}} \bar{\kappa}_{\bar{t}_i}(\tilde{\theta}^0) = 1$, for each i and $\bar{t}_i \in \bar{T}_i$. Since $\bar{\sigma}$ is an equilibrium in $U(\mathcal{M}, \bar{\mathcal{T}})$, $M_i^0(\bar{\sigma}) = \bar{\sigma}_i(\bar{T}_i) \subset BR_i(\Delta(\Theta \times \{\tilde{\theta}^0\} \times M_{-i}^0(\bar{\sigma})) \mid \mathcal{M}) = M_i^1(\bar{\sigma})$. Consequently, it is clear that for each $k : M_i^k(\bar{\sigma}) \subset M_i^{k+1}(\bar{\sigma})$. Finally, set $M_i(\bar{\sigma}) = \lim_{k \to \infty} M_i^k(\bar{\sigma}) = \bigcup\limits_{k \in \mathbb{N}} M_i^k(\bar{\sigma})$. In the sequel, for each $\bar{\sigma} \in \bar{\Sigma}$, we will note $\mathcal{M}(\bar{\sigma})$ the mechanism $(M(\bar{\sigma}), g_{|M(\bar{\sigma})})$.

[26] Bergemann and Morris (2009b) define the notion of interim rationalizable monotonicity which is necessary for full rationalizable implementation. Clearly, Theorem 1 implies that interim rationalizable monotonicity is necessary for continuous implementation.

[27] For continuous "strong" partial implementation, we believe that Bayesian monotonicity as defined by Jackson (1991) is necessary.

Notice that given any model $\mathcal{T} = (T, \kappa)$ such that $\mathcal{T} \supset \bar{\mathcal{T}}$, $\bar{\mathcal{T}}$ is a belief closed subspace in $\mathcal{T}$, i.e., for any i and $\bar{t}_i \in \bar{T}_i$: $\text{marg}_{T_{-i}}\ \kappa_{\bar{t}_i}(\bar{T}_{-i}) = 1$. Hence, for any model $\mathcal{T} \supset \bar{\mathcal{T}}$ and any equilibrium σ in $U(\mathcal{M}, \mathcal{T})$, the restriction of σ to $\bar{T}$ – denoted $\sigma_{|\bar{T}}$ – is an equilibrium in $U(\mathcal{M}, \bar{\mathcal{T}})$. A first interesting property of the family of sets $\{M(\bar{\sigma})\}_{\bar{\sigma}\in\bar{\Sigma}}$ is as follows: there is a model $\mathcal{T} \supset \bar{\mathcal{T}}$ for which any equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ has full range in $M(\sigma_{|\bar{T}})$ i.e. each message profile in $M(\sigma_{|\bar{T}})$ is played under σ at some profile of types in the model $\mathcal{T}$. More precisely, Proposition 1 is the first step of the proof of Theorem 2.

Proposition 1 *There exists a model $\mathcal{T}$ =(T, κ) such that for any $\bar{\sigma} \in \bar{\Sigma}$ and $m \in M(\bar{\sigma})$, there exists $t[\bar{\sigma}, m] \in T$ s.t. $\sigma(t[\bar{\sigma}, m]) = m$ for any equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ s.t. $\sigma_{|\bar{T}} = \bar{\sigma}$.*

Proof. We build the model $\mathcal{T}$ =(T, κ) as follows. For each equilibrium $\bar{\sigma} \in \bar{\Sigma}$, player i and integer k , we define inductively $t_i[\bar{\sigma}, k, m_i]$ for each $m_i \in M_i^k(\bar{\sigma})$ and set

$$T_i = \bigcup_{\bar{\sigma}\in\bar{\Sigma}} \bigcup_{k=1}^{\infty} \bigcup_{m_i\in M_i^k(\bar{\sigma})} t_i[\bar{\sigma}, k, m_i] \cup \bar{T}_i$$

Note that T_i is countable. In the sequel, we fix an arbitrary $\bar{\sigma} \in \bar{\Sigma}$. This equilibrium $\bar{\sigma}$ is sometimes omitted in our notations.

For each $k \geq 1$ and $m_i \in M_i^k(\bar{\sigma})$, we know that there exists $\pi_i^{k,m_i} \in \Delta(\Theta \times \{\tilde{\theta}^0\} \times M_{-i}^{k-1}(\bar{\sigma}))$ such that $m_i \in BR_i(\pi_i^{k,m_i})$. Thus we can build $\hat{\pi}_i^{k,m_i} \in \Delta(\Theta \times \tilde{\Theta} \times M_{-i}^{k-1}(\bar{\sigma}))$ such that

$$\text{marg}_{\Theta\times M_{-i}^{k-1}(\bar{\sigma})}\ \hat{\pi}_i^{k,m_i} = \text{marg}_{\Theta\times M_{-i}^{k-1}(\bar{\sigma})}\ \pi_i^{k,m_i}$$

while $\text{marg}_{\tilde{\Theta}}\hat{\pi}_i^{k,m_i} = \delta_{\tilde{\theta}^{m_i}}$. Note that $BR_i(\hat{\pi}_i^{k,m_i} \mid \mathcal{M}) = \{m_i\}$.

In the sequel, for each player i and $m_i \in M_i^0(\bar{\sigma})$, we pick one type denoted $t_i[\bar{\sigma}, 0, m_i]$ in $\bar{T}_i$ satisfying $\bar{\sigma}_i(t_i[\bar{\sigma}, 0, m_i]) = m_i$. This is well-defined because by construction, $M_i^0(\bar{\sigma}) = \bar{\sigma}_i(\bar{T}_i)$. Now, for each $m_i \in M_i^1(\bar{\sigma})$, we let $t_i[\bar{\sigma}, 1, m_i]$ be defined by[28]

$$\kappa_{t_i[\bar{\sigma},1,m_i]}(\theta, \tilde{\theta}, t_{-i}) = \begin{cases} 0 \text{ if } t_{-i} \neq t_{-i}[\bar{\sigma}, 0, m_{-i}] \text{ for each } m_{-i} \in M_{-i}^0(\bar{\sigma}) \\ \hat{\pi}_i^{1,m_i}(\theta, \tilde{\theta}, m_{-i}) \text{ if } t_{-i} = t_{-i}[\bar{\sigma}, 0, m_{-i}] \text{ for some } m_{-i} \in M_{-i}^0(\bar{\sigma}) \end{cases}.$$

This probability measure is well-defined since $\hat{\pi}_i^{1,m_i}(\Theta \times \tilde{\Theta} \times M_{-i}^0(\bar{\sigma})) = 1$. In the same way, for each $k > 1$ and $m_i \in M_i^k(\bar{\sigma})$, we define inductively $t_i[\bar{\sigma}, k, m_i]$ by:

$$\kappa_{t_i[\bar{\sigma},k,m_i]}(\theta, \tilde{\theta}, t_{-i}) = \begin{cases} 0 \text{ if } t_{-i} \neq t_{-i}[\bar{\sigma}, k-1, m_{-i}] \text{ for each } m_{-i} \in M_{-i}^{k-1}(\bar{\sigma}) \\ \hat{\pi}_i^{k,m_i}(\theta, \tilde{\theta}, m_{-i}) \text{ if } t_{-i} = t_{-i}[\bar{\sigma}, k-1, m_{-i}] \text{ for some } m_{-i} \in M_{-i}^{k-1}(\bar{\sigma}) \end{cases}.$$

Again, this probability measure is well-defined since $\hat{\pi}_i^{k,m_i}(\Theta \times \tilde{\Theta} \times M_{-i}^{k-1}(\bar{\sigma})) = 1$.

[28] Here again, we abuse notations and write $t_{-i}[\bar{\sigma}, 0, m_{-i}]$ for $(t_j[\bar{\sigma}, 0, m_j])_{j\neq i}$. Similarly, $t[\bar{\sigma}, 0, m]$ stands for $(t_i[\bar{\sigma}, 0, m_i])_{i\in\mathcal{I}}$. Similar abuse will be used along this proof.

To complete the proof, we show that for any equilibrium σ of $U(\mathcal{M}, \mathcal{T})$ such that $\sigma_{|\bar{T}} = \bar{\sigma}$, we have:

$$\sigma_i(t_i[\bar{\sigma}, k, m_i]) = m_i, \tag{1}$$

for each player i, integer k and message $m_i \in M_i^k(\bar{\sigma})$. The proof proceeds by induction on k.

First note that, by construction, of $t_i[\bar{\sigma}, 0, m_i]$, we must have for any equilibrium σ of $U(\mathcal{M}, \mathcal{T})$ such that $\sigma_{|\bar{T}} = \bar{\sigma}$:

$$\sigma_i(t_i[\bar{\sigma}, 0, m_i]) = m_i,$$

for each player i and message $m_i \in M_i^0(\bar{\sigma})$. Now, assume that Equation (1) is satisfied at rank $k-1$ and let us prove it is also satisfied at rank k. Fix any $m_i \in M_i^k(\bar{\sigma})$ and any equilibrium σ of $U(\mathcal{M}, \mathcal{T})$ such that $\sigma_{|\bar{T}} = \bar{\sigma}$. Note that $\sigma_i(t_i[\bar{\sigma}, k, m_i]) \in BR_i(\pi_i \mid \mathcal{M})$ where $\pi_i \in \Delta(\Theta \times \tilde{\Theta} \times M_{-i})$ is such that

$$\pi_i(\theta, \tilde{\theta}, m_{-i}) = \sum_{t_{-i}} \kappa_{t_i[\bar{\sigma}, k, m_i]}(\theta, \tilde{\theta}, t_{-i}) \sigma_{-i}(m_{-i} \mid t_{-i}).$$

In addition, by the inductive hypothesis and the fact that σ is an equilibrium of $U(\mathcal{M}, \mathcal{T})$ satisfying $\sigma_{|\bar{T}} = \bar{\sigma}$, we have $\sigma_{-i}(m_{-i} \mid t_{-i}[\bar{\sigma}, k-1, m_{-i}]) = 1$ for any $m_{-i} \in M_{-i}^{k-1}(\bar{\sigma})$. Hence, by construction of $\kappa_{t_i[\bar{\sigma}, k, m_i]}$, we have

$$\begin{aligned}
\pi_i(\theta, \tilde{\theta}, m_{-i}) &= \sum_{t_{-i}} \kappa_{t_i[\bar{\sigma}, k, m_i]}(\theta, \tilde{\theta}, t_{-i}) \sigma_{-i}(m_{-i} \mid t_{-i}). \\
&= \kappa_{t_i[\bar{\sigma}, k, m_i]}(\theta, \tilde{\theta}, t_{-i}[\bar{\sigma}, k-1, m_{-i}]) \\
&= \hat{\pi}_i^{k, m_i}(\theta, \tilde{\theta}, m_{-i})
\end{aligned}$$

We get that $\sigma_i(t_i[\bar{\sigma}, k, m_i]) \in BR_i(\pi_i \mid \mathcal{M}) = BR_i(\hat{\pi}_i^{m_i} \mid \mathcal{M}) = \{m_i\}$ as claimed. ■

We now give a first insight on the second step of the proof of our main result. First notice that, by construction, each $M(\bar{\sigma})$ satisfies the following closure property: taking any belief $\pi_i \in \Delta(\Theta \times \{\tilde{\theta}^0\} \times M_{-i}(\bar{\sigma}))$ such that $BR_i(\pi_i \mid \mathcal{M}) \neq \emptyset$, we must have $BR_i(\pi_i \mid \mathcal{M}) \subset M_i(\bar{\sigma})$ and hence, $BR_i(\pi_i \mid \mathcal{M}) = BR_i(\pi_i \mid \mathcal{M}(\bar{\sigma}))$.

Now pick a type $\bar{t}_i \in \bar{T}_i$ and a message $m_i \in R_i^1(\bar{t}_i | \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, it is possible to add a type $t_i^{m_i}$ to the model $\mathcal{T}$ defined in Proposition 1 satisfying the following two properties. First, $h_i^1(t_i^{m_i})$ is arbitrarily close to $h_i^1(\bar{t}_i)$; second, for any equilibrium σ with $\sigma_{|\bar{T}} = \bar{\sigma}$, $\sigma_i(t_i^{m_i}) = m_i$. Indeed, by definition of $R_i^1(\bar{t}_i | \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, there exists a belief $\pi_i^{m_i} \in \Delta(\Theta^* \times T_{-i} \times M_{-i}(\bar{\sigma}))$ where $\text{marg}_{\Theta^*} \pi_i^{m_i} = \text{marg}_{\Theta^*} \kappa_{\bar{t}_i}$ and such that $m_i \in BR_i(\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})}\ \pi_i^{m_i} \mid \mathcal{M}(\bar{\sigma}))$. Using our assumption on cost of messages, we can slightly perturb $\pi_i^{m_i}$ so that m_i becomes a unique best reply. So let us assume for simplicity that $\{m_i\} = BR_i(\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})}\ \pi_i^{m_i} \mid \mathcal{M}(\bar{\sigma}))$. Hence, we can define the type $t_i^{m_i}$

that assigns probability $\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})} \pi_i^{m_i}(\theta^*, m_{-i})$ to $(\theta^*, t_{-i}[\bar{\sigma}, m_{-i}])$ where $t_{-i}[\bar{\sigma}, m_{-i}]$ is defined as in Proposition 1 (i.e. $t_{-i}[\bar{\sigma}, m_{-i}]$ plays m_{-i} under any equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ such that $\sigma_{|\bar{T}} = \bar{\sigma}$). Now pick any equilibrium σ in $U(\mathcal{M}, \mathcal{T})$ such that $\sigma_{|\bar{T}} = \bar{\sigma}$. By construction, $Supp(\sigma_i(t_i^{m_i})) \subset BR_i(\text{marg}_{\Theta^* \times M_{-i}} \pi_i^{m_i} \mid \mathcal{M})$ and so $BR_i(\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})} \pi_i^{m_i} \mid \mathcal{M}) \neq \emptyset$. By the closure property described above, $BR_i(\text{marg}_{\Theta^* \times M_{-i}} \pi_i^{m_i} \mid \mathcal{M}) = BR_i(\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})} \pi_i^{m_i} \mid \mathcal{M}(\bar{\sigma}))$ and so we get that type $t_i^{m_i}$ plays m_i under the equilibrium σ and will satisfy the desired property. Using a similar reasoning, we show inductively the following "contagion" result.

Proposition 2 *There exists a model $\hat{\mathcal{T}} = (\hat{T}, \hat{\kappa})$ such that for each equilibrium $\bar{\sigma} \in \bar{\Sigma}$ and each player i the following holds. For all $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{\infty}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, there exists a sequence of types $\{\hat{t}_i[n]\}_{n=0}^{\infty}$ in $\hat{T}_i$ such that (1) $\hat{t}_i[n] \to_P \bar{t}_i$ and (2) $\sigma_i(\hat{t}_i[n]) = m_i$ for each equilibrium σ of $U(\mathcal{M}, \hat{\mathcal{T}})$ satisfying $\sigma_{|\bar{T}} = \bar{\sigma}$.*

Proof. See Appendix. ■

We are now in a position to complete the proof of our main Theorem.

Proof of Theorem 2. Pick $\hat{\mathcal{T}} = (\hat{T}, \hat{\kappa})$ as defined in Proposition 2. By definition of continuous implementation, there exists an equilibrium σ in $U(\mathcal{M}, \hat{\mathcal{T}})$ that continuously implements f and point (i) in Definition 7 ensures that $\sigma_{|\bar{T}}$ is a pure equilibrium. Now pick any $\bar{t} \in \bar{T}$ and $m \in R^{\infty}(\bar{t} \mid \mathcal{M}(\sigma_{|\bar{T}}), \bar{\mathcal{T}})$, we show that $g_{|M(\sigma_{|\bar{T}})}(m) = f(\bar{t})$ proving that the mechanism $\mathcal{M}(\sigma_{|\bar{T}})$ full implements f in rationalizable messages.

Applying Proposition 2, we know that there exists a sequence of types $\{\hat{t}[n]\}_{n=0}^{\infty}$ in $\hat{T}$ such that (1) $\hat{t}[n] \to_P \bar{t}$ and (2) $\sigma(\hat{t}[n]) = m$ for all n. By (1) and the fact that σ continuously implements f, we have $g \circ \sigma(\hat{t}[n]) \to f(\bar{t})$ while by (2) we have $g \circ \sigma(\hat{t}[n]) = g(m)$ for all n. Hence, we must have $g(m) = f(\bar{t})$ and so $g_{|M(\sigma_{|\bar{T}})}(m) = f(\bar{t})$ as claimed.
■

3.3 A Characterization

Our main Theorem provides a necessary condition for continuous implementation. Now, we show that if we restrict our attention to finite mechanisms, this condition is actually sufficient.

Theorem 3 *A social choice function f is continuously implementable by a finite mechanism if and only if it is fully rationalizable implementable by a finite mechanism.*

Proof of Theorem 3. The only if part is proved by Theorem 2. Let us prove the if part. Assume that $f : \bar{T} \to A$ is fully rationalizable-implementable by a finite mechanism $\mathcal{M} = (M, g)$ i.e. for all $\bar{t} \in \bar{T}$, $m \in R^{\infty}(\bar{t} \mid \mathcal{M}, \bar{\mathcal{T}}) \Longrightarrow g(m) = f(\bar{t})$.

Lemma 1 (Dekel, Fudenberg and Morris (2006)) *Fix any model $\mathcal{T} = (T, \kappa)$ such that $\mathcal{T} \supset \bar{\mathcal{T}}$ and any finite mechanism $\mathcal{M}$. (1) For any $\bar{t} \in \bar{T}$ and any sequence $\{t[n]\}_{n=0}^{\infty}$ in T, if $t[n] \to_P \bar{t}$ then, for n large enough, we have $R^{\infty}(t[n] \mid \mathcal{M}, \mathcal{T}) \subset R^{\infty}(\bar{t} \mid \mathcal{M}, \mathcal{T})$. (2) For any type $t \in T : R^{\infty}(t \mid \mathcal{M}, \mathcal{T})$ is non-empty.*

Now pick any model $\mathcal{T} = (T, \kappa)$ such that $\mathcal{T} \supset \bar{\mathcal{T}}$, we show that there exists an equilibrium that continuously implements f. Because M is finite and T is countable, standard arguments show that there exists a Bayes Nash equilibrium in $U(\mathcal{M}, \mathcal{T})$. Pick any sequence $\{t[n]\}_{n=0}^{\infty}$ in T, such that $t[n] \to_P \bar{t}$. It is clear that for each $n : \sigma(t[n]) \in R^{\infty}(t[n] \mid \mathcal{M}, \mathcal{T})$. In addition, for n large enough, we know by Lemma 1 that $R^{\infty}(t[n] \mid \mathcal{M}, \mathcal{T}) \subset R^{\infty}(\bar{t} \mid \mathcal{M}, \mathcal{T})$. Then, for n large enough, $\sigma(t[n]) \in R^{\infty}(\bar{t} \mid \mathcal{M}, \mathcal{T})$ and so $g \circ \sigma(t[n]) \in g\left(R^{\infty}(\bar{t} \mid \mathcal{M}, \mathcal{T})\right) = \{f(\bar{t})\}$ as claimed. ■

Theorem 3 allows to give a new rationale for the notion of virtual implementation where finite mechanisms are usually used.

In the sequel, we assume that A is a metric space and note d the associated metric. Given a social choice function f, for each $\delta > 0$, we note $B_\delta(f) = \{f' : \bar{T} \to A : d(f'(\bar{t}), f(\bar{t})) < \delta \text{ for all } \bar{t} \in \bar{T}\}$. A social choice function f is said to be partially virtually implementable by finite mechanisms if for each $\delta > 0$, there exists a social choice function $f' \in B_\delta(f)$ that is partially implementable by a finite mechanism (that may depend on δ). In the same way, we can extend the definition of continuous implementation.

Definition 9 *A social choice function f is virtually continuously implementable by finite mechanisms if for all $\delta > 0$, there exists a social choice function $f' \in B_\delta(f)$ that is continuously implementable by a finite mechanism.*

We also say that a social choice function f is virtually fully rationalizable implementable by finite mechanisms if for all $\delta > 0$, there exists a social choice function $f' \in B_\delta(f)$ that is fully rationalizable implementable by a finite mechanism. Using Theorem 3 above we can extend our characterization result to virtual implementation.

Proposition 3 *A social choice function f is virtually continuously implementable by finite mechanisms if and only if it is virtually fully rationalizable implementable by finite mechanisms.*

While the formulations in Proposition 3 and Theorem 3 are similar, their implications are quite different. Indeed, in Abreu and Matsushima (1992b) setting[29], Bayesian Incentive

[29] In this setting, the (finite) set of outcomes is extended to the set of lotteries over outcomes and the natural metric is used over this set.

Compatibility and a measurability condition are both necessary and sufficient for virtual implementation in rationalizable messages. The measurability condition seems weak and is generically satisfied[30]. Hence, any social choice function that is partially implementable (which is equivalent to Bayesian Incentive Compatibility) is virtually implementable in rationalizable messages. Since mechanisms used in these papers are finite, we know by Proposition 3 that they also ensures virtual continuous implementation. Hence, we believe that Proposition 3 provides a new foundation to the approach of virtual implementation in rationalizable messages.

4 Discussion

4.1 Failure of the Revelation Principle

We present a variant of the well-known Solomon's predicament and establish that the revelation principle does not hold for continuous implementation.

Each of two agents, 1 and 2, claims an object. There are two payoff types: at θ_1 (resp. θ_2), player 1 (resp. player 2) is the legitimate owner. The set of outcomes is $A = \{(x, p_1, p_2) \mid x \in \{0, 1, 2, 3\}$ and $p_1, p_2 \in \mathbb{R}_+\}$ where p_i is the level of the fine imposed on player i and the variable x correspond to the following situations. If $x = 0$, the object is not given to either player; if $x \in \{1, 2\}$, it is attributed to player x, and if $x = 3$, both players are "punished" and the none of them receive the object. The social planner wishes to give the good to the true owner, i.e. he wants to implement continuously the social choice function $f : \{\theta_1, \theta_2\} \to A$ for which $f(\theta_1) = (1, 0, 0)$ and $f(\theta_2) = (2, 0, 0)$. Utility functions are assumed to be quasi-linear and the object to have a monetary value for each player. More precisely, this value for player i is v_H if he is the legitimate owner of the object and v_L if he is not, with $v_H > v_L > 0$. Finally, the punishment outcome ($x = 3$) corresponds to a fine f_L for player i if he is the legitimate owner and to a fine f_H if he is not, with $f_H > f_L > 0$. For instance when the payoff type is θ_1, the utility of player 1 when the outcome is $(3, p_1, p_2)$ is: $u_1((3, p_1, p_2), \theta_1) = -f_L - p_1$ and when outcome $(1, p_1, p_2)$ is given: $u_1((1, p_1, p_2), \theta_1) = v_H - p_1$.

The following two claims establish the failure of the revelation principle when a continuity requirement is taken into account.

Claim 1 *f is not continuously implementable with a direct mechanism i.e. a mechanism $\mathcal{M} = (M, g)$ in which for each $i \in \{1, 2\}$, $M_i = \{\theta_1, \theta_2\}$.*

[30] As noted in the introduction, type diversity (which states that every type has distinct preferences over lotteries unconditional on others' types) is sufficient for all social choice functions to be measurable in the sense of Abreu and Matsushima (1992b).

Proof. We establish that no mechanism $\mathcal{M}' \subset \mathcal{M}$ can fully implement in NE the social choice function f. Theorem 2 completes the proof of Claim 1. Obviously, f cannot be implemented if the set of message profiles is a singleton. Now assume that the set of message profiles is a singleton for one player, say player 1, i.e.

$M_1' = \{\theta\}$ for some $\theta \in \{\theta_1, \theta_2\}$ and $M_2' = \{\theta_1, \theta_2\}$. In this case, player 2 must have a message m_2 such that $g(m_2, \theta) = (2, 0, 0)$. Then, m_2 strictly dominates any message that yields outcome $(1, 0, 0)$. Hence, $(1, 0, 0)$ cannot be an equilibrium outcome at state θ_1. Finally, we show that the direct mechanism $\mathcal{M}$ cannot fully implement in NE the social choice function f. Proceed by contradiction and assume that there exist two message profiles $m^*_{\theta_1} = (m^*_{1,\theta_1}, m^*_{2,\theta_1}) \in NE(\mathcal{M}, \theta_1)$ and $m^*_{\theta_2} = (m^*_{1,\theta_2}, m^*_{2,\theta_2}) \in NE(\mathcal{M}, \theta_2)$ such that $g(m^*_{\theta_1}) = f(\theta_1)$ and $g(m^*_{\theta_2}) = f(\theta_2)$. It is easily checked that for each player i : $m^*_{i,\theta_1} \neq m^*_{i,\theta_2}$, otherwise, at some state, one player would have an incentive to deviate from the equilibrium. Now for message profile $(m^*_{1,\theta_1}, m^*_{2,\theta_2})$, there is (at least) one player who does not receive the object. Assume without loss of generality that this is player 1 (a similar reasoning holds for player 2). Let us show that necessarily, $m^*_{\theta_2} = (m^*_{1,\theta_2}, m^*_{2,\theta_2})$ is a pure Nash equilibrium at θ_1. Since $g(m^*_{1,\theta_2}, m^*_{2,\theta_2}) = (2, 0, 0)$ is the best outcome for player 2, he has no incentive to deviate. We also know by construction that if player 1 deviates, the outcome is $g(m^*_{1,\theta_1}, m^*_{2,\theta_2}) = (x, p_1, p_2)$ where player 1 does not get the object. Hence, $u_1(g(m^*_{\theta_2}), \theta_1) = u_1((2, 0, 0), \theta_1) = 0 \geq u_1((x, p_1, p_2), \theta_1) = u_1(g(m^*_{1,\theta_2}, m^*_{2,\theta_2}), \theta_1)$ and so player 1 does not have any incentive to deviate either. Thus f is not fully implementable by $\mathcal{M}$ which completes the proof. ■

However, as we will show in the following lines, when we expand the set of messages using indirect mechanisms, f can be continuously implemented.

Claim 2 *There exists an indirect mechanism that continuously implements f.*

Proof. Consider the following indirect mechanism. Each player has three possible messages (Mine, His, and Mine+) and the outcome function is given by the matrix below, where $v_L < P < v_H$, $f_L < p < f_H$, and $\xi > p$.[31]

	Mine	His	Mine+
Mine	$(0, \xi, \xi)$	$(1, 0, 0)$	$(2, \xi, P)$
His	$(2, 0, 0)$	$(0, \xi, \xi)$	$(0, p, 0)$
Mine+	$(1, P, \xi)$	$(0, 0, p)$	$(3, 0, 0)$

At θ_1, action "His" is strictly dominated by "Mine+"for player 1. Consequently, in the second round of elimination, "Mine" and "Mine+" are strictly dominated by "His"

[31] Using the usual convention, player 1 is the row player while player 2 is the column player.

for player 2 at θ_1. Finally, in the third round, "Mine" is strictly better than "Mine+" for player 1. Hence, $(Mine, His)$ is the unique rationalizable action profile at θ_1. A symmetric reasoning applies at θ_2. By Theorem 3, we conclude that this finite indirect mechanism implements continuously the social choice function f. ■

4.2 Alternative Topology: Uniform Convergence

In this paper, we define the notion of continuous implementation using the topology of point-wise convergence. This topology is standard when working in the universal type space and has a simple interpretation. However, other topologies are interesting and so other notions of continuous implementation are worth to be investigated. One natural candidate is the topology of uniform convergence.[32] While interesting in its own right, we show that all social choice functions that are partially implementable in strict Nash equilibria with a finite mechanism are continuously implementable under this topology. This condition is much weaker than the one obtained under the topology of point-wise convergence. In particular, recall that in complete information settings, under mild conditions[33] and with more than three players, any social function is partially implementable in strict Nash equilibria with a direct mechanism (and so with a finite mechanism in our setting).

To introduce the topology of uniform convergence, we first recall the definition of the Prohorov distance that metrizes the topology of weak convergence of measures. Given a metric space (X, ρ) the Prohorov distance between any two $\mu, \mu' \in \Delta(X)$ is

$$\inf\{\delta > 0 : \mu'(A) \leq \mu(A_\delta) + \delta \text{ for every Borel set } A \subset X\}$$

where $A^\delta = \{x \in X : \inf_{y \in A} \rho(x, y) < \delta\}$.

Write $X_0 = \Theta^*$ and for each $k \geq 1 : X_k = [\Delta(X_{k-1})]^I \times X_{k-1}$. Now, let d^0 be the discrete metric on Θ^* and d^1 the Prohorov distance on 1st level beliefs $\Delta(\Theta^*)$. Then, recursively, for any $k \geq 2$, let d^k be the Prohorov distance on the kth level beliefs $\Delta(X_{k-1})$ when X_{k-1} is given the product metric induced by $d^0, d^1, ..., d^{k-1}$. We say that a sequence of types $\{t_i[n]\}_{n=0}^\infty$ converges uniformly to a type t_i, if $d_U(t_i[n], t_i) \equiv \sup_{k \geq 1} d^k(h_i^k(t_i[n]), h_i^k(t_i)) \to 0$; in this case we write $t_i[n] \to_U t_i$. We also write $t[n] \to_U t$, if, $t_i[n] \to_U t_i$, for each $i \in \mathcal{I}$. In this topology, two types are close if they have very similar first-order beliefs, second order-beliefs and so on up to infinity where the degree of

[32] Another topology in the universal type space is the strategic topology as defined in Dekel, Fudenberg and Morris (2006). Di Tillio and Faingold (2007) established the equivalence between uniform topology and strategic topology around finite types. Since $\bar{T}$ is finite, the result of this section is also true under the strategic topology.

[33] For instance in quasi-linear settings with arbitrary small transfers.

similarity is uniform over the levels of the belief hierarchy.[34]

Definition 10 *A social choice function $f : \bar{T} \to A$ is continuously implementable w.r.t. $\to_U$ if there exists a mechanism $\mathcal{M}$ such that for any model $\mathcal{T} \supset \bar{\mathcal{T}}$, there is a Bayes Nash equilibrium σ in the induced game $U(\mathcal{M}, \mathcal{T})$ where for each $\bar{t} \in \bar{T}$, (i) $\sigma(\bar{t})$ is pure and (ii) for any sequence $t[n] \to_U \bar{t}$ where for each $n : t[n] \in T$, we have $g \circ \sigma(t[n]) \to f(\bar{t})$.*

Recall that a profile of strategies $\sigma = (\sigma_1, ..., \sigma_I)$ is a strict Bayes Nash equilibrium in $U(\mathcal{M}, \mathcal{T})$ if for each $i \in \mathcal{I}$, and for each $t_i \in T_i$,

$$\{\sigma_i(t_i)\} = BR_i(\pi_i(\cdot \mid t_i, \sigma_{-i}) \mid \mathcal{M}).$$

We say that a social choice function $f : \bar{T} \to A$ is partially SNE-implementable if there exists a mechanism $\mathcal{M}$ and a strict Bayes Nash equilibrium σ in the induced game $U(\mathcal{M}, \bar{\mathcal{T}})$ where for each $\bar{t} \in \bar{T}$, $g \circ \sigma(\bar{t}) = f(\bar{t})$. We can now state a simple sufficient condition for continuous implementation w.r.t. uniform convergence.

Proposition 4 *If f is partially SNE-implementable by a finite mechanism then it is continuously implementable w.r.t. $\to_U$.*

Proof. In the sequel, we use the following notations. For any $i, \bar{t}_i \in \bar{T}_i$, and each k, we note $[\bar{t}_i]^k = \{\bar{t}'_i \in \bar{T}_i : h_i^k(\bar{t}'_i) = h_i^k(\bar{t}_i)\}$. We also note $[\bar{t}_i]$ for $\{\bar{t}'_i \in \bar{T}_i : h_i(\bar{t}'_i) = h_i(\bar{t}_i)\}$. In addition, given a model $\mathcal{T} = (T, \kappa) \supset \bar{\mathcal{T}}$, and any type $\bar{t}_i \in \bar{T}_i$, we write $C_\delta^k(\bar{t}_i)$ for the set $\{t_i \in T_i : h_i^k(t_i) \in B_\delta^k(h_i^k(\bar{t}_i))\}$ where B_δ^k is an open ball w.r.t. the distance d^k defined over $\Delta(X_{k-1})$. In a similar way, $C_\delta(\bar{t}_i)$ denotes the set $\{t_i \in T_i : h_i(t_i) \in B_\delta(h_i(\bar{t}_i))\}$ where B_δ is an open ball w.r.t. the metric $d_U \equiv \sup_{k \geq 1} d^k$ over $\mathcal{T}_i^*$. We will also use the notation $\bar{C}_\delta^k(\bar{t}_i)$ for the set $\{\bar{t}'_i \in \bar{T}_i : h_i^k(\bar{t}'_i) \in B_\delta^k(h_i^k(\bar{t}_i))\}$ and $\bar{C}_\delta(\bar{t}_i)$ for the set $\{\bar{t}'_i \in \bar{T}_i : h_i(\bar{t}'_i) \in B_\delta(h_i(\bar{t}_i))\}$.

Lemma 2 *Pick any model $\mathcal{T} = (T, \kappa) \supset \bar{\mathcal{T}}$. There exists $\bar{\delta} > 0$ such that for all $\delta < \bar{\delta}$, all $\bar{t}_i \in \bar{T}_i$ and $t_i \in C_\delta(\bar{t}_i)$:*

$$\mid \kappa_{t_i}(\{(\theta, \tilde{\theta}^0)\} \times C_\delta(\bar{t}_{-i})) - \bar{\kappa}_{\bar{t}_i}(\{\theta\} \times [\bar{t}_{-i}]) \mid \leq \delta$$

for all $\bar{t}_{-i} \in \bar{T}_{-i}$ and $\theta \in \Theta$.

Proof. For each $i, t_i \in T_i$, there exists a unique probability measure $\mu_{t_i} \in \Delta(\Theta^* \times \mathcal{T}_{-i}^*)$ whose marginal $\mu_{t_i}^k$ on X_{k-1} coincides with $h_i^k(t_i)$ for each $k \geq 1$. By definition of the Prohorov metric, for any $\delta > 0$, $\bar{t}_i \in \bar{T}_i$ and $t_i \in C_\delta(\bar{t}_i)$:

$$\mu_{\bar{t}_i}^k(\{(\theta, \tilde{\theta}^0)\} \times h_{-i}^{k-1}(\bar{t}_{-i})) - \mu_{t_i}^k(\{(\theta, \tilde{\theta}^0)\} \times B_\delta^{k-1}(h_{-i}^{k-1}(\bar{t}_{-i}))) \leq \delta, \qquad (2)$$

[34] Note that the topology of uniform convergence is an extension of notion of common p-belief (Monderer and Samet (1989)) to incomplete information environments.

and

$$\mu^k_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times B^{k-1}_\delta(h^{k-1}_{-i}(\bar{t}_{-i}))) - \mu^k_{\bar{t}_i}(\{(\theta,\tilde{\theta}^0)\}\times B^{k-1}_{2\delta}(h^{k-1}_{-i}(\bar{t}_{-i}))) \le \delta, \qquad (3)$$

for each k, $\bar{t}_{-i}$ and θ. (Where in the definition of Prohorov metric, we respectively use as Borel sets $A = \{(\theta,\tilde{\theta}^0)\}\times h^{k-1}_{-i}(\bar{t}_{-i})$ and $A = \{(\theta,\tilde{\theta}^0)\}\times B^{k-1}_\delta(h^{k-1}_{-i}(\bar{t}_{-i}))$.) Note that for all $k : \mu^k_{\bar{t}_i}(\{(\theta,\tilde{\theta}^0)\}\times h^{k-1}_{-i}(\bar{t}_{-i})) = \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]^{k-1})$ and $\mu^k_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times B^{k-1}_\delta(h^{k-1}_{-i}(\bar{t}_{-i}))) = \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i}))$. Since $\bar{T}$ is finite, there is K large enough so that for all $k \ge K$, $\bar{t}_i \in \bar{T}_i$ and $\bar{t}_{-i}\in\bar{T}_{-i} : \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]^{k-1}) = \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}])$. Hence, for all $k\ge K$, by (2) :

$$\bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]) - \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i})) \le \delta,$$

for each $\bar{t}_{-i}$ and θ. Note that $C_\delta(\bar{t}_{-i}) = \cap_k C^k_\delta(\bar{t}_{-i})$ and $\{C^k_\delta(\bar{t}_{-i})\}_k$ is a decreasing sequence. Hence, by continuity from above of probability measures, we get

$$\lim_{k\to\infty}\kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i})) = \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times \cap_k C^{k-1}_\delta(\bar{t}_{-i})) = \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C_\delta(\bar{t}_{-i})).$$

Thus, we have

$$\bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]) - \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C_\delta(\bar{t}_{-i})) \le \delta,$$

for all $\bar{t}_{-i}\in\bar{T}_{-i}$ and θ.

Now, we have to show that for $\delta > 0$ small enough, for each $\bar{t}_i \in \bar{T}_i$ and $t_i \in C_\delta(\bar{t}_i)$:

$$-\delta \le \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]) - \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C_\delta(\bar{t}_{-i})),$$

for all $\bar{t}_{-i}\in\bar{T}_{-i}$ and θ and the proof will be complete. By (3), we have for each $\bar{t}_i\in\bar{T}_i$ and $t_i \in T_i$ such that $t_i \in C_\delta(\bar{t}_i)$:

$$\kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i})) - \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times \bar{C}^{k-1}_{2\delta}(\bar{t}_{-i})) \le \delta$$

for each k, $\bar{t}_{-i}\in\bar{T}_{-i}$ and θ. Since $\bar{T}$ is finite, we have that for δ small enough and k large enough: $\bar{C}^{k-1}_{2\delta}(\bar{t}_{-i}) = [\bar{t}_{-i}]$ for all $\bar{t}_{-i}\in\bar{T}_{-i}$ which yields

$$\kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i})) - \bar{\kappa}_{\bar{t}_i}(\{\theta\}\times[\bar{t}_{-i}]) \le \delta$$

for all $\bar{t}_{-i}\in\bar{T}_{-i}$ and θ. Since $C_\delta(\bar{t}_{-i}) \subset C^{k-1}_\delta(\bar{t}_{-i})$ we have $\kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C_\delta(\bar{t}_{-i})) \le \kappa_{t_i}(\{(\theta,\tilde{\theta}^0)\}\times C^{k-1}_\delta(\bar{t}_{-i}))$, this yields the desired result. ■

Pick the finite mechanism $\mathcal{M}$ under which there exists a strict Bayes Nash equilibrium $\bar{\sigma}$ in the induced game $U(\mathcal{M},\bar{T})$ where for each $\bar{t}\in\bar{T}$, $g\circ\bar{\sigma}(\bar{t}) = f(\bar{t})$. Let us show that this mechanism continuously implements f w.r.t. $\to_U$. Since $\bar{\sigma}$ is a strict Nash equilibrium and the sets $\bar{T}$, M and so by construction Θ^* are finite, we know that there exists $\bar{\varepsilon} > 0$ such that for each $i\in\mathcal{I}$ and each $\bar{t}_i\in\bar{T}_i$,

$$\{\bar{\sigma}_i(\bar{t}_i)\} = BR_i(\pi'_i \mid \mathcal{M}). \qquad (4)$$

whenever $\|\pi'_i(\cdot) - \pi_i(\cdot \mid \bar{t}_i, \bar{\sigma}_{-i})\| \leq \bar{\varepsilon}$.[35]

Consider the induced game $U(\mathcal{M}, \mathcal{T})$ and build a modified game $U^{\delta}(\mathcal{M}, \mathcal{T})$ where for each player i, the set of strategies is restricted to

$$\Sigma_i^{\delta} = \{\sigma_i : T_i \to \Delta(M_i) \mid \text{for all } \bar{t}_i \in \bar{T}_i,\ \sigma_i(t_i) = \bar{\sigma}_i(\bar{t}_i) \text{ for all } t_i \in C_{\delta}(\bar{t}_i))\}$$

where $\delta > 0$ is assumed to be small.

Since M is finite and T is countable, standard arguments show the existence of a Bayes Nash equilibrium σ in $U^{\delta}(\mathcal{M}, \mathcal{T})$. Note that for each $\bar{t} \in \bar{T}$, (i) $\sigma(\bar{t}) = \bar{\sigma}(\bar{t})$ since $\bar{t} \in C_{\delta}(\bar{t})$ and so $\sigma(\bar{t})$ is pure and (ii) for any sequence $t[n] \to_U \bar{t}$ where for each $n : t[n] \in T$, we have $g \circ \sigma(t[n]) \to g \circ \sigma(\bar{t}) = f(\bar{t})$ since $t[n] \in C_{\delta}(\bar{t})$ for n large. Hence, it remains to show that σ is an equilibrium of the original game $U(\mathcal{M}, \mathcal{T})$. It is clear that whenever $t_i \notin \bigcup_{\bar{t}_i \in \bar{T}_i} C_{\delta}(\bar{t}_i) : m_i \in Supp(\sigma_i(t_i)) \Rightarrow m_i \in BR_i(\pi_i(\cdot \mid t_i, \sigma_{-i}) \mid \mathcal{M})$ since the set of available actions is not modified for these types from $U^{\delta}(\mathcal{M}, \mathcal{T})$ to $U(\mathcal{M}, \mathcal{T})$. Now, pick $t_i \in T_i$ and $\bar{t}_i \in \bar{T}_i$ such that $t_i \in C_{\delta}(\bar{t}_i)$. Assuming δ is small enough, by Lemma 2 and the construction of Σ_{-i}^{δ}, we obtain that $\|\pi_i(\cdot \mid t_i, \sigma_{-i}) - \pi_i(\cdot \mid \bar{t}_i, \bar{\sigma}_{-i})\| \leq \bar{\varepsilon}$ and so by (4), playing $\bar{\sigma}_i(\bar{t}_i)$ is the unique best reply. ■

4.3 Ex ante approach vs. interim approach

In this article, we formalized the notion of proximity using the interim approach due to Weinstein and Yildiz (2007) and the notion of type. In this approach, the modeler has in mind a set of hierarchies of beliefs and is interested in the strategic behavior of any type close to some type of the original model. It is possible to build another test of continuity using the ex ante approach due to Kajii and Morris (1997) and the notion of model. We briefly expose in the following lines the ex ante approach for the simplified case in which the initial model is a complete information one. While in our article types are defined using conditional beliefs, some specification of the prior distribution for each player is needed to build a perturbation under the ex ante approach. More precisely, a perturbation will be considered as close to the initial model if the set of types that are close to complete information types (as defined in our paper) has an ex ante probability that is close to one. A social choice function is *ex ante* continuously implementable if in any perturbation (arbitrarily) close to complete information, there is a Bayes Nash equilibrium such that the social choice function is implemented with an (arbitrarily) high ex ante probability. While this notion can be seen as less permissive than the one defined in our paper, Oyama and Tercieux (2005) have shown that the approach of Weinstein and Yildiz (2007) and the one of Kajii and Morris (1997) yield essentially the same qualitative results provided that

[35] Here we use the norm max, i.e:. $\|\pi_i(\cdot) - \pi'_i(\cdot)\| = \max_{(\theta^*, m_{-i})} |\pi_i(\theta^*, m_{-i}) - \pi'_i(\theta^*, m_{-i})|$.

we allow players to have heterogenous prior beliefs. Hence, we believe that our results would be maintained when considering ex ante continuous implementation if a common prior is not assumed to hold. The characterization of ex ante continuous implementation under the common prior assumption is an open question which is left for further research.

Appendix

Proof of Proposition 2

We define the set C by:

$$C := \cup_{q \in \mathbb{N}_*} \{\frac{1}{q}\} \cup \{0\}.$$

Now we build the model $\hat{\mathcal{T}} = (\hat{T}, \hat{\kappa})$ as follows. For each $\varepsilon \in C$, k, $\bar{\sigma} \in \bar{\Sigma}$, $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, we build inductively $\hat{t}_i[\varepsilon, k, \bar{\sigma}, \bar{t}_i, m_i]$ and set

$$\hat{T}_i = \bigcup_{\bar{t}_i \in \bar{T}_i} \bigcup_{k=1}^{\infty} \bigcup_{\varepsilon \in C} \bigcup_{\bar{\sigma} \in \bar{\Sigma}} \bigcup_{m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})} \hat{t}_i[\varepsilon, k, \bar{\sigma}, \bar{t}_i, m_i] \cup T_i,$$

where T_i is defined as in Proposition 1. Note that $\hat{T}_i$ is countable.

In the sequel, we fix an arbitrary $\bar{\sigma} \in \bar{\Sigma}$. This equilibrium $\bar{\sigma}$ is sometimes omitted in our notations.

We know that for each k, player i of type $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, there exists $\pi_{\bar{t}_i}^{k,m_i} \in \Delta(\Theta \times \bar{T}_{-i} \times M_{-i}(\bar{\sigma}))$ such that

$$\text{marg}_{\Theta \times \bar{T}_{-i}} \ \pi_{\bar{t}_i}^{k,m_i} = \bar{\kappa}_{\bar{t}_i};$$

$$\text{marg}_{\bar{T}_{-i} \times M_{-i}(\bar{\sigma})} \ \pi_{\bar{t}_i}^{k,m_i}(\bar{t}_{-i}, m_{-i}) > 0 \Rightarrow m_{-i} \in R_{-i}^{k-1}(\bar{t}_{-i} \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$$

and

$$m_i \in BR_i(\text{marg}_{\Theta^* \times M_{-i}(\bar{\sigma})} \ \pi_{\bar{t}_i}^{k,m_i} \mid \mathcal{M}(\bar{\sigma})).$$

For ease of exposition, we sometimes consider $\pi_{\bar{t}_i}^{k,m_i}$ as a measure over $\Theta \times \bar{T}_{-i} \times M_{-i}(\bar{\sigma})$ and sometimes as a measure over $\Theta^* \times \bar{T}_{-i} \times M_{-i}(\bar{\sigma})$ assigning probability one on $\{\tilde{\theta}^0\}$. Similar abuses will be used throughout the proof.

First, we let $\hat{t}_i[\varepsilon, 1, \bar{\sigma}, \bar{t}_i, m_i]$ be such that $\hat{\kappa}_{\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]}$ satisfies the following two conditions:

$$\text{marg}_{\tilde{\Theta}} \ \hat{\kappa}_{\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]} = \varepsilon \delta_{\tilde{\theta}^{m_i}} + (1-\varepsilon)\, \delta_{\tilde{\theta}^0} \tag{5}$$

where δ_x denotes the probability distribution that puts probability 1 on $\{x\}$. And,

$$\text{marg}_{\Theta \times \hat{T}_{-i}} \ \hat{\kappa}_{\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]} = \pi_{\bar{t}_i}^{1,m_i} \circ \left(\tau_{-i}^{\varepsilon,1}\right)^{-1} \tag{6}$$

where $\left(\tau_{-i}^{\varepsilon,1}\right)^{-1}$ stands for the preimage of the function $\tau_{-i}^{\varepsilon,1} : (\theta, \bar{t}_{-i}, m_{-i}) \longmapsto (\theta, t_{-i}[\bar{\sigma}, m_{-i}])$ and $t_{-i}[\bar{\sigma}, m_{-i}] \in T_{-i}$ is the type profile defined in Proposition 1. Recall that $\sigma_{-i}(t_{-i}[\bar{\sigma}, m_{-i}]) =$

m_{-i} for any equilibrium σ in $U(\mathcal{M},\mathcal{T})$ s.t. $\sigma_{|\bar{T}} = \bar{\sigma}$. Now for each $k \geq 2$, define $\hat{t}_i[\varepsilon, k, \bar{\sigma}, \bar{t}_i, m_i]$ inductively by

$$\text{marg}_{\tilde{\Theta}}\ \hat{\kappa}_{\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]} = \varepsilon\delta_{\tilde{\theta}^{m_i}} + (1-\varepsilon)\,\delta_{\tilde{\theta}^0},$$

and,

$$\text{marg}_{\Theta\times\hat{T}_{-i}}\ \hat{\kappa}_{\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]} = \pi_{\bar{t}_i}^{k,m_i} \circ \left(\tau_{-i}^{\varepsilon,k}\right)^{-1}$$

where $\tau_{-i}^{\varepsilon,k} : (\theta, \bar{t}_{-i}, m_{-i}) \longmapsto \big(\theta, \hat{t}_{-i}[\varepsilon, k-1, \bar{\sigma}, \bar{t}_{-i}, m_{-i}]\big)$.

Claim 3 *For each $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^\infty(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}}) : \hat{t}_i[\frac{1}{k}, k, \bar{\sigma}, \bar{t}_i, m_i] \to_P \bar{t}_i$ as $k \to \infty$.*

To prove this claim we will use the following well-known lemma.

Lemma 3 (Mertens and Zamir (1985) and Brandenburger and Dekel (1993)) *Let $\mathcal{T} = (T, \kappa)$ be any model such that $\Theta^* \times T$ is complete and separable and κ_{t_i} is a continuous function of t_i. Then, the mapping $h : T \to \mathcal{T}^*$ is continuous.*

Proof of Claim 3. In the sequel, we will note $\bar{h}$ the (continuous) mapping that projects $\bar{T}$ into $\mathcal{T}^*$ and, in a similar way, $\hat{h}$ the (continuous) mapping from $\hat{T}$ to $\mathcal{T}^*$.

For any $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, since[36] for all $k \geq 1 : \hat{t}_i[\varepsilon, k, \bar{\sigma}, \bar{t}_i, m_i] \to \hat{t}_i[0, k, \bar{\sigma}, \bar{t}_i, m_i]$ as $\varepsilon \to 0$, by Lemma 3, for all $k \geq 1 : \hat{h}_i^k\left(\hat{t}_i[\varepsilon, k, \bar{\sigma}, \bar{t}_i, m_i]\right) \to \hat{h}_i^k\left(\hat{t}_i[0, k, \bar{\sigma}, \bar{t}_i, m_i]\right)$ as $\varepsilon \to 0$.

Let us now show that for all $k \geq 1$ and $k' \geq k : \hat{h}_i^k\left(\hat{t}_i[0, k', \bar{\sigma}, \bar{t}_i, m_i]\right) = \bar{h}_i^k(\bar{t}_i)$ for all $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{k'}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$. First notice that the first order beliefs are equal, i.e. for all $k' \geq 1$, $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{k'}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$:

$$\begin{aligned}
\hat{h}_i^1\left(\hat{t}_i[0, k', \bar{\sigma}, \bar{t}_i, m_i]\right) &= \text{marg}_\Theta\ \hat{\kappa}_{\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i]} \\
&= \text{marg}_\Theta\ \pi_{\bar{t}_i}^{k',m_i} \circ \left(\tau_{-i}^{0,k'}\right)^{-1} \\
&= \text{marg}_\Theta\ \pi_{\bar{t}_i}^{k',m_i} = \text{marg}_\Theta\ \bar{\kappa}_{\bar{t}_i} = \bar{h}_i^1(\bar{t}_i)
\end{aligned}$$

where the third and the fourth equalities are by definition of $\tau_{-i}^{0,k'}$ and $\pi_{\bar{t}_i}^{k',m_i}$ respectively. Now fix some $k \geq 2$ and let L be the set of all belief profiles of players other than i at order $k-1$. Toward an induction, assume that for all $k' \geq k-1 : \hat{h}_j^{k-1}(\hat{t}_j[0, k', \bar{\sigma}, \bar{t}_j, m_j]) =$

[36] A type in $\hat{T}_i$ is either in T_i – which is endowed with the discrete topology, say τ_{T_i} – or it is in $\hat{T}_i \backslash T_i$. Any point in $\hat{T}_i \backslash T_i$ is identified with an element of the set $C \times \mathbb{N} \times \bar{\Sigma} \times M_i$ where $\mathbb{N}, \bar{\Sigma}, M_i$ are all endowed with the discrete topology while C is endowed with the usual topology on $\mathbb{R}$ induced on C. Finally, $C \times \mathbb{N} \times \bar{\Sigma} \times M_i$ is endowed with the product topology; call this topology $\tau_{\hat{T}_i \backslash T_i}$. The topology over $\hat{T}_i$ is the coarsest topology that contains $\tau_{T_i} \cup \tau_{\hat{T}_i \backslash T_i}$. It can easily be checked that under such a topology, $\hat{\mathcal{T}}$ satisfies the conditions of Lemma 3.

$\bar{h}_j^{k-1}(\bar{t}_j)$ for each j, $\bar{t}_j \in \bar{T}_j$ and $m_j \in R_j^{k'}(\bar{t}_j \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$. Then for all $k' \geq k$: $\text{proj}_{\Theta\times L} \circ \left(id_\Theta \times \hat{h}_{-i}\right) \circ \tau_{-i}^{0,k'} = \overline{\text{proj}}_{\Theta\times L} \circ \left(id_\Theta \times \bar{h}_{-i} \times id_{M_{-i}(\bar{\sigma})}\right)$ where id_Θ (resp. $id_{M_{-i}(\bar{\sigma})}$) is the identity mapping from Θ to Θ (resp. from $M_{-i}(\bar{\sigma})$ to $M_{-i}(\bar{\sigma})$) while $\text{proj}_{\Theta\times L}$ (resp. $\overline{\text{proj}}_{\Theta\times L}$) is the projection mapping from $\Theta \times \mathcal{T}^*$ to $\Theta \times L$ (resp. from $\Theta \times \mathcal{T}^* \times M_{-i}(\bar{\sigma})$ to $\Theta \times L$); hence for all $k' \geq k$, $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{k'}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$:

$$\begin{aligned}
\text{marg}_{\Theta\times L}\ \hat{\kappa}_{\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i]} \circ \left(id_\Theta \times \hat{h}_{-i}\right)^{-1} &= \text{marg}_{\Theta\times L}\ \pi_{\bar{t}_i}^{k',m_i} \circ \left(\tau_{-i}^{0,k'}\right)^{-1} \circ \left(id_\Theta \times \hat{h}_{-i}\right)^{-1} \\
&= \pi_{\bar{t}_i}^{k',m_i} \circ \left(\tau_{-i}^{0,k'}\right)^{-1} \circ \left(id_\Theta \times \hat{h}_{-i}\right)^{-1} \circ \left(\text{proj}_{\Theta\times L}\right)^{-1} \\
&= \pi_{\bar{t}_i}^{k',m_i} \circ \left(id_\Theta \times \bar{h}_{-i} \times id_{M_{-i}(\bar{\sigma})}\right)^{-1} \circ \left(\overline{\text{proj}}_{\Theta\times L}\right)^{-1} \\
&= \text{marg}_{\Theta\times L}\ \pi_{\bar{t}_i}^{k',m_i} \circ \left(id_\Theta \times \bar{h}_{-i} \times id_{M_{-i}(\bar{\sigma})}\right)^{-1} \\
&= \text{marg}_{\Theta\times L}\ \bar{\kappa}_{\bar{t}_i} \circ \left(id_\Theta \times \bar{h}_{-i}\right)^{-1}
\end{aligned}$$

Therefore,

$$\begin{aligned}
\hat{h}_i^k(\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i]) &= \delta_{\hat{h}_i^{k-1}(\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i])} \times \text{marg}_{\Theta\times L}\ \hat{\kappa}_{\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i]} \circ \left(id_\Theta \times \hat{h}_{-i}\right)^{-1} \\
&= \delta_{\bar{h}_i^{k-1}(\bar{t}_i)} \times \text{marg}_{\Theta\times L}\ \bar{\kappa}_{\bar{t}_i} \circ \left(id_\Theta \times \bar{h}_{-i}\right)^{-1} = \bar{h}_i^k(\bar{t}_i)
\end{aligned}$$

showing that $\hat{h}_i^k(\hat{t}_i[0,k',\bar{\sigma},\bar{t}_i,m_i]) = \bar{h}_i^k(\bar{t}_i)$. Thus, we have proved that for all $k \geq 1$, all $k' \geq k$: $\hat{h}_i^k\left(\hat{t}_i[\varepsilon,k',\bar{\sigma},\bar{t}_i,m_i]\right) \to \bar{h}_i^k(\bar{t}_i)$ as $\varepsilon \to 0$ for any $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{k'}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$. This implies that for all $k \geq 1$: $\hat{h}_i^k\left(\hat{t}_i[\frac{1}{k'},k',\bar{\sigma},\bar{t}_i,m_i]\right) \to \bar{h}_i^k(\bar{t}_i)$ as $k' \to \infty$ for any $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{\infty}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$ as claimed. ■

Claim 4 *For each k, $\bar{t}_i \in \bar{T}_i$ and $m_i \in R_i^{\infty}(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$, we have: $\sigma_i(\hat{t}_i[\frac{1}{k},k,\bar{\sigma},\bar{t}_i,m_i]) = m_i$ for any equilibrium σ of $U(\mathcal{M},\hat{\mathcal{T}})$ satisfying $\sigma_{|\bar{\mathcal{T}}} = \bar{\sigma}$.*

Proof. Fix a type $\bar{t}_i \in \bar{T}_i$ and an equilibrium σ of $U(\mathcal{M},\hat{\mathcal{T}})$ satisfying $\sigma_{|\bar{\mathcal{T}}} = \bar{\sigma}$. We will show by induction on k that for all $\varepsilon > 0$ and $k \geq 1$: $\sigma_i(\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]) = m_i$ for all message $m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$.

Recall that, by construction, for all $m_i \in M_i(\bar{\sigma}) : t_i[\bar{\sigma},m_i] \in T_i$ is the type in Proposition 1 such that $\sigma_i(t_i[\bar{\sigma},m_i]) = m_i$. First, fix $\varepsilon > 0$ and $m_i \in R_i^1(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$ and let us prove that $\sigma_i(\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]) = m_i$. For each $\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]$, define the belief

$$\pi_i^{\varepsilon,1} = \hat{\kappa}_{\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]} \circ \gamma^{-1} \in \Delta(\Theta^* \times \hat{T}_{-i} \times M_{-i})$$

where $\gamma : (\theta^*, t_{-i}[\bar{\sigma},m_{-i}]) \mapsto (\theta^*, t_{-i}[\bar{\sigma},m_{-i}], m_{-i})$. Note that by construction, $\pi_i^{\varepsilon,1}$ is the belief of type $\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]$ on $\Theta^* \times \hat{T}_{-i} \times M_{-i}$ when he believes that m_{-i} is played at each $(\theta^*, t_{-i}[\bar{\sigma},m_{-i}])$. Hence, for each $\varepsilon \geq 0$, $\pi_i^{\varepsilon,1}$ corresponds to beliefs of type $\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]$

when the equilibrium σ is played. Now, by Equations (5) and (6), the belief $\pi_i^{0,1}$ of type $\hat{t}_i[0,1,\bar{\sigma},\bar{t}_i,m_i]$ satisfies

$$\text{marg}_{\Theta^*\times M_{-i}}\ \pi_i^{0,1} = \text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i}\circ\left(\tau_{-i}^{0,1}\right)^{-1}\circ(\gamma_\Theta)^{-1} = \text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i}$$

where $\gamma_\Theta : (\theta, t_{-i}[\bar{\sigma}, m_{-i}]) \mapsto (\theta, \tilde{\theta}^0, t_{-i}[\bar{\sigma}, m_{-i}], m_{-i})$. Since $Supp(\sigma_i(\hat{t}_i[0,1,\bar{\sigma},\bar{t}_i,m_i])) \subset BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_i^{0,1} \mid \mathcal{M})$, we have: $BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i} \mid \mathcal{M}) \neq \emptyset$. In addition, since $\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i}(\Theta\times\{\tilde{\theta}^0\}\times M_{-i}(\bar{\sigma})) = 1$, by construction of $M_i(\bar{\sigma})$ we have $BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i} \mid \mathcal{M}) \subset M_i(\bar{\sigma})$. Thus,

$$BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i} \mid \mathcal{M}(\bar{\sigma})) = BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i} \mid \mathcal{M}).$$

Recall that, by construction of $\pi_{\bar{t}_i}^{1,m_i}$, $m_i \in BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_{\bar{t}_i}^{1,m_i} \mid \mathcal{M}(\bar{\sigma}))$. Consequently,

$$m_i \in BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_i^{0,1} \mid \mathcal{M}).$$

In addition, we have

$$\text{marg}_{\Theta\times M_{-i}}\ \pi_i^{\varepsilon,1} = \text{marg}_{\Theta\times M_{-i}}\ \pi_i^{0,1}.$$

Hence, for $\varepsilon > 0$, by construction of $\pi_i^{\varepsilon,1}$, $\{m_i\} = BR_i(\text{marg}_{\Theta^*\times M_{-i}}\ \pi_i^{\varepsilon,1} \mid \mathcal{M})$ and $\sigma_i(\hat{t}_i[\varepsilon,1,\bar{\sigma},\bar{t}_i,m_i]) = m_i$.

Now, for each $k \geq 2$, proceed by induction and assume that $\sigma_{-i}(\hat{t}_{-i}[\varepsilon,k-1,\bar{\sigma},\bar{t}_{-i},m_{-i}]) = m_{-i}$ for any $\bar{t}_{-i} \in \bar{T}_{-i}$, $m_{-i} \in R_{-i}^{k-1}(\bar{t}_{-i} \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$ and $\varepsilon > 0$. Now fix $\varepsilon > 0$ and $m_i \in R_i^k(\bar{t}_i \mid \mathcal{M}(\bar{\sigma}), \bar{\mathcal{T}})$. For each $\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]$, define the belief

$$\pi_i^{\varepsilon,k} = \hat{\kappa}_{\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]}\circ\gamma_k^{-1} \in \Delta(\Theta^*\times\hat{T}_{-i}\times M_{-i})$$

where $\gamma_k : (\theta^*, \hat{t}_{-i}[\varepsilon,k-1,\bar{\sigma},\bar{t}_{-i},m_{-i}]) \mapsto (\theta^*, \hat{t}_{-i}[\varepsilon,k-1,\bar{\sigma},\bar{t}_{-i},m_{-i}], m_{-i})$.

Note that, by construction, $\pi_i^{\varepsilon,k}$ is the belief of type $\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]$ on $\Theta^*\times\hat{T}_{-i}\times M_{-i}$ when he believes that m_{-i} is played at each $(\theta^*, \hat{t}_{-i}[\varepsilon,k-1,\bar{\sigma},\bar{t}_{-i},m_{-i}])$. Hence, by the induction hypothesis, for each $\varepsilon \geq 0$, $\pi_i^{\varepsilon,k}$ corresponds to beliefs of type $\hat{t}_i[\varepsilon,k,\bar{\sigma},\bar{t}_i,m_i]$ when the equilibrium σ is played. The end of the proof mimics the case $k = 1$. ■

References

[1] Abreu, D. and H. Matsushima, (1992a) "Virtual Implementation in Iteratively Undominated Strategies: Complete Information", *Econometrica*, 60, 993-1008.

[2] Abreu, D. and H. Matsushima, (1992b) "Virtual Implementation in Iteratively Undominated Strategies: Incomplete Information", working paper, available at http://www.princeton.edu/~dabreu/index_files/virtual%20implementation-incomplete.pdf

[3] Artemov, G., T. Kunimoto, and R. Serrano, (2007) "Robust Virtual Implementation with Incomplete Information: Towards a Reinterpretation of the Wilson Doctrine", working paper, available at http://people.mcgill.ca/files/takashi.kunimoto/AKS-RVI-Money-1-alt-space.pdf

[4] Bergemann, D. and S. Morris, (2005) "Robust Mechanism Design", *Econometrica*, 73, 1521-1534.

[5] Bergemann, D. and S. Morris, (2007) "An Ascending Auction for Interdependent Values: Uniqueness and Robustness to Strategic Uncertainty", *American Economic Review Papers and Proceedings*, 97, 125-130.

[6] Bergemann, D. and S. Morris, (2009a) "Rationalizable Implementation", *work in progress.*

[7] Bergemann, D. and S. Morris, (2009b) "Interim Rationalizable Implementation", *work in progress.*

[8] Bergemann, D. and S. Morris, (2009c) "Robust Implementation in Direct Mechanisms", forthcoming in the *Review of Economic Studies.*

[9] Bergemann, D. and S. Morris, (2009d) "Robust Virtual Implementation," *Theoretical Economics*, 4, 45-88.

[10] Bikhchandani, S. (2006) "Ex post implementation in environments with private goods", *Theoretical Economics*, 3, 369–393.

[11] Börgers, T. (1995) "A Note on Implementation and Strong Dominance", in *Social Choice, Welfare and Ethics*, ed. by W. Barnett, H. Moulin, M. Salles, and N. Schofield. Cambridge University Press, Cambridge.

[12] Brandenburger, A., and E. Dekel (1993) "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189-198.

[13] Bull, J., and J. Watson (2007) "Hard Evidence and Mechanism Design," *Games and Economic Behavior*, 58, 75-93.

[14] Carlsson, H. and E. van Damme (1993), "Global Games and Equilibrium Selection", *Econometrica*, 61, 989-1018.

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

[16] Dekel, E., D. Fudenberg, and S. Morris (2006), "Topologies on Types," *Theoretical Economics*, 1, 275-309.

[17] Dekel, E., D. Fudenberg, and S. Morris (2007), "Interim Correlated Rationalizability," *Theoretical Economics*, 2, 15-40.

[18] Deneckere, R. and S. Severinov (2008), "Mechanism Design with Partial State Verifiability", *Games and Economic Behavior*, 64, 487-513.

[19] Di Tillio, A. and E. Faingold (2007), "Uniform Topology on Types and Strategic Convergence", working paper, available at http://www.econ.yale.edu/seminars/microt/mt07/efaingold-071113.pdf

[20] Harsanyi, J. (1967) "Games with Incomplete Information Played by Bayesian Players. Part I: The Basic Model," *Management Science*, 14, 159-182.

[21] Jackson, M.O. (1991): "Bayesian Implementation," *Econometrica*, 59, 461–477.

[22] Jehiel, P., M. Meyer-ter-vehn, B. Moldovanu, and W.R. Zame (2006) "The limits of Ex-post Implementation," *Econometrica*, 74, 585-610.

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

[24] Kartik, N., and O. Tercieux (2009) "Implementation with Evidence: Complete Information", working paper, avaialble at http://www.pse.ens.fr/tercieux/implementation_evidence.pdf

[25] Kunimoto, T. (2008) "How Robust is Undominated Nash Implementation," working paper, available at http://personnel.mcgill.ca/takashi.kunimoto/?View=Publications

[26] Lipman, B. (1994) "A Note on the Implications of Common Knowledge of Rationality," *Games and Economic Behavior*, 6, 114-129

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

[28] Lipman, B. (2005) "Finite Order Implications of Common Priors in Infinite Models," mimeo, available at http://people.bu.edu/blipman/

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

[30] Matsushima, H. (2008), "A New Approach to the Implementation Problem," *Journal of Economic Theory*, 45, 128-144.

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

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

[33] Oyama, D., and O. Tercieux (2005), "Robust Equilibria under Non-Common Priors", working paper, available at http://econpapers.repec.org/paper/clalevrem/843644000000000210.htm

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

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

[36] Wilson, R., (1987) "Game-Theoretic Analyses of Trading Processes," in Advances in Economic Theory: Fifth World Congress, ed. by T. Bewley. Cambridge, U.K.: Cambridge University Press, Chap. 2, 33-70