Repeated Implementation with Incomplete Information*

Jihong Lee†
Seoul National University

Hamid Sabourian‡
University of Cambridge

September 2013

Abstract

We formulate and analyze a model of (full) repeated implementation with incomplete information. 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 and may condition behavior on his private information as well as publicly observable histories. It is shown that, with minor qualifications, a social choice function (SCF) that is efficient in the range and incentive compatible can be repeatedly implemented in Bayesian Nash equilibrium under the general information structure. When the agents' utilities are interdependent, incentive compatibility can be replaced with a payoff identifiability condition. We also show that efficiency in the range is sufficient for approximate repeated implementation in public strategies when the information structure satisfies certain statistical identifiability properties.

JEL Classification: A13, C72, C73, D78

*Some of the results in this paper appeared in an earlier working paper [20].

†Department of Economics, Seoul National University, Seoul 151-746, Korea; jihonglee@snu.ac.kr. Jihong Lee's research was supported by the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2011-330-B00063).

‡Faculty of Economics, Cambridge, CB3 9DD, United Kingdom and Fung Visiting Professor, HKUST; Hamid.Sabourian@econ.cam.ac.uk

Keywords: Repeated implementation, Incomplete information, Bayesian Nash implementation, Efficiency, Incentive compatibility, Payoff identifiability, Pairwise identifiability

1 Introduction

The literature on mechanism design and implementation seeks to provide theoretical underpinnings for the scope of various social objectives that can be achieved through individuals' strategic interactions within institutions. From public good provision and auctions to matching markets, the sphere of real world applications of this literature is large and diverse. Despite this, most of this literature to date has been concerned solely with achieving social goals in the short-term. Recently, Lee and Sabourian [21] (hereafter, LS) address the question of *full* implementation in an infinitely repeated model with randomly evolving preferences.[1] This study demonstrates a fundamental difference between one-shot and repeated implementation problems but their analysis is confined to the special case of complete information. The role of informational asymmetries among the participants of long-run institutions is yet to be addressed in full.

In this paper, we formulate and analyze a model of (full) repeated implementation with incomplete information. 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 (SCF) designates a desired outcome for each 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 always generates outcome paths implementing the desired social choice at every history of realized states. The agents are allowed to use private strategies that condition behavior on their private information as well as publicly observable histories.

We focus on the possibility of repeatedly implementing efficient SCFs in the presence

[1] See also Mezzetti and Renou [27] who consider the same repeated implementation problem in any finite or infinite horizon model.

of incomplete information. We know from the one-shot mechanism design literature that with incomplete information the task of ensuring the existence of a desired equilibrium of a mechanism can itself impose a serious hurdle on the scope of achievable social goals: incentive compatibility is necessary (Dasgupta, Hammond and Maskin [8], Myerson [29], d'Aspremont and Gerard-Varet [9], Harris and Townsend [14] among others). Full implementation in addition requires Bayesian monotonicity that is not only difficult to satisfy but also very hard to check (Postlewaite and Schmeidler [31], Palfrey and Srivastava [30], Mookherjee and Reichelstein [28] and Jackson [16]; see also Serrano [33]). Following LS who considered the special case of complete information, our results demonstrate that the monotonicity requirement continues to be neither necessary nor sufficient for the general repeated implementation environment with incomplete information. Repeated game arguments also allow us to design appropriate intertemporal incentives to construct equilibria of dynamic mechanisms that repeatedly implement efficient SCFs that need not satisfy the standard (one-shot) incentive compatibility property.

Our first main results establish that, with minor qualifications, an SCF that is *efficient in the range* and incentive compatible can be repeatedly implemented in Bayesian Nash equilibrium for the general incomplete information environment. Unlike in the complete information setup of LS, the presence of incomplete information poses a new challenge that the agents' behavior depends on their private information, and hence, the continuation payoffs must reflect their posterior beliefs on the others' past information. Nonetheless, by evaluating repeated implementation in terms of *expected* continuation payoffs computed at the beginning of a regime, and then invoking the efficiency property, we are able to pin down equilibrium continuation payoffs at every history for our regime construction without having to track beliefs. Incentive compatibility ensures the existence of a truth-telling equilibrium.

We next investigate the scope of efficient repeated implementation without incentive compatibility. In the interdependent value case, where the cost of incentive compatibility can be particularly severe against efficiency (e.g. Maskin [25] and Jehiel and Moldovanu [18]), the same set of results are obtained by replacing incentive compatibility with an intuitive condition that we call *payoff identifiablity*, if the agents are sufficiently patient. Payoff identifiability requires that when agents announce types and all but one agent reports his type truthfully then, upon learning his utility at the end of the period, some other player, upon learning his utility at the end of the period, will discover that there was a lie. Given such an identifiability property, we construct another regime that, while maintaining the desired payoff properties of its equilibrium set, admits a truth-telling equilibrium based on incentives of repeated play instead of one-shot incentive compatibility of the SCF.

Another avenue for designing correct incentives in repeated implementation is offered by the literature on (complete information) repeated games with imperfect monitoring. We adapt the linear programming arguments of Fudenberg, Levine and Maskin [12] (hereafter, FLM) to demonstrate how an efficient SCF can be repeatedly implemented in terms of expected continuation payoffs for both private and interdependent value cases if an arbitrarily small departure from exact implementation is allowed at each equilibrium history, and if the agents' discount factor is close enough to 1. This approach to obtain *approximate* (payoff) repeated implementation requires that deviations from desired behavior can be statistically distinguished; one way to ensure this is to invoke FLM's *pairwise identifiability* condition on the information structure. While FLM already showed that their techniques for repeated games with imperfect monitoring can also be used to construct equilibria in repeated adverse selection games with private values and independent types, in this paper we obtain full repeated implementation results for the case of independent private values and also beyond when the players use public strategies. In particular, with interdependent values, we show how an enforceability condition from the mechanism design literature due to Crémer and McLean [7] can be employed in tandem with pairwise identifiability in our repeated context to support an equilibrium which achieves continuation payoffs arbitrarily close to the efficient target payoffs at every history.[2]

Our paper is related to the work of Jackson and Sonnenschein [17], who consider linking a specific, *independent private values*, Bayesian implementation problem with a large, but finite, number of independent copies of itself. If the linkage takes place through time, their setup can be interpreted as a particular finitely repeated implementation problem. The authors restrict their attention to a sequence of revelation mechanisms in which each agent is budgeted in his choice of messages according to the prior distribution over his possible types. They find that for any *ex ante* Pareto efficient SCF all equilibrium payoffs of such a budgeted mechanism must approximate the target payoff profile corresponding to the SCF, as long as the agents are sufficiently patient and the horizon sufficiently long.

In contrast to Jackson and Sonnenschein [17], our setup deals with infinitely-lived

[2]We would like to thank Satoru Takahashi for a suggestion that led to this result.

agents and a fully general information structure that allows for interdependent values as well as correlated signals. In terms of the results, we are concerned with repeated implementation of outcomes and/or payoffs at every possible history of the regime, not just in terms of payoffs computed at the outset. Also, our exact implementation results do not require the discount factor to be arbitrarily large. Furthermore, these results are obtained with arguments that are very much distinct from those of [17].

In a concurrent, independent study, Renou and Tomala [32] propose an alternative concept of approximate implementation in an infinite horizon context. They ask whether it is possible to find a regime such that the probability of infinite equilibrium histories in which the desired SCF is almost repeatedly implemented approaches 1 as the players become perfectly patient. Under this concept, a sufficiency result is derived for SCFs that are efficient only among the space of payoffs given by undetectable deceptions. Their setup assumes the case of independent private values and allows for Markovian evolution of private information.

While a main body of our paper studies the issue of repeated implementation, there are also several important differences between the *approximate* implementation analysis of Renou and Tomala [32] and that of this paper. First, we require the continuation payoffs to approximate the target payoffs at *every* equilibrium history; in [32], albeit rarely, the players may still end up in equilibrium paths along which the implemented outcomes and hence continuation payoffs diverge far from the desired outcomes and payoffs after certain histories. Second, our techniques are distinct from those of [32], and they enable us to derive approximate implementation results for the case of interdependent values as well as correlated types, unlike [32] whose construction critically depends on the assumption of independent private values.[3]

The paper is organized as follows. In Section 2, we lay out the problem of repeated implementation with incomplete information. Section 3 presents our sufficiency results for repeatedly implementing SCFs that are efficient in the range as well as being incentive compatible. Section 4 then establishes results on how efficient SCFs can be repeatedly implemented without satisfying incentive compatibility. Section 5 offers some concluding remarks. Appendix contains relegated proofs.

[3] Also, we obtain our approximate implementation results without considering random mechanisms.

2 The Setup

2.1 Preliminaries

A one-shot implementation problem with incomplete information is denoted by $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i\in I}]$ with the following details:

- I is a finite, non-singleton set of agents with typical element denoted by i.
- A is a finite set of outcomes with typical element denoted by a.
- $\Theta = \Pi_{i\in I}\Theta_i$ is a finite, non-singleton 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 $\Theta_{-i} \equiv \Pi_{j\neq i}\Theta_j$.
- p is 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. We say that types are *independent* if $p(\theta) = \prod_i p_i(\theta_i)$ for i and *correlated* otherwise.
- $u_i : \Theta \times A \to \mathbb{R}$ is the *interdependent value* utility function for i. In the *private value* case, we write $u_i : \Theta_i \times A \to \mathbb{R}$. We sometimes write $u = (u_1, \ldots, u_I)$.

An SCF f in an implementation problem $\mathcal{P}$ is a mapping $f : \Theta \to A$, and the *range* of f is the set $f(\Theta) = \{a \in A : a = f(\theta) \text{ for some } \theta\}$. Let F denote the set of all possible SCFs, and for any $f \in F$, define $F(f) = \{f' \in F : f'(\Theta) \subseteq f(\Theta)\}$ as the set of all SCFs whose ranges belong to $f(\Theta)$.

For an outcome $a \in A$, define $v_i(a|\theta_i) = \sum_{\theta_{-i}\in\Theta_{-i}} p_i(\theta_{-i}|\theta_i)u_i(a, \theta_{-i}, \theta_i)$ as its *interim* (one-shot) expected utility, or payoff, to agent i of type θ_i. The corresponding ex ante payoff of outcome a is then defined by $v_i(a) = \sum_{\theta_i\in\Theta_i} v_i(a|\theta_i)p(\theta_i)$.

Similarly, with slight abuse of notation, define respectively the interim and ex ante payoffs of an SCF $f : \Theta \to A$ to agent i of type θ_i as $v_i(f|\theta_i) = \sum_{\theta_{-i}\in\Theta_{-i}} p_i(\theta_{-i}|\theta_i)u_i(f(\theta_{-i}, \theta_i), \theta_{-i}, \theta_i)$ and $v_i(f) = \sum_{\theta_i\in\Theta_i} v_i(f|\theta_i)p(\theta_i)$. Denote by $v(f) = (v_i(f))_{i\in I}$ the profile of (ex ante) payoffs associated with f . Let $V = \left\{v(f) \in \mathbb{R}^I : f \in F\right\}$ be the set of expected utility profiles of all possible SCFs. Also, for any $X \subset \mathbb{R}^I$, let $co(X)$ refer to the convex hull of X and $comp(X) = \{w \in R^I \mid w \leqq w' \text{ for some } w' \in X\}$ be the comprehensive cover of X.

LS define efficiency of an SCF in terms of the convex hull of the set of expected utility profiles of all possible SCFs since this reflects the set of (discounted average) payoffs that can be obtained in an infinitely repeated implementation problem. A payoff profile $v' = (v'_1, .., v'_I) \in co(V)$ is said to Pareto dominate another profile $v = (v_1, .., v_I)$ if $v'_i \geq v_i$ for all i with the inequality being strict for at least one agent; v' *strictly* Pareto dominates v if the inequality is strict for *all* i.

Definition 1 *(a) An SCF f is efficient if there exists no $v' \in co(V)$ that Pareto dominates $v(f)$.*

(b) An SCF f is strictly efficient if it is efficient and there exists no $f' \in F$, $f' \neq f$, such that $v(f') = v(f)$.

With each agent privately observing his type, the set of feasible payoffs in the range of a given SCF f is defined by

$$V(f) = \left\{ v : v = \sum_{\theta_1,\ldots,\theta_I \in \Theta} p(\theta_1, \ldots, \theta_I) u\left(f(\lambda_1(\theta_1), \ldots, \lambda_I(\theta_I)), \theta_1, \ldots, \theta_I\right) \text{ for some } \lambda_i : \Theta_i \to \Theta_i \right\}.$$ [4]

We next define *efficiency in the range* for our incomplete information setup.

Definition 2 *(a) An SCF f is efficient in the range if there exists no $v' \in co(V(f))$ that Pareto dominates $v(f)$.*

(b) An SCF f is strictly efficient in the range if it is efficient in the range and there exists no $f' \in F(f)$, $f' \neq f$, such that $v(f') = v(f)$.

In addition to these notions of efficiency, our incomplete information analysis below makes use of a stronger efficiency condition.

Definition 3 *An SCF f is strongly efficient (in the range) if it is strictly efficient (in the range) and $v(f)$ is an extreme point of $co(V)(co(V(f))$.*

Thus, if f is strongly efficient then there is no way of achieving exactly $v(f)$ via a convex combination of other SCFs.

[4] In the complete information setting of LS, the feasible payoffs in the range of f is defined as $V(f) = \{v = \sum_\theta p(\theta) u(f(\lambda(\theta)), \theta) \text{ for some } \lambda : \Theta \to \Theta\}$. Our results below hold also for SCFs that satisfy efficiency with respect to this alternative definition of $V(f)$.

2.2 Bayesian Repeated Implementation

An infinitely repeated implementation problem, $\mathcal{P}^\infty$, represents infinite repetitions of $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$. In each period, the state is drawn from Θ from an independent and identical probability distribution p and each agent observes only his own type. Each agent evaluates an infinite outcome sequence, $a^\infty = \left(a^{t,\theta}\right)_{t \in \mathbb{Z}_{++}, \theta \in \Theta}$, according to discounted average expected utilities with common discount factor $\delta \in (0,1)$. We assume that the structure of $\mathcal{P}^\infty$ (including the discount factor) is common knowledge among the agents and, if there is one, the planner.

In this paper, we consider deterministic mechanisms and regimes. A mechanism is defined as $g = (M^g, \psi^g)$, where $M^g = M_1^g \times \cdots \times M_I^g$ is a cross product of action/message spaces and $\psi^g : M^g \to A$ is an outcome function such that $\psi^g(m) \in A$ for any message profile $m = (m_1, \ldots, m_I) \in M^g$.[5] Let G be the set of all feasible mechanisms.

A regime specifies a history-dependent "transition rules" of mechanisms contingent on the publicly observable history of mechanisms played and the agents' corresponding actions. It is assumed that a planner, or the agents themselves, can commit to a regime at the outset.

Given mechanism $g = (M^g, \psi^g)$, define $\mathcal{E}^g \equiv \{(g, m)\}_{m \in M^g}$, and let $\mathcal{E} = \cup_{g \in G} \mathcal{E}^g$. Then, $H^t = \mathcal{E}^{t-1}$ (the $(t-1)$-fold Cartesian product of $\mathcal{E}$) represents the set of all possible *public* histories over $t-1$ periods. The initial history is empty (trivial) and denoted by $H^1 = \emptyset$. Also, let $H^\infty = \cup_{t=1}^\infty H^t$ with a typical public history denoted by $h \in H^\infty$.

We define a regime, R, as a mapping $R : H^\infty \to G$. Let $R|h$ refer to the *continuation regime* that regime R induces at history $h \in H^\infty$ (thus, $R|h(h') = R(h, h')$ for any $h, h' \in H^\infty$). We say that a regime R is *history-independent* if and only if, for any t and any $h, h' \in H^t$, $R(h) = R(h')$, and that a regime R is *stationary* if and only if, for any $h, h' \in H^\infty$, $R(h) = R(h')$.

Given a regime, an agent can condition his actions on the past history of realized private information as well as the publicly observed history of mechanisms and message profiles played. Let $\theta(t) = (\theta^1, \ldots, \theta^{t-1}) \in \Theta^{t-1}$ denote a sequence of realized states up to, but not including, period t with $\theta(1) = \emptyset$, and $q(\theta(t)) \equiv p(\theta^1) \times \cdots \times$

[5]Note that this is a normal form representation of a mechanism. If one considers a mechanism in extensive form, M_i can be interpreted as the set of player i's pure strategies in the mechanism and the outcome function, ψ, as a mapping from the set of all possible observable paths induced by the players' strategy profiles.

$p(\theta^{t-1})$. To define a strategy, define the set of all possible *full* histories at date t by $\mathbf{H}^t = H^t \times \Theta^{t-1}$, where $\mathbf{H}^1 = \emptyset$. Let $\mathbf{H}^\infty = \cup_{t=1}^\infty \mathbf{H}^t$ with its typical element denoted by $\mathbf{h}$. For any $\mathbf{h} = (g^1, m^1, \theta^1, \ldots, g^{t-1}, m^{t-1}, \theta^{t-1}) \in \mathbf{H}^t$, we denote by $\mathbf{h}_i = (g^1, m^1, \theta_i^1, \ldots, g^{t-1}, m^{t-1}, \theta_i^{t-1}) \in \mathbf{H}_i^t = H^t \times \Theta_i^{t-1}$ the private history that i observes while $h = (g^1, m^1, \ldots, g^{t-1}, m^{t-1}) \in H^t$ denotes the public history.

For any regime R and any agent i, a mixed (behavioral) strategy, σ_i, is a mapping $\sigma_i : \mathbf{H}^\infty \times G \times \Theta_i \to \cup_{g \in G} \Delta M_i^g$ such that

- $\sigma_i(\mathbf{h}, g, \theta_i) \in \Delta M_i^{R(h)}$ for any $(\mathbf{h}, g, \theta_i) \in \mathbf{H}^\infty \times G \times \Theta_i$; and
- $\sigma_i(\mathbf{h}, g, \theta_i) = \sigma_i(\mathbf{h}', g, \theta_i)$ if $\mathbf{h}_i = \mathbf{h}_i'$.

Let Σ_i be the set of all strategies, and denote a strategy profile by $\sigma = (\sigma_i)_{i=1}^I \in \Sigma \equiv \Sigma_1 \times \cdots \times \Sigma_I$. Also, we say that σ_i is a Markov (history-independent) strategy if and only if $\sigma_i(\mathbf{h}, g, \theta_i) = \sigma_i(\mathbf{h}', g, \theta_i)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ and $\theta_i \in \Theta_i$. A strategy profile σ is Markov if and only if σ_i is Markov for each i. Note that here we are considering the general case with *private* strategies, but all our results below are obtained with public strategies as well. Since our mechanisms are deterministic, we shall sometimes suppress the dependence of strategies on g (except for the first period).

For any regime R and strategy profile $\sigma \in \Sigma$, we define the following variables on the outcome path:

- $\mathbf{H}^t(\sigma, R)$ is the set of $t-1$ period full histories that occur with positive probability if all agents play R according to σ.
- $g^{\mathbf{h}}(R) = \left(M^{R(h)}, \psi^{R(h)}\right)$ denotes the mechanism played after full history $\mathbf{h} = (h, \theta(t)) \in \mathbf{H}^\infty$.
- $A^{\mathbf{h},\theta}(\sigma, R)$ denotes the set of outcomes that occur with positive probability after full history $\mathbf{h} = (h, \theta(t)) \in \mathbf{H}^\infty$ followed by realization of θ.
- For any $t \geq 1$ and $\mathbf{h} = (h, \theta(t)) \in \mathbf{H}^t(\sigma, R)$, $E_{\mathbf{h}} \pi_i^\tau(\sigma, R)$ denotes agent i's *expected* continuation payoff (i.e. discounted average expected utility) at period $\tau \geq t$ conditional on observing the public history h and private information $\theta_i(t)$ and all playing

according to σ. Thus, we have

$$
\begin{aligned}
E_{\mathbf{h}}\pi_i^\tau(\sigma, R) \;=\; & (1-\delta) \sum_{\tilde{\theta}(t)=(\theta^1,\ldots,\theta^{t-1})} \mu_i\left(\tilde{\theta}(t)|h, \theta_i(t)\right) \\
& \sum_{s\geq\tau} \sum_{\theta(s-t+1)=(\theta^t,\ldots,\theta^s)} \sum_{h'=(m^t,\ldots,m^s)} q\left(\theta(s-t+1)\right)\mathbb{P}(h' \mid h, \theta(t), \theta(s-t+1); \sigma, R) \\
& \delta^{s-\tau} u_i\left(\psi^{R(h,m^t,\ldots,m^{s-1})}(m^s), \theta^s\right),
\end{aligned}
$$

where

- $\mu_i\left(\tilde{\theta}(t)|h, \theta_i(t)\right) = \mathbb{P}\left(h, \theta_i(t) \mid \tilde{\theta}(t)\right) q\left(\tilde{\theta}(t)\right) / \mathbb{P}\left(h, \theta_i(t)\right)$ represents agent i's posterior *belief* about all agents' past types conditional on i observing $(h, \theta_i(t))$ and all playing according to σ;
- $\mathbb{P}(h' \mid h, \theta(t), \theta(s-t+1); \sigma, R)$ represents the probability that $(s-t)$-period public history h' is generated by σ in the continuation game after full history $(h, \theta(t))$ if the subsequent state realizations are $\theta(s-t+1)$.

For simplicity, let $E_{\mathbf{h}^1}\pi_i^\tau(\sigma, R) = E\pi_i^\tau(\sigma, R)$ and $E\pi_i^1(\sigma, R) = \pi_i(\sigma, R)$.

When the meaning is clear, we shall sometimes suppress some of the arguments in the above variables and refer to them simply as $\mathbf{H}(\theta(t))$, $\mathbf{H}(\mathbf{h}, \theta(t))$, $g^{\mathbf{h}}$, $A^{\mathbf{h},\theta}$ and $E_{\mathbf{h}}\pi_i^\tau$.

The standard solution concept for implementation with incomplete information is Bayesian Nash equilibrium. In our repeated setup, a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is a Bayesian Nash equilibrium of R if, for any i, $\pi_i(\sigma, R) \geq \pi_i(\sigma_i', \sigma_{-i}, R)$ for all σ_i'. Let $Q^\delta(R)$ denote the set of Bayesian Nash equilibria of regime R with discount factor δ. Similarly to LS, we propose the following notions of repeated implementation for the incomplete information case.

Definition 4 *(a) An SCF f is payoff-repeatedly implementable in Bayesian Nash equilibrium from period t if there exists a regime R such that $Q^\delta(R)$ is non-empty and every $\sigma \in Q^\delta(R)$ is such that, for every $\tau \geq t$, $E\pi_i^\tau(\sigma, R) = v_i(f)$ for any i.*

(b) An SCF f is repeatedly implementable in Bayesian Nash equilibrium from period t if there exists a regime R such that $Q^\delta(R)$ is non-empty and every $\sigma \in Q^\delta(R)$ is such that $A^{\mathbf{h}^\tau,\theta^\tau}(\sigma, R) = \{f(\theta^\tau)\}$ for any $\tau \geq t$, $\mathbf{h}^\tau \in \mathbf{H}^\tau(\sigma, R)$ and $\theta^\tau \in \Theta$.

Thus, we require that the continuation payoff profile or the implemented outcome be exactly matching the desired one at *every* equilibrium history of a regime. Note that the definition of payoff implementation above is written in terms of expected *future* payoffs evaluated at the beginning of a regime. As we shall see below, this is because we want to avoid the issue of agents' ex post private beliefs that affect their continuation payoffs at different histories.

2.3 Obtaining Target Payoffs

Before embarking on our main analysis, we present an observation about the repeated implementation setup that will play an important role in the constructive arguments below. In our setup, the planner can implement the payoff profile of a fixed outcome. Another payoff profile that can be implemented is that of a dictatorship if the agents are rational and the dictator's maximal outcome in each state generates a unique payoff profile. More generally, the planner can enforce a dictatorship over a restricted range of outcomes. Moreover, if the agents are sufficiently patient, by enforcing a non-stationary sequence of dictatorships, the planner can obtain any convex combination of the above payoff profiles.

Fix any subset of outcomes $C \subseteq A$. Let $d^i(C)$ denote i-dictatorship over C (i.e. a mechanism in which $M^i = C$ and $\psi(a) = m_i$ for all $a \in C$) and $D^i(C)$ denote a stationary regime that repeatedly enforces $d^i(C)$. If C is a singleton, $d^i(C)$ enforces a constant SCF. Let

$$v_i^i(C|\theta_i) = \max_{a \in C} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(a, \theta_{-i}, \theta_i)$$

denote the maximum interim expected payoff to agent i of type θ_i under i-dictatorship over C; also, let

$$\mathcal{A}(\theta_i, C) \equiv \arg\max_{a \in C} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(a, \theta_{-i}, \theta_i)$$

represent the set of i's best outcomes in state θ_i. The maximum ex ante payoff of agent i is then given by

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

Note that with private values, for any $f \in F$ and $i \in I$, we have $v_i^i(C) \geq v_i(f)$ if $f(\Theta) \subset C$.

Define the maximum payoff that agent i of type θ_i can obtain when agent $j \neq i$ is the dictator over $C \subseteq A$ by

$$v_i^j(C|\theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) \max_{a \in \mathcal{A}(\theta_j, C)} u_i(a, \theta_{-i}, \theta_i).$$

Then, we write

$$\begin{aligned} v_i^j(C) &= \sum_{\theta \in \Theta} p(\theta) \max_{a \in \mathcal{A}(\theta_j, C)} u_i(a, \theta) \\ &= \sum_{\theta_i \in \Theta_i} p_i(\theta_i) \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i}|\theta_i) \max_{a \in \mathcal{A}(\theta_j, C)} u_i(a, \theta_{-i}, \theta_i) \\ &= \sum_{\theta_i \in \Theta_i} p_i(\theta_i) v_i^j(C|\theta_i). \end{aligned}$$

Let $v^j(C) = (v_1^j(C), \ldots, v_I^j(C))$. Clearly, with private values, $v_i^i(C|\theta_i) \geq v_i^j(C|\theta_i)$ and $v_i^i(C) \geq v_i^j(C)$.

We next introduce the following properties that will enable us to repeatedly implement an efficient SCF with incomplete information.

Condition ω: For each $i \in I$, there exists $\tilde{a}^i \in A$ such that $v_i(\tilde{a}^i) \leq v_i(f)$ and $v(\tilde{a}^i) \in$ *comp*(*co*($V(f)$)).

This condition is similar to the corresponding condition with the same name in LS, except for an additional restriction that $v(\tilde{a}^i)$ belongs to the comprehensive cover of the set *co*($V(f)$), i.e. $v(\tilde{a}^i)$ does not Pareto dominate any efficient profile in *co*($V(f)$). If we were to consider an SCF that is (globally) efficient, instead of being efficient in the range, the additional restriction could be dropped for our results below. Note that condition ω is satisfied in many important applications (e.g. zero consumption in allocation problems).

Condition υ: For each $i \in I$, there exists $C^i \subseteq A$ with the following properties:

(a) $v_i^i(C^i) \geq v_i(f)$ and $v^i(C^i) \in$ *comp*(*co*($V(f)$)).

(b) There exist two agents i and j such that $v_i^i(C^i) > v_i(f)$ and $v_j^j(C^i) > v_j(f)$.

One natural restriction within condition υ would be to let $C^i = f(\Theta)$ for each i, in which case it immediately follows that $v^i(f(\Theta)) \in comp(co(V(f)))$; also, the first part of (a) is vacuously satisfied with private values (or with complete information). The second part of (a) is again utilized to ensure implementation of an SCF that is efficient in the range and would not be needed under an efficient SCF.

As in LS, we utilize conditions ω and υ to construct a history-independent regime for each agent from which he obtains exactly his *target* payoff from any given SCF.

Lemma 1 *Consider any $f \in F$ that satisfies conditions ω and υ. Fix any $i \in I$, and suppose that $\delta \in \left(\frac{1}{2}, 1\right)$. Then, there exists a history-independent regime, denoted by S^i, that generates a unique (discounted average) payoff to agent i equal to $v_i(f)$.*

Proof. Take $\tilde{a}^i$ from condition ω, and consider $d^i(\{\tilde{a}^i\})$, i.e. a constant mechanism that always enforces $\tilde{a}^i$. Also, consider $d^i\left(C^i\right)$, where C^i is taken from condition υ. Since $v_i(\tilde{a}^i) \leq v_i(f) \leq v_i^i\left(C^i\right)$, there exists a history-independent sequence of $d^i(\{\tilde{a}^i\})$ and $d^i\left(C^i\right)$ such that the corresponding discounted average payoff profile equals $v_i(f)$ if $\delta > \frac{1}{2}$. See Sorin [34] (or Lemma 3.7.1 of Mailath and Samuelson [24]). ■

Our results below require δ to be large enough in accordance with this Lemma, which is used for the constructive arguments. Note, however, that Lemma 1 could actually be obtained for *any* δ if random mechanisms were allowed. Furthermore, if the environment is rich enough so that we can find for each agent an outcome that gives him exactly his target expected payoff (e.g. when transfers are available), Lemma 1, and its construction involving serial dictatorships, can be entirely dispensed with. In such cases, therefore, many of our implementation results below can be strengthened accordingly.

In the rest of the paper, we invoke conditions ω and υ, and fix $\tilde{a}^i \in A$ and $C^i \subseteq A$ for any agent i in accordance with Lemma 1.

3 Efficient Implementation with Incentive Compatibility

Our first set of sufficiency results build on constructing the following regime defined for an SCF f that satisfies conditions ω and υ. First, mechanism $b^* = (M^*, \psi^*)$ is defined as follows:

(i) For all i, $M_i^* = \Theta_i \times \mathbb{Z}_+$, where $\mathbb{Z}_+$ refers to the set of non-negative integers.

(ii) For any $m = ((\theta_i, z^i))_{i \in I}$, $\psi^*(m) = f(\theta_1, \ldots, \theta_I)$.

Next, let B^* represent any regime such that $B^*(\emptyset) = b^*$ and, for any history $h = ((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 1$ and $g^{t-1} = b^*$, we have the following transition rules:

1. If $m_i^{t-1} = (\theta_i, 0)$ for all i, then $B^*(h) = b^*$.
2. If there exists some i, or "odd-one-out," 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 $B^*|h = S^i$, where S^i is obtained from Lemma 1 above.
3. If m^{t-1} is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then $B^*|h = D^i\,(C^i)$.

This regime is similar to the regimes constructed for the complete information case in LS. It starts with 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 an agent reports a non-zero integer, this odd-one-out can obtain 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 (over the subset of outcomes C^i) forever as of the next period.

Characterizing the properties of an equilibrium of this regime is made more complicated by the presence of incomplete information, compared to the corresponding task in the complete information setup. For expositional reasons, our characterization arguments in this section as well as below are based on restricting attention to pure strategies. They can be readily extended to incorporate mixed strategies; see the Supplemental Material of LS [22].

We begin by obtaining for a given history a lower bound on each agent's *expected* equilibrium continuation payoff *at the next period*. This contrasts with the corresponding result with complete information in LS, which calculates the lower bound for the actual continuation payoff at any period. With incomplete information, each player i does not know the private information held by others and, therefore, the (off-the-equilibrium) possibility of continuation regime S^i, which guarantees the continuation payoff $v_i(f)$, delivers a lower bound on equilibrium payoff only in expectation.

Lemma 2 *Fix any $\sigma \in Q^\delta(B^*)$. For any t and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$, if $g^{\mathbf{h}}(\sigma, B^*) = b^*$ then $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^*) \geq v_i(f)$ for all i.*

Proof. Suppose not; so, at some $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$, $g^{\mathbf{h}}(\sigma, B^*) = b^*$ but $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^*) < v_i(f)$ for some i. Then, consider i deviating to another strategy σ_i' identical to the equilibrium strategy σ_i at every history, except that after observing $\mathbf{h}_i$, for each current period realization of θ_i, it reports the same type as 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 others' types. As of the next period, it results in either i becoming a dictator forever (transition rule 3 of B^*) or continuation regime S^i (transition rule 2). Since $v_i^i(C^i) \geq v_i(f)$ and i can obtain continuation payoff $v_i(f)$ from S^i, the deviation is profitable, implying a contradiction. ■

Next, we invoke efficiency in the range of the SCF to pin down equilibrium continuation payoffs. Note here that, since the players' actions depend on their past private information as well as public histories, a player's expected continuation payoff at any given history depends on his posterior beliefs on the other players' past types.[6] However, by the definition of B^* and by conditions ω and υ, we know that the expected payoffs evaluated *at the beginning* of the repeated game must belong to the set *comp*(*co*($V(f)$)). Using this fact, we obtain the following lemma.

Lemma 3 *Suppose that f is efficient in the range. Fix any $\sigma \in Q^\delta(B^*)$ and any date t. Also, suppose that $g^{\mathbf{h}}(\sigma, B^*) = b^*$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$. Then, $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^*) = v_i(f)$ for any i and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$.*

Proof. Since $g^{\mathbf{h}} = b^*$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$, it immediately follows from Lemma 2 and Bayes rule that, for all i,

$$E\pi_i^{t+1} = \sum_{\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)} \Pr(\mathbf{h})\, E_{\mathbf{h}}\pi_i^{t+1} \geq v_i(f). \tag{1}$$

[6] This issue of posterior beliefs would not arise if we were to focus on public strategies.

Note that by conditions ω and υ (part (a)), and by the definition of B^*, $\left(E\pi_i^{t+1}\right)_{i\in I} \in comp(co(V(f)))$. Therefore, since f is efficient in the range, (1) implies that $E\pi_i^{t+1} = v_i(f)$ for all i. By Lemma 2, this in turn implies that $E_{\mathbf{h}}\pi_i^{t+1} = v_i(f)$ for any i and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$. $\blacksquare$

It 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 next lemma has been established.

Lemma 4 *Suppose that f is efficient in the range. Fix any $\sigma \in Q^\delta(B^*)$. Also, fix any t, and suppose that $g^{\mathbf{h}}(\sigma, B^*) = b^*$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$. Then, every agent always announces zero at any $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$.*

Proof. Suppose not; so, for some t, $g^{\mathbf{h}}(\sigma, B^*) = b^*$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$, but at some $\mathbf{h}' = (h', \theta'(t)) \in \mathbf{H}^t(\sigma, B^*)$ there exist $i \in I$, $\theta_i^t \in \Theta_i$, and $m_i \in M_i^{R(h')}$ such that i reports $m_i = (\cdot, z)$, $z \neq 0$, upon observing $\mathbf{h}'_i = (h', \theta'_i(t))$ and θ_i^t. By part (b) of condition υ, there must exist some $j \neq i$ such that $v_j^j(C^j) > v_j(f)$. Consider j deviating to another strategy identical to the equilibrium strategy, σ_j, except that, after observing $\mathbf{h}'_j = (h', \theta'_j(t))$ and any θ_j, it reports the same type as σ_j but a different integer higher than any integer that can be reported by σ_{-j} after any full history belonging to the set $\{\mathbf{h} = (h, \theta(t)) \in \mathbf{H}^t(\sigma, B^*) : h = h'\}$.

By the definition of b^*, the deviation does not alter the current outcome, regardless of the others' types. But, the continuation regime is $D^j(C^j)$ if i's realized type is θ_i^t while, otherwise, it is $D^j(C^j)$ or S^j. In the former case, j can obtain continuation payoff $v_j^j(C^j) > v_j(f)$; in the latter, he can obtain at least $v_j(f)$. Since, by Lemma 3, the equilibrium continuation payoff is $v_j(f)$, the deviation is thus profitable, implying a contradiction. $\blacksquare$

This leads to the following.

Lemma 5 *If f is efficient in the range, every $\sigma \in Q^\delta(B^*)$ is such that $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^*) = v_i(f)$ for any i, t and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$.*

Proof. Since $B^*(\emptyset) = b^*$, it follows by induction from Lemma 4 that, for any t and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$, $g^{\mathbf{h}} = b^*$. Lemma 3 then completes the proof. $\blacksquare$

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.

Incentive compatibility: An SCF f is *incentive compatible* if, for any i and θ_i, $v_i(f|\theta_i) \geq \sum_{\theta_{-i}\in\Theta_{-i}} p_i(\theta_{-i}|\theta_i) u_i(f(\theta_{-i},\theta_i'),(\theta_{-i},\theta_i))$ for all $\theta_i' \in \Theta_i$.

It is straightforward to see that, if f is incentive compatible, B^* admits a Markov Bayesian Nash equilibrium in which each agent i always reports his true type and zero integer. To see this, note first that i's equilibrium continuation payoff is equal to $v_i(f)$ at every history. Thus, reporting a non-zero integer at any history yields the same continuation payoff since it does not affect the current outcome and only activates the continuation regime S^i. Neither is it profitable to deviate by reporting a non-truthful type due to incentive compatibility and by the Markov assumption. Together with Lemmas 2-5, this allows us to obtain our first Theorem.

Theorem 1 *Fix any $I \geq 2$. If f is efficient in the range, incentive compatible and satisfies conditions ω and υ, f is payoff-repeatedly implementable in Bayesian Nash equilibrium from period 2.*

With strong efficiency in the range, we therefore obtain a stronger implementation result in terms of outcomes.

Corollary 1 *In addition to the conditions in Theorem 1, suppose that f satisfies strong efficiency in the range. Then, f is repeated-implementable in Bayesian Nash equilibrium from period 2.*

Proof. Fix any $\sigma \in Q^\delta(B^*)$, $i \in I$, $t > 1$, and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^*)$. Then, by Lemmas 4 and 5, we have

$$\begin{aligned} E_{\mathbf{h}}\pi_i^{t+1} &= \sum_{\theta^t\in\Theta} p(\theta^t) \sum_{\theta^{t+1}\in\Theta} p\left(\theta^{t+1}\right) \\ & \left[(1-\delta)u_i\left(\psi^*(m^{t+1}),\theta^{t+1}\right) + \delta E_{\mathbf{h},\theta^t,m^t,\theta^{t+1},m^{t+1}}\pi_i^{t+2}\right], \end{aligned} \tag{2}$$

where $m^t = \sigma(\mathbf{h},\theta^t)$ and $m^{t+1} = \sigma((\mathbf{h},\theta^t,m^t),\theta^{t+1})$.

But, by Lemma 5, we have

$$E_{\mathbf{h}}\pi_i^{t+1} = v_i(f),$$

and, for any $\theta^t, \theta^{t+1} \in \Theta$, and for any $m^t, m^{t+1} \in M^*$ that occur in equilibrium,

$$E_{\mathbf{h},\theta^t,m^t,\theta^{t+1},m^{t+1}} \pi_i^{t+2} = v_i(f).$$

Thus, (2) implies

$$v_i(f) = \sum_{\theta^t \in \Theta} p(\theta^t) \sum_{\theta^{t+1} \in \Theta} p\left(\theta^{t+1}\right) u_i\left(\psi^*(m^{t+1}), \theta^{t+1}\right).$$

By strong efficiency in the range this implies that $\psi^*(m^{t+1}) = f(\theta^{t+1})$ for all $\theta^{t+1} \in \Theta$ and for all $m^{t+1} \in M^*$ that occur in equilibrium. ■

4 Efficient Implementation without Incentive Compatibility

4.1 Interdependent Values and Payoff-Identifiability

Theorem 1 and its Corollary establish Bayesian repeated implementation of an efficient SCF without (Bayesian) monotonicity but they still assume incentive compatibility to ensure existence of an 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.

Many authors have identified a conflict between efficiency and incentive compatibility in the one-shot setup with interdependent values in which some agents' utilities depend on others' private information (e.g. Maskin [25] and Jehiel and Moldovanu [18]). Therefore, this case presents a particularly important challenge for understanding what is needed for implementation in dynamic settings. In this section, we obtain sufficiency results for full repeated implementation which do not involve incentive compatibility.

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

Definition 5 *An SCF f is payoff-identifiable if, for any i, $\theta_i, \theta_i' \in \Theta_i$ such that $\theta_i' \neq \theta_i$ and $\theta_{-i} \in \Theta_{-i}$, 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, payoff-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.[7] Thus, with a payoff-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. Payoff-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.

In order to ensure the existence of an equilibrium of any regime, we must have either incentive compatibility so that deviation from truth-telling is not profitable in the given period, or some form of detectability of deviation so that punishment can be built in the continuation game. Although realized utilities offer a natural signal for identifying a deviation, in general, there could also be other types of signals available. If so, we could accordingly modify Definition 5.

To allow the possibility of punishment we strengthen condition ω to allow for the existence of an outcome that is strictly worse than the SCF for every agent, which we sometimes refer to as the "bad outcome" below. Specifically, we assume the following.

Condition ω^*: There exists $\tilde{a} \in A$ such that $v_i(\tilde{a}) < v_i(f)$ for all i.

Consider an SCF f that satisfies conditions ω^* and υ. 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 arbitrary but fixed $a \in A$, and define $\tilde{b}^*$ as the following extensive form mechanism:

- Stage 1 - Each agent i announces his private information θ_i, and $f(\theta_1, \ldots, \theta_I)$ is implemented.
- Stage 2 - Once agents learn their utilities, but before a new state is drawn, 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

[7] Clearly, payoff-identifiability cannot hold with private values.

the current period but they determine the continuation play in the regime below.[8] Next define regime $\widetilde{B}^*$ to be such that $\widetilde{B}^*(\emptyset) = Z$ and satisfies the following transition rules:

1. For any $h = (g^1, m^1) \in H^2$,

   (a) if $m_i^1 = (0)$ for all i, then $\widetilde{B}^*(h) = \tilde{b}^*$;

   (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 $\widetilde{B}^*|h = S^i$ such that i's expected continuation payoff is equal to $v_i(f)$ (by conditions ω^* and υ, and by Lemma 1, regime S^i exists);

   (c) if m^1 is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then $\widetilde{B}^*|h = D^i\,(C^i)$.

2. For any $h = ((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 2$ and $g^{t-1} = \tilde{b}^*$:

   (a) if m_i^{t-1} is such that every agent reports NF and 0 in Stage 2, then $\widetilde{B}^*(h) = \tilde{b}^*$;

   (b) if m_i^{t-1} is such that at least one agent reports F in Stage 2, then $\widetilde{B}^*|h = D^i(\{\tilde{a}\})$;

   (c) if m_i^{t-1} is such that every agent reports NF and every agent except some i announces 0, then $\widetilde{B}^*|h = S^i$, where S^i is as in 1(b) above;

   (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 $\widetilde{B}^*|h = D^i\,(C^i)$.

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, then the next period's mechanism is $\tilde{b}^*$; otherwise, strategic interaction ends with the continuation regime being either S^i or $D^i\,(C^i)$ for some i.

The new mechanism $\tilde{b}^*$ sets up two reporting stages. In the first stage, the agents report their types and the current period's outcome is implemented. In the second stage, 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

[8] A similar mechanism design is considered by Mezzetti [26] in the one-shot context with interdependent values and quasi-linear utilities.

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 outcome $\tilde{a}$ which yields a payoff lower than that of the SCF for every agent.

Several comments are worth pointing out about regime $\widetilde{B}^*$. First, why do we employ a two-stage mechanism $\tilde{b}^*$? 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, while either a flag or a positive integer leads to shutdown of strategic play in the regime, it is important that integer play does not occur before the possibility of learning about a deviation. On the one hand, we want to make sure that a deviation leads to immediate punishment and hence a flag overrides any positive integer; on the other hand, as have already seen, the integer play allows us to impose a lower bound on the equilibrium continuation payoff of each player.

Third, note that the initial mechanism enforces an arbitrary outcome and only integers are reported. The integer play affects transitions such that the agents' continuation payoffs are bounded. We do not however allow for any strategic play towards first period implementation in order to avoid potential incentive or coordination problems; otherwise, one equilibrium may be that all players flag in period 1.

We show that if f is efficient in the range then Bayesian equilibria of regime $\widetilde{B}^*$ exhibit the same payoff properties as those of regime B^* in Section 3 above and always induce expected continuation payoff profile $v(f)$. This characterization result is obtained by applying similar arguments to those leading to Lemma 5 above. The additional complication in deriving the result for $\widetilde{B}^*$ is that we also need to ensure that no agent flags on an equilibrium path. This is achieved inductively by showing that expected continuation payoffs are $v(f)$ and, hence, efficient, whereas flagging induces inefficiency because the continuation payoff for each i after flagging is $v_i(\tilde{a}) < v_i(f)$.

We then establish that, with a payoff-identifiable SCF, $\widetilde{B}^*$ admits a Bayesian Nash equilibrium that attains the desired outcome path with sufficiently patient agents. In this equilibrium, every agent first announces zero integer in period 1; from then on, when playing mechanism $\tilde{b}^*$, he announces his true type in the first stage, and in the second stage, chooses $(NF, 0)$ unless he either made a false report himself in the previous stage or identifies another agent's deviation from the realization of utilities; in the latter cases, the agent chooses to "flag". Deviation is not profitable in the first stage since this would be met by "flag" and hence the permanent enforcement of the bad outcome $\tilde{a}$. In the second stage, deviating to non-zero integer clearly does not improve the continuation payoff (Rule 2(c); the same argument also applies to the behavior in period 1). Also, in the Stage 2 continuation game following a false report (off-the-equilibrium) in Stage 1, 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.[9]

Thus, we obtain our next Theorem below.

Theorem 2 *Consider the case of interdependent values. Suppose that f satisfies efficiency in the range, payoff-identifiability, and conditions ω^* and υ. Then, we have the following: there exist a regime R and $\bar{\bar{\delta}} \in (0,1)$ such that, for any $\delta \in (\bar{\delta}, 1)$ (i) $Q^{\delta}(\widetilde{B}^*)$ is non-empty and (ii) every $\sigma \in Q^{\delta}(R)$ is such that, for every $t \geq 2$, $E\pi_i^t(\sigma, R) = v_i(f)$ for any i.*

Proof. See Appendix A. ■

The above result establishes payoff repeated implementation from period 2 when $\delta \in (\bar{\delta}, 1)$. By the same reasoning as in the proof of Corollary 1 above, Theorem 2 can be extended to show outcome implementation if f satisfies strong efficiency in the range.

4.2 Approximate Repeated Implementation

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 can be possible once utilities are learned; with private values, each agent's utility depends only on his own type and hence this kind of payoff-identifiability cannot hold.

[9]Thus, the equilibrium strategy profile is sequentially rational. Note, however, that the mutual optimality of this behavior is supported by an indifference argument; outcome $\tilde{a}$ is permanently implemented regardless of one's response to another flag. This is in fact a simplification. Since, given conditions ω^* and υ, $v_i(\tilde{a}) < v_i(f) \leq v_i^i(C^i)$ for all i, the following modification to the regime makes it strictly optimal for an agent to flag given that there is another flag (regardless of his beliefs about the others' private information): If a subset of agents $J \subseteq I$ announce a flag, the continuation regime makes each player $j \in J$ the dictator in the first j-th period of the continuation game while implementing $\tilde{a}$ in every other period.

One way to identify a deviation and sustain an equilibrium 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 [24] and the references therein).

Another approach to identifying a deviation is that used in Fudenberg, Levine and Maskin [12] (hereafter FLM) to establish a folk theorem under imperfect public monitoring with complete information. This approach, in contrast with the techniques based on review strategies, deters deviations without any explicit statistical tests, by constructing appropriate continuation payoffs after each observable outcome. Of course, this method also requires that deviations can be statistically distinguished, but it does not need to adhere to a review strategies or appeal to a law of large numbers. The method uses linear programming arguments to construct appropriate continuation payoffs. While FLM's main result is for games of imperfect monitoring with complete information, FLM show that their techniques can also be used to establish a folk theorem result for repeated Bayesian games with private values and independent types.

The construction of an equilibrium in both approaches necessarily entails the possibility of "errors" since the players will sometimes fall into a punishment phase when there in fact has been no deviation. This seems to imply that neither approach can be used to obtain existence of an equilibrium whose payoffs match the desired payoffs exactly, and hence, exact implementation. However, with sufficiently patient players, the payoff consequences of any error (and hence any punishment phase) can be made arbitrarily small. Therefore, the standard (one-shot) incentive-compatibility condition may not be necessary if one wants to achieve partial implementation *approximately* in a repeated setup. In this section, we apply these construction techniques of the repeated game literature to our setup and establish *full* approximate implementation results in general incomplete information environments for the case of independent private values and beyond.

Formally, we propose the following alternative notion of repeated implementation that allow for an arbitrarily small degree of departure from exact implementation.[10]

Definition 6 *An SCF f is ϵ-payoff-repeatedly implementable in Bayesian Nash equilibrium if, for any $\epsilon > 0$, there exists a regime R and $\bar{\delta} \in (0,1)$ such that, for any $\delta \in (\bar{\delta}, 1)$, $Q^\delta(R)$ is non-empty and every $\sigma \in Q^\delta(R)$ is such that, for every $t \geq 1$ and i, $\mid E\pi_i^t(\sigma, R) - v_i(f) \mid < \epsilon$.*[11]

The above definition, as before, allows for any private strategies. In what follows, however, we restrict attention to public strategies in order to simplify analysis. Specifically, we modify the regime R^* in Theorem 1 and construct, for each $\epsilon > 0$, a regime B^ϵ that ϵ-payoff-repeatedly implements f in public Bayesian Nash equilibrium. The modifications to our previous constructions are done so that we can appeal to the techniques developed by FLM.

For each degree of approximation $\epsilon > 0$, the modified regime B^ϵ has two new features: (i) in every period, when agents are called to announce a state, each can ensure implementation of an outcome in that period that makes every agent worse off (in expectation) than the desired SCF and (ii) in contrast to R^* in which each agent could guarantee himself a continuation payoff $v_i(f)$ when he is the odd-one-out, in B^ϵ each agent i can only guarantee himself a continuation payoff $v_i(f) - \eta$, for some small $\eta > 0$ that depends on on the degree of approximation ϵ, when every other agent announces integer 0.

Formally, consider an SCF f that is efficient in the range and satisfies conditions ω^* and υ. Define mechanism $b^{**} = (M, \psi)$ as follows:

(i) For all i, $M_i = Y_i \times \mathbb{Z}_+$, where $Y_i = \Theta_i \cup \{N\}$.

(ii) For any $m = (y_i, z_i)_{i \in I}$, with $y_i \in Y_i$ and $z_i \in \mathbb{Z}_+$, $\psi(m) = f(y_1, \ldots, y_I)$ if $y_i \in \Theta_i$ for all i, and otherwise, $\psi(m) = \tilde{a}$.

Mechanism b^{**} is identical to b^* except that each player can induce the bad outcome $\tilde{a}$ unilaterally by announcing N. This means that one Nash equilibrium of b^{**} is for each player to choose N resulting in the bad outcome $\tilde{a}$.

[10] Jackson and Sonnenschein [17] adopt another definition of (approximate) repeated implementation in a finite horizon problem whereby the equilibrium payoffs are required to approximate the target payoffs just at the beginning of the repeated game. Note that, in our infinitely repeated setup, we seek a regime such that the target payoffs are almost achieved at all its equilibrium histories.

[11] Notice that this notion of approximate implementation is defined just from period 1. As $\delta \to 1$ the first period payoff becomes immaterial.

Next, we fix $\epsilon > 0$ and define regime B^ϵ. For this, note first that since f is efficient in the range there exists $\kappa_i \in (0, 1)$ for each i and $\sum_i \kappa_i = 1$ such that, for any $w = (w_1, \ldots, w_I) \in co(V(f))$,

$$\sum_{i \in I} \kappa_i v_i(f) \geq \sum_{i \in I} \kappa_i w_i.$$

By conditions ω^* and υ, there exists $\tilde{\epsilon} > 0$ such that, for all i, $v_i(\tilde{a}) < v_i(f) - \tilde{\epsilon}$ and, for two distinct agents j and k, $v_j^j(C^j) > v_j(f) + \tilde{\epsilon}$ and $v_k^k(C^k) > v_k(f) + \tilde{\epsilon}$. Also, let $\underline{\kappa} = \min_{i \in I} \{\kappa_i\}$ and $\underline{p} = \min_{\theta \in \Theta} \{p(\theta)\}$. Define $\epsilon^* = \frac{\underline{p} \min\{\epsilon, \tilde{\epsilon}\}}{2}$. Then, fix any

$$\eta \in (0, \underline{\kappa}\epsilon^*) . \tag{3}$$

By an analogous reasoning as in Lemma 1, for any $\delta > 1/2$, by condition ω^*, for each i, there exists a history-independent regime in which i obtains a payoff exactly equal to $v_i(f) - \eta$. Such a regime, which we shall refer to as $S^i(\eta)$, alternates between i-dictatorship $d^i(C^i)$ and the bad outcome $\tilde{a}$ as in condition ω^*.

Regime B^ϵ is then defined inductively as follows. First, mechanism b^{**} is played in $t = 1$. Second, if, at some date $t \geq 1$, b^{**} is the mechanism played and $m_i \in Y_i \times \mathbb{Z}_+$ is the message announced by agent i, the continuation mechanism or regime at the next period is determined by the transition rules below:

- Rule A: If $m_i = (\theta_i, 0)$ for all i, then the mechanism next period is b^{**}.
- Rule B: If there exists some i such that $m_j = (\theta_j, 0)$ for all $j \neq i$ and $m_i^{t-1} = (\theta_i, z_i)$ with $z_i \neq 0$, then the continuation regime is $S^i(\eta)$.
- Rule C: If m is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then the continuation regime is $D^i(C^i)$.

By a similar argument as in the characterization part of Theorem 1, the agents' expected continuation payoffs in any equilibrium of regime B^ϵ from period 2 is shown to be within ϵ of the target profile $v(f)$, if $\delta > 1/2$. This is ensured by our precise selection of the quantifier $\eta < \epsilon$ in the continuation regimes $\{S^i(\eta)\}$ as in (3). Since any one period payoff is insignificant when δ is near 1, it follows that the agents' expected continuation payoffs from period 1 are also within ϵ of the target profile $v(f)$, if the agents are sufficiently patient (see Appendix B.1).

Next, we show that the above regime admits an equilibrium (in public strategies) by applying the construction arguments of FLM. Our arguments are as follows. First, we consider an auxiliary stationary regime in which the agents repeatedly play a mechanism that is identical to b^{**} except that they are not allowed to report an integer. This mechanism is then just a repeated Bayesian game. By adapting FLM's techniques, we show that such a repeated Bayesian game admits an equilibrium such that all its continuation payoffs are arbitrarily close to $v(f)$ for δ sufficiently close to 1 under the required conditions (see Appendix B.2). Then, using the above equilibrium profile, we construct an equilibrium of our regime B^{ϵ} in which the players' behavior are identical to their equilibrium behavior in the auxiliary regime except that they also announce 0 at every history (see Appendix B.3).

Formally, define $\tilde{b}^{**}$ as a mechanism that is identical to b^{**} above except that the agents do not report an integer, i.e. the message set for each i is $Y_i = \Theta_i \cup \{N\}$. Each agent i's pure strategy in $\tilde{b}^{**}$ is a mapping $s_i : \Theta_i \to Y_i$, with S_i denoting the set of all such strategies and $S \equiv S_1 \times \cdots \times S_I$. The outcome function is such that if every player i chooses to report $\theta_i \in \Theta_i$ then $f(\theta_1, \ldots, \theta_I)$ is implemented; otherwise, the mechanism implements the bad outcome $\tilde{a}$. Note that one Nash equilibrium of this mechanism is for every player to always announce N, which induces expected payoff profile $v(\tilde{a})$. Let us denote by $\mathcal{B}$ the stationary regime which repeatedly enforces $\tilde{b}^{**}$. This is just a repeated Bayesian game.

Our equilibrium construction for regime $\mathcal{B}$ appeals to the threat of a Nash equilibrium that always induces the bad outcome. Thus, we first assume the following full-dimensionality condition.

Full dimensionality: The set $V^*(f) = \{v \in co(V(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$ has non-empty interior.

The crucial aspect of FLM's equilibrium construction is the players' ability to statistically identify deviations. One way to capture this is the condition of *pairwise identifiability*. Loosely speaking, a strategy profile is pairwise identifiable if, for every pair of players, the distributions over reports induced by one player's unilateral deviations are distinct from those induced by the other's deviations; that is, distributions over announcements induced by deviations by any player cannot be replicated by any (possibly random) reporting strategy of any other player.

To formally introduce this condition in our setup, let $\overline{s}_i : \Theta_i \to \Theta_i$ refer to a *revelation* strategy, and let the set of such strategies be denoted by $\overline{S}_i \subset S_i$. Let $n_i \equiv\mid \Theta_i \mid$ and $m_i \equiv\mid \overline{S}_i \mid= n_i^{n_i}$. We write $\overline{S}_i = \{\overline{s}_{i,1}, \ldots, \overline{s}_{i,m_i}\}$. Also, let $n \equiv \prod_{i=1} n_i$ denote the number of all possible type profiles; thus, $\Theta = \{\theta^1, \ldots, \theta^n\}$.

Fix any i and revelation strategy profile $\overline{s}_{-i} \in \overline{S}_{-i}$. Let $\Gamma_i(\overline{s}_{-i})$ refer to a $m_i \times n$ matrix whose (k,l)-th element corresponds to the probability of profile θ^l being reported when i plays his k-th revelation strategy $\overline{s}_{i,k}$ and the other players use $\overline{s}_{-i}$ defined by

$$p(\theta^l \mid \overline{s}_{i,k}, \overline{s}_{-i}) = \sum_{\{\theta' : \overline{s}_{i,k}(\theta'_i)=\theta^l_i, \overline{s}_{-i}(\theta'_{-i})=\theta^l_{-i}\}} p(\theta').$$

Also, for all $\overline{s}$ and all $i, j \in I$, let $\Gamma_{ij}(\overline{s})$ be a $(m_i + m_j) \times n$ matrix formed by stacking matrices $\Gamma_i(\overline{s}_{-i})$ and $\Gamma_j(\overline{s}_{-j})$; thus,

$$\Gamma_{ij}(\overline{s}) = \begin{pmatrix} \Gamma_i(\overline{s}_{-i}) \\ \Gamma_j(\overline{s}_{-j}) \end{pmatrix}.$$

We say that a revelation strategy profile $\overline{s}$ is pairwise identifiable if, for all $i, j \in I$,

$$rank(\Gamma_{ij}(\overline{s})) = rank(\Gamma_i(\overline{s}_{-i})) + rank(\Gamma_j(\overline{s}_{-j})) - 1.$$ [12]

Note that pairwise identifiability is a property of the information structure (and not a particular SCF). We say that a revelation strategy profile $\overline{s}$ induces a payoff profile $w \in V(f)$ if for all i

$$\sum_{\theta} p(\theta) u_i(f(\overline{s}_i(\theta_i), \overline{s}_{-i}(\theta_{-i})), \theta) = w_i.$$

Definition 7 *The information structure is pairwise identifiable with respect to $X \subseteq V(f)$ if every revelation strategy profile $\overline{s}$ that induces an efficient payoff profile in X is pairwise identifiable for every pair of players.*[13]

[12] Pairwise identifiability requires that the stacked matrix $\Gamma_{ij}(\overline{s})$ have the largest rank possible given the ranks of the constituent matrices $\Gamma_i(\overline{s}_{-i}))$ and $\Gamma_j(\overline{s}_{-j})$. This is because the row of $\Gamma_i(\overline{s}_{-i})$ corresponding to player i choosing $\overline{s}_i$ must be the same as that of $\Gamma_j(\overline{s}_{-j})$ corresponding to j choosing $\overline{s}_j$.

[13] If $\mid \Theta_i \mid \leq \prod_{k \neq i,j} \mid \Theta_k \mid$ for every i and j, for any generic probability distribution on Θ, the truthful reporting profile is pairwise identifiable for every pair of players (see Fudenberg, Levine and Maskin [13]). Therefore, if the players have the same number of possible types, then for any generic probability distribution on Θ pairwise identifiability is satisfied with three or more players. More generally, this conclusion holds if the number of players is sufficiently large compared to the variation in the number of types across players.

Recall that pairwise identifiability ensures that the two players' deviations can be distinguished statistically. Another condition that plays a similar role in one-shot mechanism design problems is the *weak identifiability* condition introduced by Kosenok and Severinov [19].[14] We conjecture that our results will continue to hold if pairwise identifiability is replaced by weak identifiability.[15]

Our first approximate implementation result is obtained for the case of private values.

Theorem 3 *Consider the case of private values, and fix an SCF f that satisfies efficiency in the range and conditions ω^* and υ. Suppose also that the information structure is pairwise identifiable with respect to $V(f)$ and that $V^*(f) = \{v \in co(V(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$ has non-empty interior. Then, f is ϵ-payoff-repeatedly implementable in public Bayesian Nash equilibrium.*

Proof. See Appendix B. ■

As noted by FLM, when types are independently distributed in the private value case, pairwise identifiability of efficient revelation profiles is vacuously satisfied. Thus, we immediately obtain the following corollary.

Corollary 2 *Consider the case of independent private values. Then, if an SCF f satisfies efficiency in the range and conditions ω^* and υ, and if $V^*(f)$ has non-empty interior, f is approximate-payoff-repeatedly implementable.*

Before turning to the case of interdependent values, let us offer a brief sketch of how FLM's techniques are adapted to our implementation setup for equilibrium construction. Note that in our regime B^ϵ the players' equilibrium continuation payoffs must always be bounded below by $v(f) - \epsilon$. Therefore, in order to construct an equilibrium, we have to ensure that the corresponding payoff profile is sustained by continuation payoffs that all lie arbitrarily close to the target profile $v(f)$. FLM show that this is possible if a smooth set W of feasible payoffs is *decomposable on tangent hyperplanes*; that is, for every boundary

[14]Both pairwise identifiability and weak identifiability in one-shot mechanism design serve to ensure that any enforceable mechanism is balanced. See Fudenberg, Maskin and Levine [13] and Kosenok and Severinov [19].

[15]Hörner, Takahashi and Vieille [15] generalize the methods of FLM as well as those of Fudenberg and Levine [11] to obtain a folk theorem result for general repeated Bayesian games using the weak identifiability condition.

point of W, there exists a stage game strategy profile inducing a payoff vector outside W that is *enforceable* by continuation payoffs on the tangent hyperplane.[16] This is then ensured by enforceability and pairwise identifiability of the revelation strategies that yield payoffs on the efficient frontier (together with the Nash equilibrium inducing $v(\tilde{a})$).

FLM note that, with private values, any efficient revelation profile is enforceable and, therefore, the pairwise identifiability condition suffices to deliver a public perfect equilibrium. With interdependent values, this is no longer the case. We show that, when types are correlated, one-to-one revelation profiles are enforceable if the information structure satisfies a condition due to Crémer and McLean [7].

Definition 8 *The information structure satisfies condition CM if, for each i, any θ_i and any $\mu_i : \Theta_i \to \mathbb{R}_+$, we have*

$$p_{-i}(\theta_{-i} \mid \theta_i) \neq \sum_{\theta_i' \neq \theta_i} \mu_i(\theta_i') p_{-i}(\theta_{-i} \mid \theta_i').$$

The above condition rules out the possibility that player i of type θ_i could generate the same conditional probabilities on the types of the other players through a random untruthful reporting strategy. It clearly holds for any generic probability distribution on Θ.

In mechanism design, condition CM is invoked to ensure enforceability (via appropriately designed transfers) of *truth-telling*. We extend these arguments to show that, in our setup, condition CM allows for enforceability of revelation strategies that are one-to-one.[17] Thus, to extend the FLM construction with interdependent values, we require the pairwise identifiability condition with respect to

$$\widetilde{V}(f) = \{v : v = \sum_{\theta} u(f(\lambda_1(\theta_1), \ldots, \lambda_I(\theta_I)), \theta) p(\theta) \text{ for some 1-to-1 functions } \lambda_i : \Theta_i \to \Theta_i\},$$

and full dimensionality with respect to

$$\widetilde{V}^*(f) = \{v \in co(\widetilde{V}(f)) : v > v(\tilde{a})\} \cap v(\tilde{a}).$$

[16] The notion of enforceability is standard: a stage game strategy profile is enforceable with respect to some set of feasible payoffs if one can find signal-contingent continuation payoffs in the set such that the current strategies are mutual best responses given the expected continuation payoffs.

[17] In mechanism design, the focus is to find enforceability of truth-telling, which is of course one-to-one. In our setup, though we only need to construct one equilibrium, we have to ensure that every continuation payoff profile is itself appropriately enforceable. The latter is why the mechanism design arguments are extended to accommodate all one-to-one revelation strategies here.

We present our final result below.

Theorem 4 *Consider the case of interdependent values, and fix an SCF f that satisfies efficiency in the range and conditions ω^* and υ. Suppose also that the information structure satisfies condition CM and pairwise identifiability with respect to $\widetilde{V}(f)$, and that $\widetilde{V}^*(f)$ has non-empty interior. Then, f is ϵ-payoff-repeatedly implementable in public Bayesian Nash equilibrium.*

Proof. See Appendix B. ■

5 Concluding Discussion

This paper sets up the general problem of repeated implementation with incomplete information and demonstrates the extent to which social choices that satisfy efficiency in the range can be repeatedly implemented in (Bayesian) Nash equilibrium. As similarly established by LS for the special case of complete information, a fundamental difference exists between one-shot implementation and repeated implementation with incomplete information. Our sufficiency results demonstrate that neither (Bayesian) monotonicity nor incentive compatibility is necessary for repeated implementation. Instead, they suggest that efficiency remains the critical property of social objectives that enable successful implementation in dynamic contexts.

Let us conclude with some remarks on our results.

Necessity With complete information, LS have shown that, with minor qualifications, efficiency in the range is necessary for repeated implementation. We are yet to derive a corresponding result with incomplete information. While intuition would suggest that whenever there is an opportunity to make every player strictly better off the possibility to coordinate on past histories would enable infinitely-lived players to arrive at an alternative collusive equilibrium, with incomplete information, this is a difficult proposition to establish. The main reasons are two-fold. First, a deviation from collusive strategies needs to be detected for punishment but incomplete information naturally imposes barriers to build such threats. Second, the regime may possess complex transitional structures off-the-equilibrium which make any collusion difficult to sustain. Nonetheless, we conjecture that efficiency in the range is a necessary condition for repeated implementation in terms of outcomes at least for the case of independent private values.

Solution concept and construction method The solution concept invoked 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 [4], Baliga [1] and Brusco [5][6].[18]

Our construction involves integer games which have sometimes drawn criticism for their unbounded nature. They can be replaced with finite constructions, such as the modulo game, but then unwanted equilibria emerge in mixed strategies. In a recent paper, Lee and Sabourian [23] attempt to address this issue in the context of repeated implementation with complete information. They show that it is indeed possible to obtain efficient repeated implementation using only simple finite mechanisms if one invokes a mild refinement argument. Similar approaches may also work with incomplete information.

Informational requirement 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 [2][3]). To see this, recall Lemma 2 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.

We can therefore extend our incomplete information analysis by adopting ex post equilibrium as solution concept, thereby requiring the agents' strategies to be mutually optimal for every possible realization of past and present types and not just for some given distribution. As the precise nature of the agents' private information are not relevant in our constructive arguments, in fact, the more stringent restrictions imposed by ex post equilibrium would enable us to derive a sharper set of results that are closer to the results with complete information; see our previous working paper [20].

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

[18]Nonetheless, we show that our regimes admit equilibria satisfying sequential rationality.

but not necessarily the marginal distributions; and (ii) how such a payoff can be generated 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.

Markovian preferences In this paper, as well as in LS, we consider preferences that follow an i.i.d. process. An extension would be to consider a Markov process. Recently, folk theorem results have been extended to repeated Bayesian games where information follows a Markov process (e.g. Escobar and Toikka [10] and Hörner, Takahashi and Vieille [15]). Given these results, we conjecture that our construction methods and implementation results can be enriched to such Markovian setups.

A Proof of Theorem 2

This Theorem is proved by the following two lemmas.

Lemma 6 *If f is efficient in the range and satisfies conditions ω^* and υ, then every $\sigma \in Q^\delta(\widetilde{B}^*)$ is such that $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, \widetilde{B}^*) = v_i(f)$ for any i, t and $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$.*

Proof. Fix any $\sigma \in Q^\delta(\widetilde{B}^*)$. We proceed with the following claims.

Claim 1: Fix any $t > 1$. Assume that, for any $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$, $g^{\mathbf{h}} = \tilde{b}^*$ and also that, after any realization of θ^t at this date, every agent reports "no flag" in Stage 2. Then, for any i, $E_{\mathbf{h}}\pi_i^{t+1} = v_i(f)$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$ and hence $E\pi_i^{t+1} = v_i(f)$.

Since, in the above claim, it is assumed that every agent reports "no flag" in Stage 2 after any θ^t, the claim can be proved analogously to Lemmas 2 and 3 above.

Claim 2: Fix any $t > 1$. Assume that, for any $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$, $g^{\mathbf{h}} = \tilde{b}^*$ and also that, after any realization of θ^t at this date, every agent reports "no flag" in Stage 2. Then, for any $\mathbf{h}' \in \mathbf{H}^{t+1}(\sigma, \widetilde{B}^*)$, the following two properties hold: (i) $g^{\mathbf{h}'} = \tilde{b}^*$ and (ii) every agent will always report "no flag" in period $t+1$.

Again, since in the above claim it is assumed that every agent reports "no flag" in Stage 2 at any θ^t, part (i) of the claim can be established analogously to Lemma 4 above.

To prove part (ii), suppose otherwise; then some agent reports "flag" in Stage 2 of period $t+1$ after some sequence of states $\hat{\theta}^1, \ldots, \hat{\theta}^{t+1}$. (Recall that we assume pure strategies to simplify the arguments.)

For any s and $\theta(s) \in \Theta^{s-1}$ define $f^{\theta(s)}$ as the SCF such that, for each θ, $f^{\theta(s)}(\theta)$ is the outcome implemented in equilibrium at $(\theta(s), \theta)$. Then $E\pi_i^{t+1}$ can be written as

$$E\pi_i^{t+1} = (1-\delta)\left[\sum_{\theta(t+1)} q(\theta(t+1))v_i(f^{\theta(t+1)}) + \delta \sum_{s \geq t+2} \sum_{\theta(s)} \delta^{s-t-2} q(\theta(s))v_i(f^{\theta(s)})\right]. \quad (4)$$

Next, for any $s > t+1$, let $\widehat{\Theta}^{s-1} = \left\{\theta(s) = (\theta^1, \ldots, \theta^{s-1}) : \theta^\tau = \hat{\theta}^\tau \forall \tau \leq t+1\right\}$ be the set of $(s-1)$ states that are consistent with $\hat{\theta}^1, \ldots, \hat{\theta}^{t+1}$. Since there is flagging in period $t+1$ after the sequence of states $\hat{\theta}^1, \ldots, \hat{\theta}^{t+1}$ it follows that, for all i,

$$(1-\delta) \sum_{s \geq t+2} \sum_{\theta(s) \in \widehat{\Theta}^{s-1}} \delta^{s-t-2} q(\theta(s))v_i(f^{\theta(s)}) = \alpha v_i(\tilde{a}), \quad (5)$$

where $\alpha = (1-\delta)\sum_{s \geq t+2} \sum_{\theta(s) \in \widehat{\Theta}^{s-1}} \delta^{s-t-2} q(\theta(s))$. Also, by the previous claim, $E\pi_i^{t+1} = v_i(f)$ for all i. Therefore, it follows from (4), (5) and $v(\tilde{a}) \ll v(f)$ that, for all i,

$$v_i(f) < \frac{(1-\delta)}{(1-\delta\alpha)}\left\{\sum_{\theta(t+1) \in \Theta^t} q(\theta(t+1))v_i\left(f^{\theta(t+1)}\right) + \delta \sum_{s \geq t+2} \sum_{\theta(s) \notin \widehat{\Theta}^{s-1}} \delta^{s-t-2} q(\theta(s))v_i\left(f^{\theta(s)}\right)\right\}. \quad (6)$$

Since $\sum_{\theta(t+1)} q(\theta(t+1)) = 1$ and $\alpha + (1-\delta)\sum_{s \geq t+2} \sum_{\theta(s) \notin \widehat{\Theta}^{s-1}} \delta^{s-t-2} q(\theta(s)) = 1$ by definition, it follows that

$$\sum_{\theta(t+1)} q(\theta(t+1)) + \delta \sum_{s \geq t+2} \sum_{\theta(s) \notin \widehat{\Theta}^{s}} \delta^{s-t-2} q(\theta(s)) = \frac{(1-\delta\alpha)}{(1-\delta)}.$$

Therefore, by (6), f is not efficient in the range; but this is a contradiction.

Claim 3: (i) $E\pi_i^2 = v_i(f)$ for all i; (ii) $g^{\mathbf{h}} = \tilde{b}^*$ for any $\mathbf{h} \in \mathbf{H}^2(\sigma, \widetilde{B}^*)$; and (iii) every agent will report "no flag" in Stage 2 of period 2 after any $\mathbf{h}^2 \in \mathbf{H}(\sigma, \widetilde{B}^*)$ and θ^2.

Since $\widetilde{B}^*(\emptyset) = Z$, we can establish (i) and (ii) by applying similar arguments as those in Lemmas 2-4 to period 1. Also, by similar reasoning for Claim 2 just above, it must be that no player flags in period 2 at any $\mathbf{h}^2 \in \mathbf{H}(\sigma, \widetilde{B}^*)$ and θ^2; otherwise the continuation payoff profile will be $v(\tilde{a})$ from period 3 and this is inconsistent with $E\pi_i^2 = v_i(f) > v_i(\tilde{a})$ for all i.

Finally, to complete the proof of Lemma 6, note that, by induction, it follows from Claims 1-3 that, for any t and $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$, $g^{\mathbf{h}} = \tilde{b}^*$ and, moreover, every agent will always report "no flag"in Stage 2 of period t. But then, Claims 1 and 3 imply that $E_{\mathbf{h}}\pi_i^{t+1} = v_i(f)$ for all i, t and $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$. $\blacksquare$

Lemma 7 *Suppose that f is payoff-identifiable and satisfies conditions ω^* and υ. If δ is sufficiently large, there exists $\sigma^* \in Q^\delta(\widetilde{B}^*)$ such that, for any $t > 1$, $\mathbf{h} \in \mathbf{H}^t(\sigma, \widetilde{B}^*)$ and θ^t, (i) $g^{\theta(t)}(\sigma^*, \widetilde{B}^*) = \tilde{b}^*$; and (ii) $A^{\mathbf{h},\theta^t}(\sigma^*, \widetilde{B}^*) = \{f(\theta^t)\}$.*

Proof. By condition ω^* there exists some $\epsilon > 0$ such that, for each i, $v_i(\tilde{a}) < v_i(f) - \epsilon$. Define

$$\rho \equiv \max_{i,\theta,a,a'} [u_i(a,\theta) - u_i(a',\theta)] \tag{7}$$

and $\bar{\delta} \equiv \frac{\rho}{\rho+\epsilon}$. Fix any $\delta \in \left(\bar{\delta}, 1\right)$. Consider the following symmetric strategy profile $\sigma^* \in \Sigma$: for each i, σ_i^* is such that:

- for any θ_i, $\sigma_i^*(\emptyset, Z, \theta_i) = 0$;
- for any $t > 1$ and 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 1(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 on the equilibrium path. Deviation can take place in two stages:

(i) Stage 1 - Announce a false type. But then, due to payoff-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 $(1-\delta)\max_{a,\theta} u_i(a,\theta)+\delta(v_i(f)-\epsilon)$, while the equilibrium payoff is at least $(1-\delta)\min_{a,\theta} u_i(a,\theta)+\delta v_i(f)$. Since $\delta > \bar{\delta}$, the latter exceeds the former, and the deviation is not profitable.

(ii) Stage 2 - Flag or announce a non-zero integer following a stage 1 at which a no agent provides a false report. But given transition rules 2(b) and 2(c), such deviations cannot make i better off than no "flag" and zero integer. ■

B Proofs of Theorem 3 and Theorem 4

Both theorems are proved by considering regime B^{ϵ} defined in Section 4.2. We offer a brief outline of our arguments. First, we characterize the equilibrium payoff properties of this regime for both private and interdependent value cases (Lemma 8). Second, we consider the auxiliary stationary regime $\mathcal{B}$ for which the techniques of FLM are applied to obtain existence. Here, the required conditions depend on whether values are private or interdependent. Finally, using the equilibrium constructions from the auxiliary regime, we show existence of an equilibrium in our regime B^{ϵ}.

B.1 Characterization of Equilibrium Payoffs

Lemma 8 *If f is efficient in the range and satisfies conditions ω^* and υ, every $\sigma \in Q^{\delta}(B^{\epsilon})$ is such that, for any t and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^{\epsilon})$, we have the following:*

(a) $g^{\mathbf{h}}(\sigma, B^{\epsilon}) = b^{**}$.

(b) $\mid E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^{\epsilon}) - v_i(f) \mid < \epsilon^* \in (0,\epsilon)$ for any i.

Proof. Fix any $\sigma \in Q^{\delta}(B^{\epsilon})$. We proceed in the following steps.

Step 1: For any t and $\mathbf{h} \in \mathbf{H}^t(\sigma, B^{\epsilon})$, if $g^{\mathbf{h}}(\sigma, B^{\epsilon}) = b^{**}$, $E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^{\epsilon}) \geq v_i(f) - \eta$ for all i.

Proof of Step 1. This follows from arguments similar to those for Lemma 2 above.

Step 2: Fix any t, and suppose that $g^{\mathbf{h}}(\sigma, B^{\epsilon}) = b^{**}$ for all $\mathbf{h} \in \mathbf{H}^t(\sigma, B^{\epsilon})$. Then, $\mid E_{\mathbf{h}}\pi_i^{t+1}(\sigma, B^{\epsilon}) - v_i(f) \mid < \epsilon^*$ for any i and $\theta(t)$.

Proof of Step 2. To show $E_{\mathbf{h}}\pi_i^{t+1} < v_i(f) + \epsilon^*$ for all i, suppose otherwise. So, $E_{\mathbf{h}}\pi_j^{t+1} \geq v_j(f) + \epsilon^*$ for some j. Then, by Step 1 above, and since we consider public strategies, we have

$$\begin{aligned}\sum_{i\in I}\kappa_i E_{\mathbf{h}}\pi_i^{t+1} &\geq \kappa_j(v_j(f)+\epsilon^*)+\sum_{i\neq j}\kappa_i(v_i(f)-\eta)\\ &= \sum_{i\in I}\kappa_i v_i(f)+\kappa_j\epsilon^*-\sum_{i\neq j}\kappa_i\eta\\ &> \sum_{i\in I}\kappa_i v_i(f), \end{aligned}\tag{8}$$

where the last inequality follows from the definition of η in (3) above. But, since $\left(E_{\mathbf{h}}\pi_i^{t+1}\right)_{i\in I}\in comp(co(V(f)))$, (8) contradicts that f is efficient in the range.

Step 3: Fix any t, and suppose that $g^{\mathbf{h}}(\sigma,B^\epsilon)=b^{**}$ for all $\mathbf{h}\in\mathbf{H}^t(\sigma,B^\epsilon)$. Then, $g^{\mathbf{h}'}(\sigma,B^\epsilon)=b^{**}$ for any $\mathbf{h}'\in\mathbf{H}^{t+1}(\sigma,B^\epsilon)$.

Proof of Step 3. Suppose not; so, for some t, $g^{\mathbf{h}}(\sigma,B^*)=b^*$ for all $\mathbf{h}\in\mathbf{H}^t(\sigma,B^*)$, but at some $\tilde{\mathbf{h}}=(\tilde{h},\theta(t))\in\mathbf{H}^t(\sigma,B^*)$ there exist $i\in I$, $\theta_i^t\in\Theta_i$, and $m_i\in M_i^{R(\tilde{h})}$ such that i reports $m_i=(\cdot,z)$, $z\neq 0$, upon observing $\tilde{h}$ and θ_i^t. By part (b) of condition v, there must exist some $j\neq i$ such that $v_j^j(C^j)>v_j(f)$. Consider j deviating to another strategy identical to the equilibrium strategy, σ_j, except that, after observing $\tilde{h}$ and any θ_j, it reports the same type as σ_j but a different integer higher than any integer that can be reported by σ_{-j} after the given public history $\tilde{h}$.

By the definition of b^*, the deviation does not alter the current outcome, regardless of the others' types. But, the continuation regime is $D^j(C^j)$ if i's realized type is θ_i^t while, otherwise, it is $D^j(C^j)$ or S^j. In the former case, j can obtain continuation payoff $v_j^j(C^j)>v_j(f)+\tilde{\epsilon}$; in the latter, he can obtain at least $v_j(f)-\eta$. Thus, the deviation continuation payoff at the next period is at least

$$\underline{p}\left(v_j(f)+\tilde{\epsilon}\right)+(1-\underline{p})\left(v_j(f)-\eta\right).$$

By Step 2, the equilibrium continuation payoff next period is at most $v_j(f)+\epsilon^*$. Thus, the deviation is profitable if

$$\begin{aligned}\underline{p}\tilde{\epsilon}-(1-\underline{p})\eta &> \epsilon^*\\ \underline{p}\tilde{\epsilon}-\epsilon^* &> (1-\underline{p})\eta\\ \underline{p}\tilde{\epsilon}-\epsilon^* &> \eta\\ \epsilon^* &> \eta,\end{aligned}$$

where the last inequality follows from the definition of $\epsilon^* = \frac{\underline{p}\min\{\epsilon,\tilde{\epsilon}\}}{2}$. This is true by (3), implying a contradiction. ■

B.2 Auxiliary Stationary Regime $\mathcal{B}$

We can write each player's stage game payoff function in $\widetilde{b}^{**}$ as follows. First, for any $y \in Y$, define the payoff of i when y occurs in state θ by

$$\widetilde{u}_i(f(y),\theta) \equiv \begin{cases} u_i(f(y),\theta) & \text{if } y_i \in \theta_i \text{ for all } i \\ v_i(\tilde{a}) & \text{otherwise.} \end{cases}$$

Define i's payoff function w_i in $\widetilde{b}^{**}$ by

$$w_i(s_i, s_{-i}, y) = \frac{\sum_{\theta\in s^{-1}(y)} u_i(f(y),\theta)p(\theta)}{\sum_{\theta\in s^{-1}(y)} p(\theta)}.$$

The distribution of announcement of i corresponding to s_i is given by

$$q(y_i \mid s_i) = \begin{cases} \sum_{\theta_i\in s_i^{-1}(y_i)} p(\theta) & \text{if } s_i^{-1}(y_i) \text{ is not empty;} \\ 0 & \text{otherwise,} \end{cases}$$

and the distribution of announcements of all (which can be thought of as stage outcome) corresponding to strategy profile $s \in S$ is given by

$$q(y \mid s) = \begin{cases} \sum_{\theta\in s^{-1}(y)} p(\theta) & \text{if } s^{-1}(y) \text{ is not empty;} \\ 0 & \text{otherwise.} \end{cases}$$

Since the distribution of announcements depend on the strategy profile, i's payoff can also be written only as a function of strategy profile s as follows:

$$\begin{aligned} \overline{w}_i(s) &= \sum_m w_i(s,y)q(y \mid s) \\ &= \sum_y \sum_{\theta\in s^{-1}(y)} \widetilde{u}_i(f(y),\theta)\left[\frac{p(\theta)}{\sum_{\theta\in s^{-1}(m)} p(\theta)}\right] q(m \mid s) \\ &= \sum_y \sum_{\theta\in s^{-1}(y)} \widetilde{u}_i(f(y),\theta)p(\theta). \end{aligned}$$

It is clear from above that with private values $w_i(s_i, s_{-i}, y) = w(s_i, s'_{-i}, y)$ for all s'_{-i} and hence it can be written only as a function of own strategy and the public announcement profile y. Therefore, in this case the repeated Bayesian game defined by regime $\mathcal{B}$ can be thought of as a repeated game with imperfect monitoring and hence the result of KLM applies to this setting.

When $w_i(\cdot,\cdot,\cdot)$ depends on s_{-i}, as will be the case with interdependent values, we can still use KLM techniques to analyze regime $\mathcal{B}$. In particular, consider the set of equilibria of regime $\mathcal{B}$ such that strategies depend only on past public announcements. Then such equilibria correspond to public perfect equilibria in KLM and hence we can define the concept of enforceability, decomposability, (local) self-decomposability and decomposability on a tangent hyperplane as in KLM.

B.2.1 Equilibrium with private values

Recall that $V^*(f) = \{v \in co(V(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$.

Lemma 9 (Extension of KLM) *Consider the case of private values. Suppose that the information structure is pairwise identifiable with respect to f. Then, the following statements hold:*

(a) Any smooth set $W \subset Int(V^*(f))$ is decomposable on tangent hyperplanes.

(b) For any any smooth set $W \subset Int(V^*(f))$ there exists $\overline{\delta} < 1$ such that for any $\delta > \overline{\delta}$, W belongs to the set of payoffs that can be generated by some public perfect equilibrium (PPE) of regime $\mathcal{B}$.

(c) Assume full-dimensionality (i.e. $V^*(f)$ has non-empty interior). Then there exists $\overline{\delta} < 1$ such that for any $\delta > \overline{\delta}$ the regime $\mathcal{B}$ admits a PPE $\tilde{\sigma} \in Q^{\delta}(\mathcal{B})$, such that every continuation payoff w of $\tilde{\sigma}$ satisfies

$$v(f) > w > v(f) - \eta,$$

where η is defined by (3) *above.*

Proof. (a) This follows from the proof of Theorem 6.1 of KLM.

(b) This follows from part (a) and Theorem 4.1 of KLM.

(c) Consider any smooth set $W \subset Int(V^*(f))$ such that, for any $w \in W$, $v(f) > w > v(f) - \eta$, where η is defined by (3) above. (Since $Int(V^*(f))$ is non-empty, $v(f)$ belongs to the efficient frontier of $V^*(f)$ and $v(f) > v(\tilde{a})$, such a smooth set exists.) Then, by part (b) above, there exists $\overline{\delta} < 1$ such that for any $w^* \in W$ and any $\delta > \overline{\delta}$, the repeated game, or stationary regime, $\mathcal{B}$ admits a PPE $\tilde{\sigma} \in Q^{\delta}(\mathcal{B})$ with payoffs w^* such that all equilibrium continuation payoffs belong to the set W. ■

B.2.2 Equilibrium with interdependent values

Recall that $\widetilde{V}^*(f) \equiv \{v \in co(\widetilde{V}(f)) : v > v(\tilde{a})\} \cap v(\tilde{a})$, where

$$\widetilde{V}(f) = \{v : v = \sum_{\theta} u(f(\lambda_1(\theta_1), .., \lambda_I(\theta_I)), \theta)p(\theta) \text{ for some 1-to-1 functions } \lambda_i : \Theta_i \to \Theta_i\}.$$

Lemma 10 (Enforceability with interdependent values) *With correlated types, any strategy profile s such that $s_i : \Theta_i \to \Theta_i$ is 1-to-1 for every i, is enforceable if condition CM holds.*

Proof. For any i and θ_i, let $p_{-i}(\theta_i) = \{p_{-i}(\theta_{-i} \mid \theta_i)\}_{\theta_{-i} \in \Theta_{-i}}$. Fix any i and θ_i', and define

$$A(\theta_i') = \left\{ \sum_{\theta_i \neq \theta_i'} \mu_i(\theta_i) p_{-i}(\theta_i) : \text{ for any } \{\mu_i(\theta_i)\}_{\theta_i \neq \theta_i'} \in \mathbb{R}_+^{n_i - 1} \right\}.$$

That is, $A(\cdot)$ is the positive cone generated in $\mathbb{R}^{\Pi_{j \neq i} n_j}$ induced by the vectors $\{p_{-i}(\theta_i)\}_{\theta_i \neq \theta_i'}$.

Condition CM states that $A(\theta_i') \cap p_{-i}(\theta_i')$ is empty. Therefore, we can apply the separating hyperplane theorem to conclude the existence of $b(\theta_i') = \{b(\theta_{-i} \mid \theta_i')\}_{\theta_{-i}} \in \mathbb{R}^{\Pi_{j \neq i} n_j}$ and $c(\theta_i') \in \mathbb{R}$ such that

$$b(\theta_i') \cdot p_{-i}(\theta_i) = \sum_{\theta_{-i}} b(\theta_{-i} \mid \theta_i') p_{-i}(\theta_{-i} \mid \theta_i) > c(\theta_i') \quad \forall \theta_i \neq \theta_i' \tag{9}$$

and

$$b(\theta_i') \cdot p_{-i}(\theta_i') = \sum_{\theta_{-i}} b(\theta_{-i} \mid \theta_i') p_{-i}(\theta_{-i} \mid \theta_i') = c(\theta_i'). \tag{10}$$

In one shot mechanism design $b_i(\theta_{-i} \mid \theta_i') - c_i(\theta_i')$ is interpreted as a transfer from i if (θ_i', θ_{-i}) is announced. In such games, players are only asked to report their types and then the transfers above are enacted. The above construction then ensures that the players have strict incentives to report their true types as multiplying $(b_i(\theta_{-i} \mid \theta_i') - c_i(\theta_i'))$ by a large enough positive constant K overwhelms any other incentives that might arise from the original incentive incompatible mechanism. To show enforceability we do something similar here.

Fix any 1-to-1 revelation strategy profile $\overline{s} \in \overline{S}$. Let $r_i : \Theta_i \to \Theta_i$ be the inverse of $\overline{s}_i$. Let the continuation payoff if $\tilde{\theta}$ is announced be given by

$$W_i(\tilde{\theta}) = \frac{K}{\delta} \{b(r_{-i}(\tilde{\theta}_{-i}) \mid r_i(\tilde{\theta}_i)) - c(r_i(\tilde{\theta}_i))\}$$

for some large constant K. Then the average expected payoff when $(\widetilde{s}_i, \overline{s}_{-i})$ is implemented is given by

$$\begin{aligned} & (1-\delta)\overline{w}_i(\widetilde{s}_i, \overline{s}_{-i}) + \delta \sum_{\theta_i} p_i(\theta_i) \sum_{\theta_{-i}} W_i(\widetilde{s}_i(\theta_i), \overline{s}_{-i}(\theta_{-i})) p_i(\theta_{-i} \mid \theta_i) \\ = \; & (1-\delta)\overline{w}_i(\widetilde{s}_i, \overline{s}_{-i}) + K \sum_{\theta_i} p_i(\theta_i) \sum_{\theta_{-i}} p_i(\theta_{-i} \mid \theta_i) \left\{ b(r_{-i}(\overline{s}_{-i}(\theta_{-i})) \mid r_i(\widetilde{s}_i(\theta_i)) - c(r_i(\widetilde{s}_i(\theta_i)) \right\} \\ = \; & (1-\delta)\overline{w}_i(\widetilde{s}_i, \overline{s}_{-i}) + K \sum_{\theta_i} p_i(\theta_i) \left\{ -c(r_i(\widetilde{s}_i(\theta_i)) + \sum_{\theta_{-i}} b(\theta_{-i} \mid r_i(\widetilde{s}_i(\theta_i)) p(\theta_{-i} \mid \theta_i) \right\}. \end{aligned}$$

Hence the gain from deviating from $(\overline{s}_i, \overline{s}_{-i})$ to $(\widetilde{s}_i, \overline{s}_{-i})$ is given by

$$\begin{aligned} & (1-\delta)\rho + K \sum_{\theta_i} p_i(\theta_i) \left\{ -c(r_i(\widetilde{s}_i(\theta_i)) + \sum_{\theta_{-i}} \{ b(\theta_{-i} \mid r_i(\widetilde{s}_i(\theta_i)) p_i(\theta_{-i} \mid \theta_i) \} \right\} \\ & -K \sum_{\theta_i} p_i(\theta_i) \left\{ -c(\theta_i) + \sum_{\theta_{-i}} \{ b(\theta_{-i} \mid \theta_i) p_i(\theta_{-i} \mid \theta_i) \} \right\} \\ = \; & (1-\delta)\rho + K \sum_{\theta_i} p_i(\theta_i) \left\{ -c(r_i(\widetilde{s}_i(\theta_i)) + \sum_{\theta_{-i}} \{ b(\theta_{-i} \mid r_i(\widetilde{s}_i(\theta_i)) p_i(\theta_{-i} \mid \theta_i) \} \right\}, \end{aligned}$$

where ρ is the largest possible one-period deviation gain as defined in (7) and the last equality follows from (10). But, for any $\widetilde{s}_i \neq \overline{s}_i$, we have $r_i(\widetilde{s}_i(\theta_i)) \neq \theta_i$ for some θ_i and, by (9),

$$\sum_{\theta_i} p_i(\theta_i) \left\{ -c(r_i(\widetilde{s}_i(\theta_i)) + \sum_{\theta_{-i}} \{ b(\theta_{-i} \mid r_i(\widetilde{s}_i(\theta_i)) p_i(\theta_{-i} \mid \theta_i) \} \right\} < 0.$$

By choosing K sufficiently large it follows that any deviation induces a loss. Hence, any $\overline{s}$ that is 1-to-1 is enforceable. ∎

Lemma 11 (Extension of KLM for the interdependent value case) *Consider the case of interdependent values. Suppose that the information structure satisfies condition CM and strong pairwise identifiability with respect to f. Then, the following statements hold:*

(a) Any smooth set $W \subset Int(\widetilde{V}^*(f))$ is decomposable on tangent hyperplanes.

(b) For any any smooth set $W \subset Int(\widetilde{V}^*(f))$ there exists $\overline{\delta} < 1$ such that for any $\delta > \overline{\delta}$, W belongs to the set of payoffs that can be generated by some public perfect equilibrium (PPE) of regime $\mathcal{B}$.

(c) Assume that $\widetilde{V}^*(f)$ has non-empty interior. Then there exists $\overline{\delta} < 1$ such that for any $\delta > \overline{\delta}$ the regime $\mathcal{B}$ admits a PPE $\tilde{\sigma} \in Q^{\delta}(\mathcal{B})$, such that every continuation payoff w of $\tilde{\sigma}$ satisfies

$$v(f) > w > v(f) - \eta,$$

where η is defined by (3) *above.*

Proof. It suffices to prove part (a). The remainder of the claim will then follow via the corresponding arguments of KLM as stated in the proof of Lemma 9 above.

With private values, every efficient revelation profile is enforceable, which was used in the corresponding proof of KLM to show part (a). With interdependent values, this is no longer the case. However, by Lemma 10, any 1-to-1 revelation profile $\overline{s}$ is enforceable and this allows us to establish the same conclusion as in the private value case below.

Formally, consider any v on the boundary of W, and let P_v be the hyperplane tangent to W at v. Suppose first that P_v is orthogonal to the i-th axis. Then v either minimizes or maximizes player i's payoff on the set W. In the former case, note that the strategy of always announcing N by all players is a Nash equilibrium and therefore enforceable on P_v. In the latter case, consider $s^i \in \arg\max_{\{s \in S\}} \overline{w}_i(s)$, a strategy profile of mechanism $\tilde{b}^{**}$ which we know from FLM to be enforceable on P_v.

Note that, since $W \subset Int(\widetilde{V}^*(f))$, in both cases above the profile in question is separated from W by P_v.

Next, suppose that P_v is a regular hyperplane. Because $W \subset Int(\widetilde{V}^*(f))$, there exists profile $s \in S$, with $\overline{w}_i(s)$ separated from W by P_v such that either (i) s is a pure 1-to-1 strategy profile that induces payoffs on the Pareto frontier of $\widetilde{V}^*(f)$, or (ii) s_i always announces N for all i. In case (i), Lemma 10 and Theorem 5.1 of KLM imply that s is enforceable with respect to P_v, and the same is trivially true for case (ii). Hence, W is also decomposable on tangent hyperplanes in this case as well. ■

B.3 Equilibrium of Regime B^{ϵ}

Lemma 12 *Consider the case of private values. Suppose that f satisfies full dimensionality and the information structure is pairwise identifiable with respect to f. Then, for* *any $\epsilon > 0$, there exists a strategy profile σ^* in regime B^ϵ and $\overline{\delta}$ such that $\sigma^* \in Q^\delta(B^\epsilon)$ for any $\delta > \overline{\delta}$.*

Proof. By Lemma 9, there exists $\overline{\delta} < 1$ such that for any $\delta > \overline{\delta}$ the regime $\mathcal{B}$ admits a PPE $\tilde{\sigma} \in Q^\delta(\mathcal{B})$, such that every continuation payoff w of $\tilde{\sigma}$ satisfies

$$v(f) > w > v(f) - \eta,$$

where η is defined by (3) above.

Consider regime B^ϵ. For this regime, let σ^* be a strategy profile that at any history at which b^{**} is the mechanism it plays according to $\tilde{\sigma}$ and announces zero: for any i and $\mathbf{h}_i \in \mathbf{H}_i^\infty$ at which g^* is the mechanism, $\sigma_i^*(\mathbf{h}_i, b^{**}, \theta_i) = (\tilde{\sigma}_i(\mathbf{h}_i, \tilde{b}^{**}, \theta_i), 0)$. Since $\tilde{\sigma}$ is an equilibrium of $\mathcal{B}$ given δ, it then follows that σ^* is a Bayesian Nash equilibrium of B^ϵ, given δ.

To see this, note that since $\tilde{\sigma}_i$ is a best response to $\tilde{\sigma}_{-i}$ in $\mathcal{B}$ and since σ^*_{-i} prescribes playing zero integer at every history, it follows that in regime B^ϵ, σ_i^* must be a best response to b^*_{-i} amongst all strategies that choose zero integers at all histories. Furthermore, since σ^*_{-i} prescribes playing zero integer at every history, choosing a positive integer by i at any history induces a payoff of $v_i(f) - \eta$ (Rule B). On the other hand, σ_i^* induces a payoff of $w^* \in W$. Since $w_i^* > v_i(f) - \eta$, it follows that such a deviation is not profitable either. Hence, $\sigma^* \in Q^\delta(B^\epsilon)$. ■

Lemma 13 *Consider the case of interdependent values. Suppose that f satisfies strong full dimensionality and the information structure satisfies condition CM and strong pairwise identifiability with respect to f. Then, for any $\epsilon > 0$, there exists a strategy profile σ^* in regime B^ϵ and $\overline{\delta}$ such that $\sigma^* \in Q^\delta(B^\epsilon)$ for any $\delta > \overline{\delta}$.*

Proof. Given Lemma 11, this claim follows via identical arguments as in Lemma 12 above. ■

References

[1] Baliga, S. (1999): "Implementation in Economic Environments with Incomplete Information: The Use of Multi-Stage Games," *Games and Economic Behavior*, 27, 173-183.

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

[3] Bergemann, D. and S. Morris (2008): “Ex Post Implementation,” *Games and Economic Behavior*, 63, 527-566.

[4] Bergin, J. and A. Sen (1998): “Extensive Form Implementation in Incomplete Information Environments,” *Journal of Economic Theory*, 80, 222-256.

[5] Brusco, S. (1995): “Perfect Bayesian Implementation,” *Economic Theory*, 5, 419-444.

[6] Brusco, S. (2006): “Perfect Bayesian Implementation in Economic Environments,” *Journal of Economic Theory*, 129, 1-30.

[7] Crémer, J. and R. P. McLean (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247-1257.

[8] Dasgupta, P., P. Hammond and E. Maskin (1979): “Implementation of Social Choice Rules: Some General Results on Incentive Compatibility,” *Review of Economic Studies*, 46, 195-216.

[9] d’Aspremont, C. and L-A. Gerard-Varet (1979): “Incentives and Incomplete Information,” *Journal of Public Economics*, 11, 25-45.

[10] Escobar, P. and J. Toikka (2013): “Efficiency in Games with Markovian Private Information,” *Econometrica*, 81, 1887-1934.

[11] Fudenberg, D. and D. K. Levine (1994): “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, 62, 103135.

[12] Fudenberg, D., D. K. Levine and E. Maskin (1994): “The Folk Theorem in Repeated Games with Imperfect Public Information,” *Econometrica*, 62, 103-135.

[13] Fudenberg, D., D. K. Levine and E. Maskin (1996): “The Balanced-Budget Mechanisms with Incomplete Information,” *Mimeo.*

[14] Harris, M. and R. Townsend (1981): “Resource Allocation with Asymmetric Information,” *Econometrica*, 49, 33-64.

[15] Hörner, J., S. Takahashi and N. Vieille (2013): "Truthful Equilibria in Dynamic Bayesian Games," *Mimeo.*

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

[17] Jackson, M. O. and H. F. Sonnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75, 241-257.

[18] Jehiel, P. and B. Moldovanu (2001): "Strictly Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237-1259.

[19] Kosenok, G. and S. Severinov (2008): "Individually Rational, Balanced-Budget Bayesian Mechanism and Surplus Allocation," *Journal of Economic Theory*, 140, 126-161.

[20] Lee, J. and H. Sabourian (2009): "Efficient Repeated Implementation," Cambridge Working Paper in Economics 0948.

[21] Lee, J. and H. Sabourian (2011a): "Efficient Repeated Implementation," *Econometrica*, 79, 1967-1994.

[22] Lee, J. and H. Sabourian (2011b): "Supplement to 'Efficient Repeated Implementation'," *Econometrica Supplemental Material*, 79.

[23] Lee, J. and H. Sabourian (2012): "Repeated Implementation with Finite Mechanisms and Complexity," *Mimeo.*

[24] Mailath, G. J. and L. Samuelson (2006): *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press.

[25] Maskin, E. (1992): "Auctions and Privatization," in *Privatization*, ed. by H. Siebert, Institut für Weltwertschaft an der Universität Kiel.

[26] Mezzetti, C. (2004): "Mechanism Design with Interdependent Valuations: Efficiency," *Econometrica*, 72, 1617-1626.

[27] Mezzetti, C. and L. Renou (2012): "Repeated Nash Implementation," *Mimeo.*

[28] Mookherjee, D. and S. Reichelstein (1990): "Implementation via Augmented Revelation Mechanisms," *Review of Economic Studies*, 57, 453-475.

[29] Myerson, R. B. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47, 61-73.

[30] Palfrey, T. R. and S. Srivastava (1989): "Implementation with Incomplete Information in Exchange Economies," *Econometrica*, 57, 115-134.

[31] Postlewaite, A. and D. Schmeidler (1986): "Implementation with Differential Information Economies," *Journal of Economic Theory*, 39, 14-33.

[32] Renou, L. and T. Tomala (2013): "Approximate Implementation in Markovian Environments," *Mimeo.*

[33] Serrano, R. (2004): "The Theory of Implementation of Social Choice Rules," *SIAM Review*, 46, 377-414.

[34] Sorin, S. (1986): "On Repeated Games with Complete Information," *Mathematics of Operations Research*, 11, 147-160.