

Efficient Repeated Implementation II: Incomplete Information*

Jihong Lee

Birkbeck College, London and Yonsei University[†]

Hamid Sabourian

University of Cambridge[‡]

September 2008

Abstract

This paper extends Lee and Sabourian [14] and examines repeated implementation of a social choice function in a general incomplete information environment where agents are infinitely-lived and their preferences are determined stochastically in each period. Our main result establishes how any efficient social choice function can be repeatedly implemented in Bayesian Nash equilibrium. Neither incentive compatibility nor (Bayesian) monotonicity is necessary for repeated implementation.

JEL Classification: A13, C72, C73, D78

Keywords: Bayesian repeated implementation, Bayesian Nash implementation, Efficiency, Incentive compatibility

*The authors are grateful to Matt Jackson, Qingmin Liu, Eric Maskin and Roberto Serrano for their helpful comments and suggestions.

[†]School of Business and Economics, Yonsei University, Seodaemun-gu, Seoul 120-749, Korea, Jihong.Lee@yonsei.ac.kr

[‡]Faculty of Economics, Sidgwick Avenue, Cambridge, CB3 9DD, United Kingdom, Hamid.Sabourian@econ.cam.ac.uk

1 Introduction

In this paper, we develop and analyze a model of repeated implementation with incomplete information, extending the complete information counterpart of our companion paper [14]. A group of infinitely-lived agents possess state-dependent utilities over a set of outcomes and in each period a state is drawn independently from an identical prior distribution. Each agent privately observes some partial contents of a realized state, referred to as his type. No restrictions are imposed on the information structure; an agent's utility may or may not depend on others' private information and information signals may or may not be correlated across agents. A social choice function designates a desired outcome for each state/type profile. We introduce a history-dependent sequence of mechanisms, referred to as a *regime*, and ask if a regime can be found such that every Bayesian Nash equilibrium of the regime generates outcome paths implementing the desired social choice at every realized state.

Extending the arguments of [14], it is established that efficiency continues to be critical for repeated implementation of a social choice function in the incomplete information setup. We then demonstrate how, for a social choice function that is efficient, one can indeed construct a regime that attains the desired repeated implementation properties. This result, therefore, confirms the fundamental departure from one-shot (Bayesian) implementation. It requires neither incentive compatibility (Dasgupta, Hammond and Maskin [8], Myerson [19], d'Aspremont and Gerard-Varet [9], Harris and Townsend [10] among others) nor Bayesian monotonicity (Postlewaite and Schmeidler [21], Palfrey and Srivastava [20], Mookherjee and Reichelstein [18] and Jackson [11]).

Furthermore, our results broaden the scope of efficient repeated implementation in the incomplete information setup beyond the finite horizon, private value case considered by Jackson and Sonnenschein [12]. These authors consider a sequence of linked direct revelation mechanisms in which each agent is *budgeted* in the sense that the distribution of his type announcements over the entire horizon must exactly match the underlying prior distribution. Their result is similar to ours: if the agents are sufficiently patient, and if the horizon is sufficiently long, then every equilibrium payoff profile of the budgeted mechanism approximates the payoff profile of the target social choice function as long as the target is efficient. In contrast to their setup, this paper considers an infinite horizon problem with a general information structure. Also, we derive precise, rather than approximate, repeated

implementation of an efficient social choice function (in terms of payoffs). In addition, our results do not require the discount factor to be arbitrarily large.

We also address the question of existence in the regime constructed to derive our characterization results above. It is straightforward to observe that incentive compatibility ensures existence of an equilibrium that repeatedly implements the social choice function. However, incentive compatibility by no means poses a necessary condition to this end. Restricting our attention to the case of *interdependent* values in which agents' utilities may depend on others' information, the following is established. Instead of incentive compatibility, we assume that the social choice function is *identifiable*. This property means that if implementation occurs according to agents' type announcements and all but one agent reports his type truthfully then, upon learning his utility at the end of the period, someone other than the untruthful odd-man-out will discover that there was a lie. Then, one can build a regime that admits an equilibrium in pure strategies and, moreover, maintains the desired equilibrium properties if the social choice function also satisfies efficiency.¹

The paper is organized as follows. In Section 2, we lay out some basic definitions and notation associated with an implementation problem. Using these, we then describe a Bayesian repeated implementation problem in Section 3. Main findings of our analysis will be presented in Sections 4 and 5. Finally, Section 6 offers some concluding remarks. Appendix contains relegated proofs.

2 Basic definitions and notation

A (one-shot) Bayesian implementation problem is denoted by a collection $\tilde{\mathcal{P}} = [I, A, \Theta, p, (u_i)_{i \in I}]$, where:

- I is a finite set of agents (with some abuse of notation, we shall also use I to represent the cardinality of this set);
- A is a finite set of outcomes;
- $\Theta = \prod_{i \in I} \Theta_i$ is a finite set of states, where Θ_i denotes the finite set of agent i 's *types*; let $\theta_{-i} \equiv (\theta_j)_{j \neq i}$ and $\Theta_{-i} \equiv \prod_{j \neq i} \Theta_j$;

¹The equilibrium constructed here uses standard incentives in repeated play and requires the discount factor to be sufficiently large. The characterization part does not, however, require this, as before.

- p denotes a probability distribution defined on Θ (such that $p(\theta) > 0$ for each θ); for each i , let $p_i(\theta_i) = \sum_{\theta_{-i}} p(\theta_{-i}, \theta_i)$ be the marginal probability of type θ_i and $p_i(\theta_{-i}|\theta_i) = p(\theta_{-i}, \theta_i)/p_i(\theta_i)$ be the conditional probability of θ_{-i} given θ_i ; and
- $u_i : \Theta \times A \rightarrow \mathbb{R}$ is a state-dependent utility function of agent i .

A *social choice function* (SCF), f , in a Bayesian implementation problem $\tilde{\mathcal{P}} = [I, A, \Theta, p, (u_i)_{i \in I}]$ is a mapping $f : \Theta \rightarrow A$ such that $f(\theta) \in A$ for any $\theta \in \Theta$. The *image* of an SCF f is the set

$$f(\Theta) = \{a \in A : a \in f(\theta) \text{ for some } \theta \in \Theta\} .$$

Also, let F denote the set of all possible SCFs and, for an $f \in F$, define

$$F(f) = \{f' \in F | f'(\Theta) \subseteq f(\Theta)\} .$$

Given u_i , the *interim* expected utility of outcome a to agent i of type θ_i is given by

$$v_i(a|\theta_i) \equiv \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(a, \theta_{-i}, \theta_i),$$

and let $v_i(a) \equiv \sum_{\theta_i \in \Theta_i} v_i(a|\theta_i) p(\theta_i)$. Similarly, with slight abuse of notation, define the interim expected utility of an SCF f to agent i of type θ_i as

$$v_i(f|\theta_i) \equiv \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(f(\theta_{-i}, \theta_i), \theta_{-i}, \theta_i),$$

and, hence, let $v_i(f) \equiv \sum_{\theta_i \in \Theta_i} v_i(f|\theta_i) p(\theta_i)$.

Also, define

$$v_i^i \equiv \sum_{\theta_i \in \Theta_i} p_i(\theta_i) \left[\max_{a \in A} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(a, \theta_{-i}, \theta_i) \right] .$$

Notice here that v_i^i is agent i 's maximal expected utility only if he knows his most preferred outcome at the interim stage, which is of course true with *private* values (i.e. when the state-dependent utility can be written as $u_i : \Theta_i \times A \rightarrow \mathbb{R}$ for each i) but not necessarily so with *interdependent* values when each agent's utility may depend on other agents' private information as well as his own.

A *mechanism* (or game form), g , is defined by $g = (M^g, \psi^g)$, where $M^g = M_1^g \times \dots \times M_I^g$ is a cross product of message spaces and $\psi^g : M^g \rightarrow A$ is an outcome function such that

$\psi^g(m) \in A$ for any message profile $m = (m_1, \dots, m_I) \in M^g$. Let G be the set of all feasible mechanisms.

A Bayesian Nash equilibrium of mechanism $g = (M^g, \psi^g)$ is a profile of strategies $s = (s_i)_{i \in I}$ where $s_i : \Theta_i \rightarrow M_i^g$ such that, for all i and all $\theta_i \in \Theta_i$,

$$v_i(\psi^g(s)|\theta_i) \geq v_i(\psi^g(s_{-i}, s'_i)|\theta_i) \text{ for all } s'_i : \Theta_i \rightarrow M_i^g .$$

Let $\mathcal{B}_g(\theta) \subseteq M^g$ be the set of Bayesian Nash equilibrium message profiles of mechanism g in state θ . Define

$$\mathcal{O}_g^{\mathcal{B}}(\theta) = \{a \in A : \exists m \in \mathcal{B}_g(\theta) \text{ such that } \psi^g(m) = a\} .$$

We then say that an SCF f is *Bayesian implementable* if there exists a mechanism g such that $\mathcal{O}_g^{\mathcal{B}}(\theta) = f(\theta)$ for all $\theta \in \Theta$.

Two necessary conditions for one-shot Bayesian implementation are *incentive compatibility* and *Bayesian monotonicity*. An SCF f is incentive compatible if, for any i and any $\theta_i \in \Theta_i$,

$$v_i(f|\theta_i) \geq \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(f(\theta_{-i}, \theta'_i), (\theta_{-i}, \theta_i)) \text{ for all } \theta'_i \in \Theta_i .$$

Bayesian monotonicity extends monotonicity to the incomplete information setup and justifies a change in the social choice change on the basis of a preference reversal around it.²

An *efficient* SCF is defined as in [14]. Let $V = \{(v_i(f))_{i \in I} \in \mathbb{R}^I : f \in F\}$ denote the set of expected utility profiles of all possible SCFs. We write $co(V)$ for the convex hull of V .

Definition 1. *An SCF f is efficient if there exists no $v = (v_1, \dots, v_I) \in co(V)$ such that $v_i \geq v_i(f)$ for all i and $v_i > v_i(f)$ for some i ; f is strictly efficient if, in addition, there exists no $f' \in F$, $f' \neq f$, such that $v_i(f') = v_i(f)$ for all i .*

3 Bayesian repeated implementation

An infinitely repeated Bayesian implementation problem is denoted by $\tilde{\mathcal{P}}^\infty$, representing infinite repetitions of the Bayesian implementation problem $\tilde{\mathcal{P}} = [I, A, \Theta, p, (u_i)_{i \in I}]$. Peri-

²The severe restrictiveness of Bayesian monotonicity is well documented. See, for instance, the survey of Serrano [23].

ods are indexed by $t \in \{1, 2, \dots, \infty\} \equiv \mathbb{Z}_{++}$ and the agents hold a common discount factor $\delta \in (0, 1)$. In each period, the state is drawn from Θ from an independent and identical probability distribution p ; each agent observes only his own type.

An (uncertain) infinite sequence of outcomes is denoted by $a^\infty = (a^{t,\theta})_{t \in \mathbb{Z}_{++}, \theta \in \Theta}$, where $a^{t,\theta} \in A$ is the outcome implemented in period t and state θ . Let A^∞ denote the set of all such sequences. Agents' preferences over alternative infinite sequences of outcomes are represented by discounted average expected utilities, that is, a mapping $\pi_i : A^\infty \rightarrow \mathbb{R}$ such that

$$\pi_i(a^\infty) = (1 - \delta) \sum_{t \in \mathbb{Z}_{++}} \sum_{\theta \in \Theta} \delta^{t-1} p(\theta) u_i(a^{t,\theta}, \theta).$$

It is assumed that the structure of $\tilde{\mathcal{P}}^\infty$ (including the discount factor) is common knowledge among the agents and, if there is one, the social planner.

For a mechanism $g = (M^g, \psi^g)$, define $\mathcal{E}^g \equiv \{(g, m)\}_{m \in M^g}$, and let $\mathcal{E} = \cup_{g \in G} \mathcal{E}^g$, $H^t = \mathcal{E}^{t-1}$ ($H^1 = \emptyset$) and $H^\infty = \cup_{t=1}^\infty H^t$ with its typical element denoted by h . A regime, R , is then a mapping $R : H^\infty \rightarrow G$; $R|h$ refers to the continuation regime that regime R induces at history $h \in H^\infty$. A regime R is *history-independent* if and only if, for any $h, h' \in H^t$ and any t , $R(h) = R(h') \in G$. Notice that, in such a history-independent regime, the specified mechanisms may change over time in a pre-determined sequence. We say that a regime R is *stationary* if and only if, for any $h, h' \in H^\infty$, $R(h) = R(h') \in G$.

In order to define a strategy, let $\mathbf{H}_i^t = (\mathcal{E} \times \Theta_i)^{t-1}$, where $\mathbf{H}_i^1 = \emptyset$, and $\mathbf{H}_i^\infty = \cup_{t=1}^\infty \mathbf{H}_i^t$ with its typical element denoted by \mathbf{h}_i . Then, for any regime R , each agent i 's strategy, σ_i , is a mapping $\sigma_i : \mathbf{H}_i^\infty \times G \times \Theta_i \rightarrow \cup_{g \in G} M_i^g$ such that $\sigma_i(\mathbf{h}_i, g, \theta_i) \in M_i^g$ for any $(\mathbf{h}_i, g, \theta_i) \in \mathbf{H}_i^\infty \times G \times \Theta_i$.³ Let Σ_i be the set of all such strategies, and let $\Sigma \equiv \Sigma_1 \times \dots \times \Sigma_I$. A strategy profile is denoted by $\sigma \in \Sigma$. We say that σ_i is a *Markov* strategy if and only if $\sigma_i(\mathbf{h}, g, \theta) = \sigma_i(\mathbf{h}', g, \theta)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}_i^\infty$, any $g \in G$ and any $\theta \in \Theta$. A strategy profile $\sigma = (\sigma_1, \dots, \sigma_I)$ is Markov if and only if σ_i is Markov for each i . Note here that we are considering *private* strategies; our results below also stand with *public* strategies that depend only on publicly observable histories.

Next, let $\theta(t) = (\theta^1, \dots, \theta^{t-1}) \in \Theta^{t-1}$ denote a sequence of realized states up to, but not including, period t with $\theta(1) = \emptyset$. Let $q(\theta(t)) \equiv p(\theta^1) \times \dots \times p(\theta^{t-1})$. Fix a regime R

³We restrict our attention to pure strategies purely for expositional clarity. In Section 6 below, we discuss behavioral strategies.

and a strategy profile $\sigma \in \Sigma$. We first define the following variables on the outcome path:⁴

- $\mathbf{h}_i(\theta(t), \sigma, R) \in \mathbf{H}_i^t$ denotes the $t - 1$ period history that agent i observes if all agents play R according to σ and the state/type profile realizations are $\theta(t) \in \Theta^{t-1}$; let $\mathbf{h}(\theta(t), \sigma, R) = [\mathbf{h}_i(\theta(t), \sigma, R)]_{i \in I}$.
- $\mu_i(\tilde{\theta}(t) | \mathbf{h}_i(\theta(t), \sigma, R))$ denotes agent i 's posterior *belief* about the other agents' past types, $\tilde{\theta}(t)$, conditional on $\mathbf{h}_i(\theta(t), \sigma, R)$.
- $g^{\theta(t)}(\sigma, R)$ denotes the mechanism played at $\mathbf{h}(\theta(t), \sigma, R)$.
- $m^{\theta(t), \theta^t}(\sigma, R)$ denotes the message profile reported at $\mathbf{h}(\theta(t), \sigma, R)$ when the realized state is θ^t .
- $a^{\theta(t), \theta^t}(\sigma, R)$ denotes the outcome implemented at $\mathbf{h}(\theta(t), \sigma, R)$ when the realized state is θ^t .
- $E\pi_i^{\theta(t), \tau}(\sigma, R)$ denotes agent i 's *expected* continuation payoff at period $\tau \geq t$ conditional on history $\mathbf{h}(\theta(t), \sigma, R)$; that is, for any $t \geq 1$ and any $\tau \geq t$,

$$E\pi_i^{\theta(t), \tau} = (1-\delta) \sum_{s \geq \tau} \sum_{\theta(s-t+1)} \sum_{\theta^s} \sum_{\tilde{\theta}(t)} \delta^{s-\tau} q(\theta(s-t+1), \theta^s) \mu_i(\tilde{\theta}(t) | \mathbf{h}_i(\theta(t))) u_i(a^{\tilde{\theta}(t), \theta(s-t+1), \theta^s}, \theta^s).$$

For simplicity, let $E\pi_i^{\theta(1), \tau}(\sigma, R) \equiv E\pi_i^\tau(\sigma, R)$, $E\pi_i^{\theta(t), t}(\sigma, R) \equiv \pi_i^{\theta(t)}(\sigma, R)$ and $\pi_i(\sigma, R) \equiv \pi_i^{\theta(1)}(\sigma, R)$.

A strategy profile $\sigma = (\sigma_1, \dots, \sigma_I)$ is a Bayesian Nash equilibrium of regime R if, for each i , $\pi_i(\sigma, R) \geq \pi_i(\sigma'_i, \sigma_{-i}, R)$ for all $\sigma'_i \in \Sigma_i$. Let $\mathcal{B}^\delta(R) \subseteq \Sigma$ denote the set of (pure strategy) Bayesian Nash equilibria of regime R with discount factor δ . We introduce the following notion of Bayesian implementation.

Definition 2. *An SCF f is Bayesian-repeated-implementable from period τ if there exists a regime R such that (i) $\mathcal{B}^\delta(R)$ is non-empty; and (ii) every $\sigma \in \mathcal{B}(R, \delta)$ is such that $a^{\theta(t), \theta^t}(\sigma, R) = f(\theta^t)$ for any $t \geq \tau$, any $\theta(t)$ and any θ^t .*

⁴When the meaning is clear, we shall refer to these variables without their arguments.

For simplicity, when we say that f is Bayesian-repeated-implementable, we mean that it is Bayesian-repeated-implementable from period 1. Also, when we say that f is Bayesian-repeated-implementable with sufficiently large discount factor, we mean that there exist R and $\bar{\delta}$ such that, for any $\delta \in (\bar{\delta}, 1)$, the properties in Definition 2 hold.

4 Main results

Let us first establish what is necessary for Bayesian repeated implementation in our setup. It is in fact straightforward to adapt the arguments behind Theorem 1 of [14] to the incomplete information setup. If the agents are sufficiently patient and an SCF f is Bayesian-repeated-implementable, then there cannot be another SCF whose *image* also belongs to that of f such that all agents strictly prefer it to f in expectation. Otherwise, there must be a collusive equilibrium in which the agents obtain higher payoffs but this is a contradiction.

Theorem 1. *There exists $\bar{\delta}$ such that, for any $\delta \in (\bar{\delta}, 1)$, the following holds: if an SCF f is Bayesian-repeated-implementable, then there exists no $f' \in F(f)$ such that $v_i(f') > v_i(f)$ for all i .*

Proof. See the proof of Theorem 1 of [14]. □

We now proceed to examine how an efficient SCF can be Bayesian-repeated-implemented in our setup. Our results below make use of two minor conditions:

Condition $\tilde{\omega}$. For each i , $v_i(f) \leq v_i^i$, and there exists some \tilde{a}^i such that $v_i(f) \geq v_i(\tilde{a}^i)$.

Condition $\tilde{\nu}$. There exist $i, j \in I$, $i \neq j$, such that $v_i(f) < v_i^i$ and $v_j(f) < v_j^j$.

The first property extends condition ω that appears in [14]. In addition to a lower bound on each agent's expected utility from the SCF, we require that the expected utility be bounded above by what the agent could obtain if he were a dictator. Note that this latter property is satisfied vacuously with private values, but the same cannot be said with interdependent values. The second property requires that there be at least two agents who prefer to be dictators than have the SCF implemented.

Condition $\tilde{\omega}$ enables us to extend to the incomplete information setup Lemma 1 of [14], which applies Sorin's [22] observation. That is, given any SCF f satisfying condition $\tilde{\omega}$ and any i , one can construct a history-independent regime such that agent i 's (unique) payoff is

exactly $v_i(f)$, as long as $\delta > \frac{1}{2}$. Such a regime appropriately alternates agent i dictatorship and unconditional enforcement of some outcome. Let the regime continue to be called S^i . Also, let D^i be a stationary regime in which agent i is the dictator and Φ^a a stationary regime in which outcome a is unconditionally implemented forever. In what follows, we shall fix δ above one half unless stated otherwise.

Our main sufficiency result builds on constructing the following regime defined for an SCF f that satisfies condition $\tilde{\omega}$. First, mechanism $b^* = (M, \psi)$ is defined such that (i) for all i , $M_i = \Theta_i \times \mathbb{Z}_+$; and (ii) for any $m = ((\theta_i, z^i))_{i \in I}$, $\psi(m) = f(\theta_1, \dots, \theta_I)$. Then, let B^* represent any regime satisfying the following transition rules:

1. $B^*(\emptyset) = b^*$;
2. For any $h = ((g^1, m^1), \dots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 1$ and $g^{t-1} = b^*$:
 - (a) if $m_i^{t-1} = (\theta_i, 0)$ for all i , then $B^*(h) = b^*$;
 - (b) if there exists some i such that $m_j^{t-1} = (\theta_j, 0)$ for all $j \neq i$ and $m_i^{t-1} = (\theta_i, z^i)$ with $z^i \neq 0$, then $B^*|h = S^i$;
 - (c) if m^{t-1} is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then $B^*|h = D^i$.

This regime is similar to the regimes constructed for the complete information case. The agents in the above regime start by playing mechanism b^* in which each agent reports his type and a non-negative integer. Strategic interaction is maintained, and mechanism b^* continues to be played, only if every agent reports zero integer. If only one agent reports a non-zero integer, that odd-man-out obtains a continuation payoff at the next period exactly equal to what he would obtain from implementation of the SCF. If two or more agents report non-zero integers then the one announcing the highest integer becomes a dictator forever as of the next period.

Characterizing the properties of an equilibrium of the regime is made more complicated by the presence of incomplete information, compared to the corresponding task in the complete information setup. The reason is two-fold.

First, as in the complete information counterpart, we employ the integer play to guarantee a lower bound on each agent's equilibrium continuation payoff. But here, this is done

with incomplete information, and we obtain lower bounds in terms of *expected* continuation payoff at the *next* period.

Second, since we allow for private strategies that depend on each individual's past types, continuation payoffs depend on the agents' posterior beliefs about other agents' past types. This implies that, unlike in the complete information case, we cannot directly apply efficiency and impose upper bounds on continuation payoffs since the agents may not share common posteriors at a given history. Instead, we invoke Bayes' rule, evaluate expected continuation payoffs *at the beginning of* $t = 1$ and then impose efficiency to find upper bounds accordingly. From this, we can show that each expected continuation payoffs evaluated at any other history also has to be bounded above.

Lemma 1. *Fix any $\sigma \in \mathcal{B}^\delta(B^*)$. For any t and any $\theta(t)$, if $g^{\theta(t)}(\sigma, B^*) = b^*$, then $E\pi_i^{\theta(t), t+1}(\sigma, B^*) \geq v_i(f)$ for all i .*

Proof. We prove the claim by contradiction. So, suppose that, at some $\theta(t)$, $g^{\theta(t)}(\sigma, B^*) = b^*$ but $E\pi_i^{\theta(t), t+1}(\sigma, B^*) < v_i(f)$ for some i . Then, consider i deviating to another strategy σ'_i which is identical to the equilibrium strategy σ_i at every history, except at $\mathbf{h}_i(\theta(t), \sigma, B^*)$ where, for each current period realization of θ_i , it reports the same type as in σ_i but a different integer which is higher than any integer that can be reported by σ at such a history.

By the definition of b^* , such a deviation does not alter the current period's implemented outcome, regardless of the other agents' types. As of the next period, it results in either i becoming a dictator forever (transition rule 2(c)) or continuation regime S^i (transition rule 2(b)). By condition $\tilde{\omega}$, we know that $v_i^i \geq v_i(f)$. On the other hand, S^i yields a continuation payoff equal to $v_i(f)$ as of the next period. Thus, the deviation is profitable, and we have a contradiction. \square

Lemma 2. *Suppose that f is efficient (in addition to satisfying condition $\tilde{\omega}$). Fix any $\sigma \in \mathcal{B}^\delta(B^*)$. Also, fix any t , and suppose that $g^{\theta(t)}(\sigma, B^*) = b^*$ for all $\theta(t)$. Then, $E\pi_i^{\theta(t), t+1}(\sigma, B^*) = v_i(f)$ for all i and any $\theta(t)$.*

Proof. Since $g^{\theta(t)} = b^*$ for all $\theta(t)$, it immediately follows from Lemma 1 and Bayes' rule that, for all i ,

$$E\pi_i^{t+1} = \sum_{\theta(t) \in \Theta^{t-1}} q(\theta(t)) E\pi_i^{\theta(t), t+1} \geq v_i(f) . \quad (1)$$

Since $(E\pi_i^{t+1})_{i \in I} \in co(V)$ and f is efficient, (1) then implies that $E\pi_i^{t+1} = v_i(f)$ for all i . By Lemma 1, this in turn implies that $E\pi_i^{\theta(t), t+1} = v_i(f)$ for all i and any $\theta(t)$. \square

Lemma 2 ties down expected continuation payoffs at any given history along the equilibrium path *if* the mechanism played in the corresponding period is b^* for *all* possible state realizations up to the period. It, therefore, remains to be shown that mechanism b^* must always be played along any equilibrium path. Since the regime begins with b^* , we can apply induction to derive this, once the next Lemma has been established.

Lemma 3. *Suppose either (i) that f is efficient and satisfies condition \tilde{v} , or (ii) that f is strictly efficient. Fix any $\sigma \in \mathcal{B}^\delta(B^*)$. Also, fix any t , and suppose that $g^{\theta(t)}(\sigma, B^*) = b^*$ for all $\theta(t)$. Then, $m_i^{\theta(t), \theta^t}(\sigma, B^*) = (\cdot, 0)$ for all i , any $\theta(t)$ and any θ^t .*

Proof. Suppose not. So, suppose that, for some t , $g^{\theta(t)}(\sigma, B^*) = b^*$ for all $\theta(t)$ but $m_i^{\theta(t), \theta^t}(\sigma, B^*) = (\cdot, z)$, $z \neq 0$, for some i , $\theta(t)$ and θ^t .

(i) By condition \tilde{v} , there must exist some $j \neq i$ such that $v_j^j > v_j(f)$. Consider this agent deviating to another strategy which is identical to the equilibrium strategy except at $\mathbf{h}_j(\theta(t), \sigma, B^*)$ where, for each current period realization of θ_j , it reports the same type as in the original equilibrium but a different integer higher than any integer reported by σ at such a history.

By the construction of b^* , the deviation does not alter the current period's outcome, regardless of the other agents' types. However, as of the next period, the continuation regime is as follows: if the realized type profile is θ^t , it is D^j ; otherwise, it is D^j or S^j . In the former case, j obtains a continuation payoff strictly above $v_j(f)$ while, in the latter case, he obtains exactly $v_j(f)$. Since, by Lemma 2, the equilibrium continuation payoff is $v_j(f)$, the deviation is therefore profitable, and we have a contradiction.

(ii) Given the transition rules, the continuation regime will involve dictatorial and/or trivial mechanisms. Unless f can be implemented by non-contingent repetition of a dictatorial or trivial mechanism (the problem would then be trivial), this means that, since f is strictly efficient, $E\pi^{t+1}(\sigma, B^*) \neq v_i(f)$ for some i . But, this contradicts Lemma 2. \square

We are now ready to state the equilibrium properties of regime B^* . If the SCF is efficient/strictly efficient and satisfies the minor additional conditions, every equilibrium of B^* must deliver desired expected continuation payoffs/outcomes at every history beyond the first period.

Proposition 1. Fix any $I \geq 2$.

1. If f is efficient and satisfies conditions $\tilde{\omega}$ and $\tilde{\nu}$, then every $\sigma \in \mathcal{B}^\delta(B^*)$ is such that $E\pi_i^{\theta(t), t+1}(\sigma, B^*) = v_i(f)$ for all i , any t and any $\theta(t)$.
2. If f is strictly efficient and satisfies condition $\tilde{\omega}$, then every $\sigma \in \mathcal{B}^\delta(B^*)$ is such that $a^{\theta(t), \theta^t}(\sigma, B^*) = f(\theta^t)$ for any $t \geq 2$, any $\theta(t)$ and any θ^t .

Proof. 1. Fix any $\sigma \in \mathcal{B}^\delta(B^*)$. Since $B^*(\emptyset) = b^*$, and by Lemma 3, induction implies that, for any t and $\theta(t)$, $g^{\theta(t)}(\sigma, B^*) = b^*$. Lemma 2 then completes the proof.

2. It is straightforward to derive the claim from the same induction arguments and the definition of strict efficiency. \square

Our objective of Bayesian repeated implementation is now achieved if regime B^* admits an equilibrium. One natural sufficiency condition that will guarantee existence in our setup is incentive compatibility. It is straightforward to see that, if f is incentive compatible, regime B^* admits a Markov equilibrium in which each agent always reports his true type and zero integer. Thus, from Proposition 1 above, we can state the following:

Theorem 2. If f is strictly efficient, incentive compatible and satisfies condition $\tilde{\omega}$, then f is Bayesian-repeated-implementable from period 2.

5 More on incentive compatibility

The purpose of incentive compatibility in Theorem 2 is merely to ensure existence of a “good” equilibrium in which every agent always reports his true type. We now explore if it is in fact possible to construct a regime that keeps the desired equilibrium properties and admits such an equilibrium without incentive compatibility.

Our focus is on the case of *interdependent* values, in which some agents’ utilities depend on others’ private information. Here, in the one-shot setup, many authors have identified a conflict between efficiency and incentive compatibility (for instance, Maskin [16] and Jehiel and Moldovanu [13]). It is therefore of particular interest that we establish non-necessity of incentive compatibility for repeated implementation.

Let us assume that the agents know their utilities from the implemented outcomes at the end of each period, and define *identifiability* as follows.

Definition 3. An SCF f is identifiable if, for any i , any $\theta_i, \theta'_i \in \Theta_i$ such that $\theta'_i \neq \theta_i$ and any $\theta_{-i} \in \Theta_{-i}$, there exists some $j \neq i$ such that

$$u_j(f(\theta'_i, \theta_{-i}), \theta'_i, \theta_{-i}) \neq u_j(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i}) .$$

In words, identifiability requires that, whenever there is one agent lying about his type while all others report their types truthfully, there exists another agent who obtains a (one-period) utility different from what he would have obtained under everyone behaving truthfully.⁵ Thus, with an identifiable SCF, if an agent deviates from an equilibrium in which all agents report their types truthfully then there will be at least one other agent who can detect the lie at the end of the period. Notice that the above definition does not require that the detector knows who has lied; he only learns that *someone* has.

Identifiability will enable us to build a regime which admits a truth-telling equilibrium based on incentives of repeated play, instead of one-shot incentive compatibility of the SCF. Such incentives involve punishment when someone misreports his type. To this end, we also introduce a *bad outcome*; that is, there exists a bad outcome \tilde{a} such that $v_i(\tilde{a}) < v_i(f)$ for all i .

Let us now lay out the regime used for our next result. Consider an SCF satisfying $\tilde{\omega}$. Also, suppose that there exists a bad outcome. Define Z as a mechanism in which (i) for all i , $m_i = \mathbb{Z}_+$; and (ii) for all m , $\psi(m) = a$ for some a , and define \tilde{b}^* as the following extensive form mechanism:

- Stage 1 - Each agent i announces his private information θ_i , and $f(\theta_1, \dots, \theta_I)$ is implemented.
- Stage 2 - Once agents learn their utilities, but before a new period begins, each of them announces a report belonging to the set $\{NF, F\} \times \mathbb{Z}_+$, where NF and F refer to “no flag” and “flag” respectively.

The agents’ actions in Stage 2 do not affect the outcome implemented and payoffs in the current period but they determine the continuation play as of the next period in the regime below.⁶

⁵Notice that identifiability cannot hold with private values.

⁶A similar mechanism design is considered by Mezzetti [17] in the one-shot context with interdependent values and quasi-linear utilities. There, in the first stage, the agents are asked to report their private information, followed by immediate outcome implementation; then, in the second stage, the agents report their payoffs and transfers are arranged subsequently.

Let \tilde{B}^* represent any regime satisfying the following transition rules:

1. $\tilde{B}^*(\emptyset) = Z$;
2. For any $h = (g^1, m^1) \in H^2$:
 - (a) if $m_i^1 = (0)$ for all i , then $\tilde{B}^*(h) = \tilde{b}^*$;
 - (b) if there exists some i such that $m_j^1 = (0)$ for all $j \neq i$ and $m_i^1 = z^i$ with $z^i \neq 0$, then $\tilde{B}^*|h = S^i$;
 - (c) if m^1 is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then $\tilde{B}^*|h = D^i$;
3. For any $h = ((g^1, m^1), \dots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 2$ and $g^{t-1} = \tilde{b}^*$:
 - (a) if m_i^{t-1} is such that every agent reports NF and zero integer in Stage 2, then $\tilde{B}^*(h) = \tilde{b}^*$;
 - (b) if m_i^{t-1} is such that at least one agent reports F in Stage 2, then $\tilde{B}^*|h = \Phi^{\tilde{a}}$;
 - (c) if m_i^{t-1} is such that every agent reports NF and every agent except some i announces zero integer, then $\tilde{B}^*|h = S^i$;
 - (d) if m^{t-1} is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then $\tilde{B}^*|h = D^i$.

This regime begins with a simple integer mechanism and non-contingent implementation of an arbitrary outcome in the first period. If all agents report zero integer, then the next period's mechanism is \tilde{b}^* ; otherwise, strategic interaction ends with the continuation regime being either S^i or D^i for some i .

The new mechanism \tilde{b}^* sets up two reporting stages. In particular, each agent is endowed with an opportunity to report detection of a lie by raising “flag” (though he may not know who the liar is) after an outcome has been implemented and his own within-period payoff learned. The second stage also features integer play, with the transitions being the same as before as long as every agent reports “no flag”. But, only one “flag” is needed to overrule the integers and activate permanent implementation of the bad outcome.

Several comments are worth pointing out about regime \tilde{B}^* . First, why do we employ a two-stage mechanism? This is because we want to find an equilibrium in which a deviation

from truth-telling can be identified and subsequently punished. This can be done only after utilities are learned, via the choice of “flag”.

Second, the agents report an integer also in the second stage of mechanism \tilde{b}^* . Note that either a positive integer or a flag leads to shutdown of strategic play in the regime. This means that if we let the integer play to occur before the choice of flag an agent may deviate from truth-telling by reporting a false type and a positive integer, thereby avoiding subsequent detection and punishment altogether.

Third, note that the initial mechanism enforces an arbitrary mechanism and only integers are reported. The integer play affects transitions such that the agents’ continuation payoffs are bounded. We do not however want to endow the agents with the choice of flag immediately because we want to avoid the possibility of *co-ordination failure* where the agents all raise flag.⁷ Moreover, the absence of flags means that we cannot let the agents reveal their types in the first period. This is because of the lack of incentive compatibility; otherwise, one may deviate from truth-telling, obtain a one-period gain and subsequently enjoy the same continuation payoff as the equilibrium.

We now establish that, with an identifiable SCF, regime \tilde{B}^* admits an equilibrium that attains the desired outcome path as long as the agents are sufficiently patient.

Proposition 2. *Suppose that f satisfies identifiability. If δ is sufficiently large, there exists $\sigma^* \in \mathcal{B}^\delta(\tilde{B}^*)$ such that, for any $t > 1$, any $\theta(t)$ and any θ^t , (i) $g^{\theta(t)}(\sigma^*, \tilde{B}^*) = \tilde{b}^*$; and (ii) $a^{\theta(t), \theta^t}(\sigma^*, \tilde{B}^*) = f(\theta^t)$.*

Proof. There exists some $\epsilon > 0$ such that, for each i , $v_i(\tilde{a}) < v_i(f) - \epsilon$. Define

$$\rho \equiv \max_{i \in I, \theta \in \Theta, a, a' \in A} [u_i(a, \theta) - u_i(a', \theta)]$$

and

$$\bar{\delta} \equiv \frac{\rho}{\rho + \epsilon} .$$

Fix any $\delta \in (\bar{\delta}, 1)$.

Consider the following symmetric strategy profile $\sigma^* \in \Sigma$: for each i , σ_i^* is such that:

- $\sigma_i^*(\emptyset, Z, \theta_i) = 0$ for any θ_i ;

⁷We shall use induction arguments to ensure that flags cannot be raised in equilibrium beyond the first period.

- for any $t > 1$ and any corresponding history, if \tilde{b}^* is played in the period,
 - in Stage 1, it always reports the true type;
 - in Stage 2, it reports NF and zero integer if the agent has not detected a false report from another agent or has not made a false report himself in Stage 1; otherwise, report F .

From these strategies, each agent i obtains continuation payoff $v_i(f)$ at the beginning of each period $t > 1$. Let us now examine deviation by any agent i .

First, consider $t = 1$. Given the definition of mechanism Z and transition rule 2(b), announcing a positive integer alters neither the current period's outcome/payoff nor the continuation payoff at the next period.

Second, consider any $t > 1$ and any corresponding history. Deviation can take place in two stages:

(i) Stage 1 - Announce a false type. But then, due to identifiability, another agent will raise “flag” in Stage 2, thereby activating permanent implementation of \tilde{a} as of the next period. The corresponding continuation payoff cannot exceed

$$(1 - \delta) \max_{a, \theta} u_i(a, \theta) + \delta(v_i(f) - \epsilon) , \quad (2)$$

while the equilibrium payoff is at least

$$(1 - \delta) \min_{a, \theta} u_i(a, \theta) + \delta v_i(f) . \quad (3)$$

Since $\delta > \bar{\delta}$, (3) exceeds (2), and the deviation is not profitable.

(ii) Stage 2 - Given transition rules 3(b) and 3(c), neither “flag” nor non-zero integer improves continuation payoff.⁸ □

Note from the proof that, in the above equilibrium, punishment for a lie is activated by two flags, one from the detector and also one from the deviator himself. The mutual optimality of this behavior (given a history of deviation) is supported by an indifference argument; the bad outcome is permanently implemented regardless of one's response to another flag. This is in fact a simplification. Since $v_i(\tilde{a}) < v_i(f)$ for the bad outcome

⁸Notice that the proposed equilibrium strategy profile is also sequentially rational; in particular, in the Stage 2 continuation game following a false report (off-the-equilibrium), it is indeed a best response by either the liar or the detector to “flag” given that there will be another agent also doing the same.

\tilde{a} , for instance, the following modification to the regime makes it strictly optimal for an agent to flag given that there is another flag. If there is only one flag then the continuation regime simply implements \tilde{a} forever; if two or more agents flag then the continuation regime alternates between enforcement of \tilde{a} and dictatorships of those who flag such that those agents obtain payoffs greater than what they would obtain from \tilde{a} forever.

Given Proposition 2, it remains to characterize the properties of an equilibrium of regime \tilde{B} . This task involves similar induction arguments to those behind Theorem 2 above, using only efficiency and condition $\tilde{\omega}$. It also does not depend on δ being arbitrarily large, as is the case for existence.

Theorem 3. *Consider the case of interdependent values. If an SCF f is strictly efficient, identifiable and satisfies condition $\tilde{\omega}$, and if there is a bad outcome, then f is Bayesian repeated-implementable from period 2 with sufficiently large discount factor.*

Proof. See Appendix. □

The case of private values In order to use the incentives of repeated play to overcome incentive compatibility, someone in the group must be able to detect a deviation and subsequently enforce punishment. With interdependent values, this is possible once utilities are learned; with private values, each agent's utility depends only on his own type and hence identifiability cannot hold.

One way to identify a deviation in the private value setup is to observe the distribution of an agent's type reports over a long horizon. By a law of large numbers, the distribution of the actual type realizations of an agent must approach the true prior distribution as the horizon grows. Thus, at any given history, if an agent has made type announcements that differ too much from the true distribution, it is highly likely that there have been lies, and it may be possible to build punishments accordingly such that the desired outcome path is supported in equilibrium.

Similar methods, based on *review strategies*, have been proposed to derive a number of positive results in repeated games (see Chapter 11 of Mailath and Samuelson [15] and the references therein). Extending these techniques to our setup may lead to a fruitful outcome.

The budgeted mechanism of Jackson and Sonnenschein [12] in some sense adopts a similar approach. However, they use budgeting to derive a characterization of equilibrium

properties that every equilibrium payoff profile must be arbitrarily close to the efficient target profile if the discount factor is sufficiently large and the horizon long enough.⁹ Their setup is a finite one and, therefore, standard existence arguments are invoked to show that the budgeted mechanism must have an equilibrium (in mixed strategies).

6 Concluding discussion

In this paper, we have extended the complete information analysis of Lee and Sabourian [14] to the general incomplete information setup. Here again, it is shown that the problem of repeated implementation is fundamentally different from its one-shot counterpart. Efficiency, not incentive compatibility or Bayesian monotonicity, plays a critical role in achieving Bayesian repeated implementation.

Let us conclude with some remarks on our results. First, the solution concept used in our setup is that of Bayesian Nash equilibrium. Thus, our results do not rely on sequential rationality or particular restrictions on off-the-equilibrium belief formation, as in Bergin and Sen [5], Baliga [2] and Brusco [6][7].¹⁰

Second, it is straightforward to incorporate behavioral strategies into our setup. To see this, recall Lemmas 1-3 and notice that their proofs are unaffected by allowing for players' randomization.

Finally, our results are not sensitive to the precise knowledge that each agent possesses about the distribution of other agents' types, a criticism frequently raised against Bayesian implementation in general (where each agent is required to behave optimally against others for the given distribution of types) by, for instance, the literature on *ex post* implementation (see, for example, Bergemann and Morris [3][4]). To see this, recall Lemma 1 and its proof: the lower bound on each agent's equilibrium continuation payoff is established by a deviation argument that is actually independent of others' private information.

On the other hand, constructing the continuation regime S^i invokes knowledge of two things: (i) the target payoff, $v_i(f)$, which requires information on the joint distribution but not necessarily the marginal distributions; and (ii) how such a payoff can be generated

⁹Recall that our approach delivers a sharper equilibrium characterization (i.e. precise, rather than approximate, matching between equilibrium and target payoffs obtained independently of δ) for a general incomplete information environment.

¹⁰Nonetheless, we show that our regimes admit equilibria satisfying sequential rationality.

precisely in a non-strategic way. Whether this task is performed by a planner or by the agents themselves, as in a contractual setting, the potential danger of incorrect specification in our setup does not appear to be as serious as in other Bayesian implementation contexts. A large departure from the true distribution and utility functions allow for a wide range of equilibrium payoffs, but still, they must all be bounded below by the profile corresponding to $\{S^i\}$.

7 Appendix

Proof of Theorem 3 Consider regime \tilde{B}^* defined above. Fix any $\sigma \in \mathcal{B}^\delta(\tilde{B}^*)$. We proceed with the following Lemmas.

Lemma 4. *Fix any $t > 1$. Suppose that $g^{\theta(t)} = \tilde{b}^*$ for all $\theta(t)$ and also that, at any θ^t , every agent reports “no flag” in Stage 2. Then, $E\pi_i^{\theta(t), t+1} = v_i(f)$ for any $\theta(t)$ and all i .*

Proof. The arguments are essentially the same as the proof of Lemma 2 above and, hence, omitted. \square

Lemma 5. *Fix any $t > 1$. Suppose that $g^{\theta(t)} = \tilde{b}^*$ for all $\theta(t)$ and also that, at any θ^t , every agent reports “no flag” in Stage 2. Then, for any $\theta(t)$ and any θ^t ,*

1. $g^{\theta(t), \theta^t} = \tilde{b}^*$; and
2. every agent will report “no flag” at the next period $t + 1$, for any θ^{t+1} .

Proof. The first part can be proved by exactly the same arguments as in Lemma 3 above. To prove the second part, suppose otherwise; so, suppose that some agent reports “flag” in Stage 2 of period $t + 1$ at some $\theta(t)$, θ^t , θ^{t+1} .

Since f is strictly efficient and \tilde{a} is such that $v_i(\tilde{a}) < v_i(f)$ for all i , it must be that $E\pi_i^{t+1} < v_i(f)$ for some i . But, this contradicts the previous Lemma, which implies that $E\pi_i^{t+1} = v_i(f)$ for all i . \square

Since $\tilde{B}^*(\emptyset) = Z$, the same arguments clearly apply also to period 1; that is, it must be that $E\pi_i^2 = v_i(f)$ for all i , $g^{\theta(2)} = \tilde{b}^*$ for any $\theta(2) \in \Theta$ and every agent will report “no flag” in Stage 2 of period 2 for any $\theta(2)$ and any θ^2 .

By Lemma 5, induction then implies that flag is never raised and \tilde{b}^* always played along any equilibrium path beyond the first period; Lemma 4 in turn implies that the desired social choice must always be implemented along an equilibrium path.

References

- [1]
- [2] Baliga, S., "Implementation in Economic Environments with Incomplete Information: The Use of Multi-Stage Games," *Games and Economic Behavior*, 27 (1999), 173-183.
- [3] Bergemann, D. and S. Morris, "Robust Mechanism Design," *Econometrica*, 73 (2005), 1771-1813.
- [4] Bergemann, D. and S. Morris, "Ex Post Implementation," forthcoming in *Games and Economic Behavior*.
- [5] Bergin, J. and A. Sen, "Extensive Form Implementation in Incomplete Information Environments," *Journal of Economic Theory*, 80 (1998), 222-256.
- [6] Brusco, S., "Perfect Bayesian Implementation," *Economic Theory*, 5 (1995), 419-444.
- [7] Brusco, S., "Perfect Bayesian Implementation in Economic Environments," *Journal of Economic Theory*, 129 (2006), 1-30.
- [8] Dasgupta, P., P. Hammond and E. Maskin, "Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46 (1979), 195-216.
- [9] d'Asprement, C. and L-A. Gerard-Varet, "Incentives and Incomplete Information," *Journal of Public Economics*, 11 (1979), 25-45.
- [10] Harris, M. and R. Townsend, "Resource Allocation with Asymmetric Information," *Econometrica*, 49 (1981), 33-64.
- [11] Jackson, M. O., "Bayesian Implementation," *Econometrica*, 59 (1991), 461-477.
- [12] Jackson, M. O. and H. F. Sonnenschein, "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75 (2007), 241-257.
- [13] Jehiel, P. and B. Moldovanu, "Efficient Design with Interdependent Valuations," *Econometrica*, 69 (2001), 1237-1259.

- [14] Lee, J. and H. Sabourian, “Efficient Repeated Implementation I: Complete Information,” *mimeo* (2008).
- [15] Mailath, G. J. and L. Samuelson, *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press (2006).
- [16] Maskin, E., “Auctions and Privatization,” in *Privatization*, ed. by H. Siebert. Institut für Weltwirtschaft an der Universität Kiel.
- [17] Mezzetti, C., “Mechanism Design with Interdependent Valuations: Efficiency,” *Econometrica*, 72 (2004), 1617-1626.
- [18] Mookherjee, D. and S. Reichelstein, “Implementation via Augmented Revelation Mechanisms,” *Review of Economic Studies*, 57 (1990), 453-475.
- [19] Myerson, R. B., “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47 (1979), 61-73.
- [20] Palfrey, T. R. and S. Srivastava, “Implementation with Incomplete Information in Exchange Economies,” *Econometrica*, 57 (1989), 115-134.
- [21] Postlewaite, A. and D. Schmeidler, “Implementation with Differential Information Economies,” *Journal of Economic Theory*, 39 (1986), 14-33.
- [22] Sorin, S., “On Repeated Games with Complete Information,” *Mathematics of Operations Research*, 11 (1986), 147-160.
- [23] Serrano, R., “The Theory of Implementation of Social Choice Rules,” *SIAM Review*, 46 (2004), 377-414.