Efficient Repeated Implementation I: Complete Information*

Jihong Lee
Birkbeck College, London and Yonsei University†

Hamid Sabourian
University of Cambridge‡

September 2008

Abstract

This paper examines repeatedly implementing a social choice function in a general complete information environment where agents are infinitely-lived and their preferences are determined stochastically in each period. We first demonstrate a necessary role that efficiency plays for repeated implementation. We then establish how any efficient social choice function can be repeatedly implemented in Nash equilibrium.

JEL Classification: A13, C72, C73, D78
Keywords: Repeated implementation, Nash implementation, Efficiency

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

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

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

1 Introduction

1.1 Motivation and overview

Many real world institutions, from voting and markets to contracts, are used repeatedly by their participants. However, despite numerous characterizations of implementable social choice rules in static/one-shot settings in which agents interact only once, implementation theory has yet to offer much to the question of what is generally implementable in repeated contexts (see, for example, the survey of Jackson [8]).

This paper examines repeated implementation of social choice functions in environments in which agents' preferences are determined stochastically across periods and, therefore, a sequence of mechanisms need to be devised in order to repeatedly implement desired social choices. In our setup, the agents are infinitely-lived and their preferences are represented by state-dependent utilities with the state being drawn independently in each period from an identical prior distribution.[1] The realizations of states are assumed to be complete information among the agents. The case of incomplete information is investigated in a companion paper [13].

As is the case between one-shot and repeated games, a repeated implementation problem introduces fundamental differences to what we have learned about implementation in the one-shot context. In particular, one-shot implementability does not imply repeated implementability if the agents can co-ordinate on histories, thereby creating other, possibly unwanted, equilibria.

To gain some intuition, consider a set of agents repeatedly playing the mechanism proposed by Maskin [15] defined for a social choice function satisfying monotonicity and no veto power. Each agent possesses a state-dependent utility function and in each period a state is drawn independently from a fixed distribution.[2] This is simply a repeated game with random states. Since, in the stage game, every Nash equilibrium outcome corresponds to the desired outcome in each state, this repeated game has an equilibrium in which each agent plays the desired action at each period/state regardless of past histories. However, we also know from the study of repeated games, by Dutta [6] and others, that unless minmax expected utility profile of the stage game lies on the efficient payoff frontier of the repeated game we will in general obtain a folk theorem and, therefore, there will be many equilibrium

[1] Utilities are not necessarily transferable.

[2] A detailed example is provided in Section 3 below.

paths along which unwanted outcomes are implemented. Clearly, monotonicity and no veto power are not sufficient to guarantee otherwise. Our results will show that they are not necessary either.

Our understanding of repeated interactions and multiple equilibria gives us several clues about repeated implementation. First, a critical condition for repeated implementation is likely to be some form of efficiency of the social choices, that is, the payoff profile of the social choice function ought to lie on the efficient frontier of the corresponding repeated game/implementation payoffs. Second, we want to devise a sequence of mechanisms such that, roughly speaking, the agents' individually rational payoffs also coincide with the efficient payoff profile of the social choice function.

While repeated play introduces the possibility of co-ordinating on histories, it also allows for more structure in the mechanisms that the planner can enforce. We introduce a sequence of mechanisms, or a *regime*, such that the mechanism played in a given period depends on the past history of mechanisms played and the agents' corresponding actions. This way the infinite future gives the planner additional leverage: the planner can alter the future mechanisms in a way that rewards desirable behavior while punishing the undesirable.

The following notion of repeated implementation is adopted. A social choice function is payoff-repeated-implementable in Nash equilibrium if there exists a regime such that in any of its pure strategy Nash equilibria each agent's continuation payoff (discounted average expected utilities) at any possible history is precisely the payoff that he would obtain if the desired social choices were implemented; similarly, a social choice function is repeated-implementable if the desired outcome is implemented at every history and at every possible state realization in the period.

We first establish a necessary role that efficiency plays for repeated implementation in our setup. If the agents are sufficiently patient and a social choice function is repeated-implementable in Nash equilibrium, then there cannot be another social choice function whose *image* belongs to that of the desired one such that every agent obtains a strictly higher payoff. As the theory of repeated game suggests, the agents can indeed "collude" in our repeated implementation setup if there is a possibility of collective benefits.

We then characterize how an efficient social choice function can be repeatedly implemented. Our sufficiency results are obtained by constructing a regime in which, at any history along an equilibrium path, each agent's continuation payoff has a lower bound equal to the target payoff. The exact match between the equilibrium continuation payoff and the desired payoff of the social choice function here is achieved by applying Sorin's [22] observation that with infinite horizon any payoff can be generated exactly by the discounted average payoff from some sequence of pure action profiles, as long as the discount factor is sufficiently large. (In our setup, the required threshold on discount factor is one half and, therefore, the sufficiency results do not in fact depend on an arbitrarily large discount factor.) Thus, by allowing for history-dependent mechanisms, one can push the individually rational payoff profile of the repeated game to precisely the desired profile. It then follows that if the desired payoffs are located on the efficient frontier of the repeated game payoffs, it is not possible for the agents to sustain any collusion away from the desired payoff profile and outcomes.

Several further comments are worth making about our results. First, it is important to mention that the solution concept is that of Nash equilibrium. Our positive findings do not rely on imposing credibility *off-the-equilibrium*, which have been adopted elsewhere to sharpen predictions (Moore and Repullo [18] and Abreu and Sen [2]). Moreover, in many cases, it is possible to construct regimes such that off-the-equilibrium outcomes are themselves efficient, thereby avoiding *renegotiation* possibilities (Maskin and Moore [16]). Finally, our sufficiency results can be sharpened if the agents have a preference for less *complex* strategies at the margin (Abreu and Rubinstein [1], Kalai and Stanford [11] and Lee and Sabourian [12]).

1.2 Related literature

To this date, only few papers address the problem of repeated implementation. Kalai and Ledyard [10] and Chambers [4] ask the question of implementing an infinite sequence of outcomes when the agents' preferences are fixed. Kalai and Ledyard [10] find that, if the social planner is more patient than the agents and, moreover, is interested only in the long-run implementation of a sequence of outcomes, he can elicit the agents' preferences truthfully in dominant strategies. Chambers [4] applies the intuitions behind the virtual implementation literature to demonstrate that, in a continuous time, complete information setup, any outcome sequence that realizes every feasible outcome for a positive amount of time satisfies monotonicity and no veto power and, hence, is Nash implementable.

In these models, however, there is only one piece of information to be extracted from the agents who, therefore, do not interact repeatedly themselves. More recently, Jackson and Sonnenschein [9] 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, with sufficiently patient agents and sufficiently long horizon, all equilibrium payoffs must approximate the target payoff profile for each agent if the target corresponds to an *ex ante* Pareto efficient social choice function. In contrast to [9], our setup deals with infinitely-lived agents and complete information.[3] In terms of the results, we derive precise, rather than approximate, repeated implementation of an efficient social choice function that does not require the discount factor to be arbitrarily large.

1.3 Plan

The paper is organized as follows. Section 2 introduces the basic definitions and notation associated one-shot implementation. Section 3 then describes the infinitely repeated implementation problem. We present our main results in Section 4 and discuss an extension based on complexity considerations in Section 5. Some concluding remarks are offered in Section 6. Appendix contains some proofs omitted from the main text for expositional reasons.

2 Preliminaries

2.1 Basic definitions and notation

Let I be a finite, non-singleton set of *agents*; with some abuse of notation, let I also denote the cardinality of this set. Let A be a finite set of *outcomes*. Let Θ be a finite, non-singleton set of the possible *states* and p denote a probability distribution defined on Θ such that $p(\theta) > 0$ for all $\theta \in \Theta$. Agent i's state-dependent utility function is given by $u_i : A \times \Theta \to \mathbb{R}$. An *implementation problem*, $\mathcal{P}$, is a collection $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$.

A *social choice function* (SCF), f, in an implementation problem $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$

[3] See our companion paper [13] for further discussion.

is a mapping $f : \Theta \to A$ such that $f(\theta) \in A$ for any $\theta \in \Theta$. The *image* of an SCF f is the set

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

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

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

A *mechanism* (or game form), g, is defined by $g = (M^g, \psi^g)$, where $M^g = M_1^g \times \cdots \times M_I^g$ is a cross product of message spaces and $\psi^g : M^g \to A$ is an outcome function such that $\psi^g(m) \in A$ for any message profile $m = (m_1, \ldots, m_I) \in M^g$. Let G be the set of all feasible mechanisms.

As a benchmark, let us first specify the problem of Nash implementation in the one-shot context. Given a mechanism $g = (M^g, \psi^g)$, we denote by $\mathcal{N}_g(\theta) \subseteq M^g$ the set of Nash equilibria of the game induced by g in state θ. Define

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

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

An SCF f is *monotonic* if, for any $\theta, \theta' \in \Theta$ and $a = f(\theta)$ such that $a \neq f(\theta')$, there exist some i and some $b \in A$ such that $u_i(a, \theta) \geq u_i(b, \theta)$ and $u_i(a, \theta') < u_i(b, \theta')$. An SCF f satisfies *no veto power* if, whenever i, θ and a are such that $u_j(a, \theta) \geq u_j(b, \theta)$ for all $j \neq i$ and all $b \in A$, then $a = f(\theta)$.

The seminal result on Nash implementation is due to Maskin [15]:*(i) If an SCF f is Nash implementable, f satisfies monotonicity. (ii) If an SCF f satisfies monotonicity and no veto power, it is Nash implementable.*

2.2 Efficiency

For an outcome $a \in A$, define $v_i(a) \equiv \sum_\theta p(\theta) u_i(a, \theta)$ as its expected utility to agent i. Similarly, though with some abuse of notation, for an SCF f define $v_i(f) \equiv \sum_\theta p(\theta) u_i(f(\theta), \theta)$. Let $V = \left\{(v_i(f))_{i \in I} \in \mathbb{R}^I : f \in F\right\}$ denote the set of expected utility profiles of all possible SCFs and, for a given $f \in F$, let $V(f) = \left\{(v_i(f'))_{i \in I} \in \mathbb{R}^I : f' \in F(f)\right\}$ be the set of expected utility profiles of all SCFs whose images belong to the image of f. We write $co(V)$ and $co(V(f))$ for the convex hulls of the two sets, respectively.

We now define an *efficient* SCF.

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

Our notion of efficiency is similar to *ex ante Pareto efficiency* used by Jackson and Sonnenschein [9]. The difference is that we define efficiency over the *convex hull* of the set of expected utility profiles of all possible SCFs. As will shortly become clear, this reflects the set of payoffs (discounted average expected utility profiles) of an infinitely repeated implementation problem.[4]

We also define efficiency *on the image* as follows.

Definition 2. *An SCF f is efficient on the image if there exists no $v \in co(V(f))$ such that $v_i \geq v_i(f)$ for all i and $v_i > v_i(f)$ for some i; f is strictly efficient on the image if, in addition, there exists no $f' \in F(f)$, $f' \neq f$, such that $v_i(f') = v_i(f)$ for all i.*

Let us briefly illustrate the conflict between efficiency and monotonicity at this juncture.[5] Consider an implementation problem where $I = \{1, 2, 3, 4\}$, $A = \{a, b, c\}$, $\Theta = \{\theta', \theta''\}$ and the agents' state-contingent utilities are given below:

	θ'				θ''			
	$i =$	$i = 2$	$i = 3$	$i = 4$	$i = 1$	$i = 2$	$i = 3$	$i = 4$
a	3	2	1	3	3	2	1	3
b	1	3	2	2	2	3	3	2
c	2	1	3	1	1	1	2	1

The SCF f is such that $f(\theta') = a$ and $f(\theta'') = b$. Notice that f is *utilitarian* (i.e. maximizes the sum of agents' utilities) and, hence, (strictly) efficient; moreover, in a voting context, such social objectives can be interpreted as representing a scoring rule, such as the Borda count. However, the social choice function is not monotonic. The position of outcome a does not change in any agent's preference ordering across the two states and yet a is $f(\theta')$ but not $f(\theta'')$.

Consider another example. The problem is the same as before except that the state-contingent utilities are given by:

[4] Clearly an efficient f is ex post Pareto efficient in that, for each θ, $f(\theta)$ is Pareto efficient in θ. An ex post Pareto efficient SCF needs not however be efficient.

[5] Some formal results showing the restrictiveness of monotonicity can be found in Mueller and Satterthwaite [20], Dasgupta, Hammond and Maskin [5] and Saijo [21].

	θ'			θ''		
	$i=1$	$i=2$	$i=3$	$i=1$	$i=2$	$i=3$
a	30	0	0	10	0	0
b	0	10	0	0	30	0
c	0	0	20	0	0	20

This is an example of auctions (without transfers). Outcomes a, b and c represent the object being awarded to agent 1, 2, and 3, respectively, and each agent derives positive utility if and only if he obtains the object. In this case, the relative ranking of outcomes does not change for any agent but the social choice may vary with agents' preference intensity such that $f(\theta') = a$ and $f(\theta'') = b$. Here, such a social choice function, which is clearly efficient, has no hope of satisfying monotonicity, or even *ordinality*, which allows for virtual implementation (Matsushima [17] and Abreu and Sen [3]).[6]

3 Repeated implementation

3.1 A motivating example

We begin by discussing an illustrative example. Consider the following case: $I = \{1, 2, 3\}$, $A = \{a, b, c\}$, $\Theta = \{\theta', \theta''\}$ and the agents' state-contingent utilities are given below:

	θ'			θ''		
	$i=1$	$i=2$	$i=3$	$i=1$	$i=2$	$i=3$
a	4	2	2	3	1	2
b	0	3	3	0	4	4
c	0	0	4	0	2	3

The SCF, f, is such that $f(\theta') = a$ and $f(\theta'') = b$. This SCF is monotonic and satisfies no veto power; f is also efficient. The Maskin mechanism, $\mathcal{M} = (M, \psi)$, is such that, for all i, $M_i = \Theta \times A \times \mathbb{Z}_+$ (where $\mathbb{Z}_+$ refers to the set of non-negative integers) and ψ is defined as follows:

[6] An SCF f is ordinal if, whenever $f(\theta) \neq f(\theta')$, there exist some individual i and two outcomes (lotteries) $a, b \in A$ such that $u_i(a, \theta) \geq u_i(b, \theta)$ and $u_i(a, \theta') < u_i(b, \theta')$.

1. if $m_i = (\theta, f(\theta), 0)$ for all i, then $\psi(m) = f(\theta)$;

2. if there exists some i such that $m_j = (\theta, f(\theta), 0)$ for all $j \neq i$ and $m_i = (\cdot, b, \cdot) \neq m_j$, then

$$\begin{aligned} \psi(m) = b \quad & \text{ if } \; u_i(f(\theta), \theta) \geq u_i(b, \theta) \\ \psi(m) = f(\theta) \quad & \text{ if } \; u_i(f(\theta), \theta) < u_i(b, \theta) \; ; \end{aligned}$$

3. if $m = \left((\theta^i, a^i, z^i)\right)_{i \in I}$ is of any other type and i is lowest-indexed agent among those who announce the highest integer, then $\psi(m) = a^i$.

By monotonicity and no veto power of f, for each θ, the unique Nash equilibrium outcome of this mechanism corresponds to $f(\theta)$.

Now, consider the infinitely repeated Maskin mechanism, $\mathcal{M}^\infty$, in this setup. Clearly, this repeated game admits an equilibrium in which the agents play the unique Nash equilibrium in each state regardless of past history, thereby implementing f in each period. However, if the agents are sufficiently patient, there will be other equilibria, for instance, one in which, in each period, outcome b is implemented in both states θ' and θ''.

To see this, consider the following repeated game strategies. Each agent reports $(\theta'', b, 0)$ in each state/period with the following punishment schemes: (i) if either agent 1 or 2 deviates then each agent ignores the deviation and continues to report the same; (ii) if agent 3 deviates then they play the stage game Nash equilibrium in each state/period thereafter independently of subsequent history. It is easy to see that neither agent 1 nor agent 2 has an incentive to deviate. Although agent 1 would prefer a over b in both states, the rules of $\mathcal{M}$ do not allow implementation of a from his unilateral deviation; on the other hand, agent 2 is getting his most preferred outcome in each state. If the discount factor is sufficiently large, agent 3 does not want to deviate either. Agent 3 can indeed alter the implemented outcome in state θ' and obtain c instead of b (after all, this is the agent whose preference reversal supports monotonicity). However, such a deviation would be met by (credible) punishment in which his continuation payoff is a convex combination of 2 (in θ') and 4 (in θ''), which is less than the equilibrium continuation payoff.

In fact, it is straightforward to observe that the Nash equilibrium payoffs (corresponding to the desired social choices) are different from the minmax payoffs of the Maskin mechanism. For instance, agent 1's minmax utility in θ' is equal to 0, which is less than 3, his utility from $f(\theta') = a$; in θ'', minmax utilities of agents 2 and 3 are below their respective utilities from $f(\theta'') = b$. This is the reason why one can obtain numerous equilibrium paths/payoffs with sufficiently patient agents.

The above example highlights the fundamental difference between repeated and one-shot implementation, and suggests that one-shot implementability, characterized by monotonicity and no veto power of an SCF, may be irrelevent for repeated implementability. Our understanding of repeated interactions and the multiplicity of equilibria gives us two clues. First, a critical condition for repeated implementation is likely to be some form of efficiency of the social choices, that is, the payoff profile of the SCF ought to lie on the efficient frontier of the repeated game/implementation payoffs. Second, we want to devise a sequence of mechanisms such that, roughly speaking, the agents' individually rational payoffs also coincide with the efficient payoff profile of the SCF. In what follows, we shall demonstrate that these intuitions are indeed correct and, moreover, achievable.

3.2 Definitions

An *infinitely repeated implementation problem* is denoted by $\mathcal{P}^\infty$, representing infinite repetitions of the implementation problem $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i\in I}]$. Periods are indexed by $t \in \mathbb{Z}_{++}$. In each period, the state is drawn from Θ from an independent and identical probability distribution p.

An (uncertain) infinite sequence of outcomes is denoted by $a^\infty = \left(a^{t,\theta}\right)_{t\in\mathbb{Z}_{++},\theta\in\Theta}$, where $a^{t,\theta} \in A$ is the outcome implemented in period t and state θ. Let A^∞ denote the set of all such sequences. Agents' preferences over alternative infinite sequences of outcomes are represented by discounted average expected utilities. Formally, $\delta \in (0,1)$ is the agents' common discount factor, and agent i's discounted average expected utilities are given by a mapping $\pi_i : A^\infty \to \mathbb{R}$ such that

$$\pi_i(a^\infty) = (1-\delta) \sum_{t\in\mathbb{Z}_{++}} \sum_{\theta\in\Theta} \delta^{t-1} p(\theta) u_i(a^{t,\theta}, \theta) .^7$$

It is assumed that the structure of an infinitely repeated implementation problem (including the discount factor) is common knowledge among the agents and, if there is one,

[7] The assumption of common discount factor is purely for simplifying the exposition. It is also straightforward to extend the analysis to the no discounting case using the Limit of Means specification.

the social planner. The realized state in each period is complete information among the agents but unobservable to an outsider.

We want to repeatedly implement a social choice function in each period by devising a mechanism for each period. We consider the planner choosing a *regime* that specifies a sequence of mechanisms contingent on the history of mechanisms played and the agents' corresponding actions.

Some notation is needed to formally define a regime. Given a mechanism $g = (M^g, \psi^g)$, define $\mathcal{E}^g \equiv \{(g, m)\}_{m \in M^g}$, and let $\mathcal{E} = \cup_{g \in G} \mathcal{E}^g$. Let $H^t = \mathcal{E}^{t-1}$ (the $(t-1)$-fold Cartesian product of $\mathcal{E}$) represent the set of all possible histories of mechanisms played and the agents' corresponding actions over $t-1$ periods. The initial history is empty (trivial) and denoted by $H^1 = \emptyset$. Also, let $H^\infty = \cup_{t=1}^\infty H^t$. A typical (ordered) history of mechanisms and message profiles played is denoted by $h \in H^\infty$.

A regime, R, is then a mapping, or a set of *transition rules*, $R : H^\infty \to G$. 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 $(h, h') \in H^\infty$. A regime R is *history-independent* if and only if, for any $h, h' \in H^t$ and any t, $R(h) = R(h') \in G$. Notice that, in such a history-independent regime, the specified mechanisms may change over time in a pre-determined sequence. We say that a regime R is *stationary* if and only if, for any $h, h' \in H^\infty$, $R(h) = R(h') \in G$.

Given a regime, a (pure) strategy for an agent depends on the sequences of realized states as well as the histories of mechanisms and message profiles played. Define $\mathbf{H}^t$ as the $(t-1)$-fold Cartesian product of the set $\mathcal{E} \times \Theta$, and let $\mathbf{H}^1 = \emptyset$ and $\mathbf{H}^\infty = \cup_{t=1}^\infty \mathbf{H}^t$ with its typical element denoted by $\mathbf{h}$.

Then, each agent i's corresponding strategy, σ_i, is a mapping $\sigma_i : \mathbf{H}^\infty \times G \times \Theta \to \cup_{g \in G} M_i^g$ such that $\sigma_i(\mathbf{h}, g, \theta) \in M_i^g$ for any $(\mathbf{h}, g, \theta) \in \mathbf{H}^\infty \times G \times \Theta$.[8] Let Σ_i be the set of all such strategies, and let $\Sigma \equiv \Sigma_1 \times \cdots \times \Sigma_I$. A strategy profile is denoted by $\sigma \in \Sigma$. We say that σ_i is a *Markov* strategy if and only if $\sigma_i(\mathbf{h}, g, \theta) = \sigma_i(\mathbf{h}', g, \theta)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$, any $g \in G$ and any $\theta \in \Theta$. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is Markov if and only if σ_i is Markov for each i.

Next, 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 $\theta(1) = \emptyset$. Let $q(\theta(t)) \equiv p(\theta^1) \times \cdots \times p(\theta^{t-1})$. Suppose that R is the regime and σ the strategy profile chosen by the agents. Let us define the following

[8] Although we restrict our attention to pure strategies, it is possible to extend the analysis to allow for mixed strategies. See Section 6 below.

variables *on the outcome path*:

- $\mathbf{h}(\theta(t), \sigma, R) \in \mathbf{H}^t$ denotes the $t-1$ period history generated by σ in R over state realizations $\theta(t) \in \Theta^{t-1}$.

- $g^{\theta(t)}(\sigma, R) \equiv (M^{\theta(t)}(\sigma, R), \psi^{\theta(t)}(\sigma, R))$ refers to the mechanism played at $\mathbf{h}(\theta(t), \sigma, R)$.

- $m^{\theta(t),\theta^t}(\sigma, R) \in M^{\theta(t)}(\sigma, R)$ refers to the message profile reported at $\mathbf{h}(\theta(t), \sigma, R)$ when the realized state is θ^t.

- $a^{\theta(t),\theta^t}(\sigma, R) \equiv \psi^{\theta(t)}\left(m^{\theta(t),\theta^t}(\sigma, R)\right) \in A$ refers to the outcome implemented at $\mathbf{h}(\theta(t), \sigma, R)$ when the realized state is θ^t.

- With slight abuse of notation, $\pi_i^{\theta(t)}(\sigma, R)$ denotes agent i's continuation payoff at $\mathbf{h}(\theta(t), \sigma, R)$; that is,

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

  For notational simplicity, let $\pi_i(\sigma, R) \equiv \pi_i^{\theta(1)}(\sigma, R)$.

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

A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is a Nash equilibrium of regime R if, for each i, $\pi_i(\sigma, R) \geq \pi_i(\sigma_i', \sigma_{-i}, R)$ for all $\sigma_i' \in \Sigma_i$. Let $\Omega^\delta(R) \subseteq \Sigma$ denote the set of (pure strategy) Nash equilibria of regime R with discount factor δ.

We are now ready to define the following notions of Nash repeated implementation in the present setup.

Definition 3. *An SCF f is payoff-repeated-implementable in Nash equilibrium from period τ if there exists a regime R such that (i) $\Omega^\delta(R)$ is non-empty; and (ii) every $\sigma \in \Omega^\delta(R)$ is such that $\pi_i^{\theta(t)}(\sigma, R) = v_i(f)$ for all i, any $t \geq \tau$ and any $\theta(t)$.*

Definition 4. *An SCF f is repeated-implementable in Nash equilibrium from period τ if there exists a regime R such that (i) $\Omega^\delta(R)$ is non-empty; and (ii) every $\sigma \in \Omega^\delta(R)$ is such that $a^{\theta(t),\theta^t}(\sigma, R) = f(\theta^t)$ for any $t \geq \tau$, any $\theta(t)$ and any θ^t.*

The first notion represents repeated implementation in terms of payoffs, while the second asks for repeated implementation of outcomes and, therefore, is a stronger concept. Repeated implementation requires a regime in which every Nash equilibrium delivers the correct continuation payoff profile or the correct outcomes for each period/each implementation problem at every possible history.

For simplicity, when we say that f is payoff-repeated-implementable (repeated-implementable), we mean that it is payoff-repeated-implementable (repeated-implementable) from period 1. Also, when we say that f is payoff-repeated-implementable (repeated-implementable) with sufficiently large discount factor, we mean that there exist R and $\bar{\delta}$ such that, for any $\delta \in (\bar{\delta}, 1)$, the properties in Definition 3 (Definition 4) hold.

4 Main results

4.1 Necessity

Our understanding of repeated games suggests that some form of efficiency ought to play a necessary role towards repeated implementation. Our first result establishes that, if the agents are sufficiently patient and an SCF f is repeated-implementable, then there cannot be another SCF whose *image* also belongs to that of f such that all agents strictly prefer it to f in expectation. Otherwise, there must be a collusive equilibrium in which the agents obtain higher payoffs but this is a contradiction.

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

Proof. Suppose not. So, for all δ, we have the following: f is payoff-repeated-implementable in Nash equilibrium but there exists some $f' \in F(f)$ such that $v_i(f') > v_i(f)$ for all i.

Define

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

and

$$\bar{\delta} = \frac{2\rho}{2\rho + \max_i [v_i(f') - v_i(f)]}.$$

Fix any $\delta \in (\bar{\delta}, 1)$. Let R^* be a regime that repeated-implements f. For any i, t, $\theta(t)$ and θ^t, let $\mathcal{A}_i(\mathbf{h}(\theta(t), \sigma, R^*), \theta^t)$ denote the set of outcomes that can result from a unilateral deviation by i from a strategy profile σ at (on-the-outcome-path) history $\left(\mathbf{h}(\theta(t), \sigma, R^*), \theta^t\right)$. Also, with some abuse of notation, for any $a' \in \mathcal{A}_i(\mathbf{h}(\theta(t), \sigma, R^*), \theta^t)$, let $\pi_i^{\theta(t),\theta^t}|a'(\sigma, R^*)$ represent agent i's next period continuation payoff should he make the corresponding deviation at the same history.

Now, fix any $\sigma^* \in \Omega^\delta(R^*)$. Since σ^* is a Nash equilibrium that attains repeated-implementation of f, the following must be true about the equilibrium path: for any i, t, $\theta(t)$, θ^t and $a' \in \mathcal{A}_i(\mathbf{h}(\theta(t), \sigma^*, R^*), \theta^t)$,

$$(1-\delta)u_i(a^{\theta(t),\theta^t}(\sigma^*, R^*), \theta^t) + \delta v_i(f) \geq (1-\delta)u_i(a', \theta^t) + \delta \pi_i^{\theta(t),\theta^t}|a'(\sigma^*, R^*),$$

which implies that

$$\delta \pi_i^{\theta(t),\theta^t}|a'(\sigma^*, R^*) \leq (1-\delta)\rho + \delta v_i(f). \tag{1}$$

Next, since $f' \in F(f)$, there exists a mapping $\phi : \Theta \to \Theta$ such that $f'(\theta) = f(\phi(\theta))$ for all θ. Consider the following strategy profile σ': for any i, $\mathbf{h}$, g, and θ, $\sigma'_i(\mathbf{h}, g, \theta) = \sigma^*_i(\mathbf{h}, g, \phi(\theta))$ if $\mathbf{h}$ is such that there has been no deviation from σ'; $\sigma'_i(\mathbf{h}, g, \theta) = \sigma^*_i(\mathbf{h}, g, \theta)$ otherwise.

Fix any i, t, $\theta(t)$ and θ^t. Also, fix any $a' \in \mathcal{A}_i(\mathbf{h}(\theta(t), \sigma', R^*), \theta^t)$. The agent's continuation payoff from σ' amounts to

$$(1-\delta)u_i(a^{\theta(t),\theta^t}(\sigma', R^*), \theta^t) + \delta v_i(f'). \tag{2}$$

On the other hand, the corresponding payoff from a unilateral deviation resulting in implementation of a' in the current period is

$$(1-\delta)u_i(a', \theta^t) + \delta \pi_i^{\theta(t),\theta^t}|a'(\sigma', R^*). \tag{3}$$

Notice that, by the construction of σ', there exist some $\tilde{\theta}(t)$ such that $\mathbf{h}(\theta(t), \sigma', R^*) = \mathbf{h}(\tilde{\theta}(t), \sigma^*, R^*)$ and, hence, $\mathcal{A}_i(\mathbf{h}(\theta(t), \sigma', R^*), \theta^t) = \mathcal{A}_i(\mathbf{h}(\tilde{\theta}(t), \sigma^*, R^*), \phi(\theta^t))$. Moreover, after a deviation, σ' induces the same continuation strategies as σ^*. Thus, we have

$$\pi_i^{\theta(t),\theta^t}|a'(\sigma', R^*) = \pi_i^{\tilde{\theta}(t),\phi(\theta^t)}|a'(\sigma^*, R^*).$$

Then, by (1) above, the deviation payoff (3) is less than or equal to

$$(1-\delta)\left[u_i(a', \theta^t) + \rho\right] + \delta v_i(f),$$

and, therefore, we derive that, since $v_i(f') > v_i(f)$ and $\delta > \bar{\delta}$, (2) exceeds (3). But, this implies that $\sigma' \in \Omega^\delta(R^*)$ and we have a contradiction against the assumption that R^* repeated-implements f. □

4.2 Sufficiency

Let us now investigate if an efficient SCF can indeed be repeatedly implemented. We begin with some additional definitions and an important general observation. Let

$$v_i^i \equiv \sum_{\theta \in \Theta} p(\theta) \max_{a \in A} u_i(a, \theta)$$

denote agent i's maximal one-period expected utility/payoff, which is unique by finiteness of the set A. Also, define, for $j \neq i$,

$$v_i^j \equiv \sum_{\theta \in \Theta} p(\theta) \max_{a \in A^j(\theta)} u_i(a, \theta) \ ,$$

where $A^j(\theta) \equiv \{\arg\max_{a \in A} u_j(a, \theta)\}$ represents the set of agent j's maximal outcomes in state θ. Thus, v_i^j is the maximum expected utility that agent i can obtain should agent j pick his maximal outcome in each state.

We make the following assumption throughout.

(A) There exist some i, j such that $v_i^i > v_i^j$ and $v_j^j > v_j^i$.

Assumption (A) incurs no loss of generality since we are concerned with repeated implementation of strictly efficient SCFs; otherwise, we can simply let any agent choose the outcome in each period to obtain efficient repeated implementation.[9]

Let $d(i)$ denote a *dictatorial mechanism* in which agent i is the dictator; formally, $d(i) = (M, \psi)$ is such that $M_i = A$, $M_j = \{\emptyset\}$ for all $j \neq i$ and $\psi(m) = m_i$ for all $m \in M$. From $d(i)$, agent i obtains a unique expected utility equal to v_i^i. Let D^i denote a stationary regime in which the dictatorial mechanism $d(i)$ is repeated forever.

We call a *trivial mechanism* one that enforces a single outcome. Formally, $\phi(a) = (M, \psi)$ is such that $M_i = \{\emptyset\}$ for all i and $\psi(m) = a \in A$ for all $m \in M$. Let Φ^a denote a stationary regime in which the trivial mechanism $\phi(a)$ is repeated forever.

[9] This last statement is not quite true if some agents have multiple maximal outcomes in some states that can be Pareto-ranked. Even so, one may expect that renegotiation will ensure a dictatorial outcome to be always efficient.

Let $\mathcal{S}(i,a)$ be the set of all possible history-independent regimes in which the enforced mechanisms are either $d(i)$ or $\phi(a)$ only. Note that any such regime yields a unique payoff to agent i; for any $S \in \mathcal{S}(i,a)$, let $\pi_i(S)$ denote this agent's payoff.[10]

Our results on efficient repeated implementation below are based on the following relatively innocuous auxiliary condition.

Condition ω. For each i, there exists some $\tilde{a}^i \in A$ such that $v_i(f) \geq v_i(\tilde{a}^i)$.

This property says that, for *each* agent, the expected utility that he derives from the SCF is bounded below by that of some constant SCF.[11] Note that the property does not require that there be a single constant social choice function to provide the lower bound for every agent. In many applications, condition ω is naturally satisfied.

Our first Lemma applies the result of Sorin [22] to our setup. If an SCF satisfies condition ω, any individual's corresponding payoff can be generated precisely by a sequence of appropriate dictatorial and trivial mechanisms, as long as the discount factor is greater than a half.

Lemma 1. *Consider an SCF f which satisfies condition ω. Fix any i. For any $\delta \in \left(\frac{1}{2}, 1\right)$, there exist some $a \in A$ and some $S^i \in \mathcal{S}(i,a)$ such that $\pi_i(S^i) = v_i(f)$.*

Proof. If f satisfies condition ω, then, for each i, there exists some outcome $\tilde{a}^i$ such that $v_i(f) \in [v_i(\tilde{a}^i), v_i^i]$ and, hence, $v_i(f)$ is a convex combination of $v_i(\tilde{a}^i)$ and v_i^i. Then, if $\delta > \frac{1}{2}$, the algorithm of Sorin [22] can be applied to generate $v_i(f)$ exactly. Note that the payoff $v_i(f)$ is a convex combination of just two other payoffs, $v_i(\tilde{a})$ and v_i^i, and this gives rise to the requirement that $\delta > \frac{1}{2}$. See Lemma 3.7.1 of Mailath and Samuelson [14] for a full proof. □

For the remainder of the paper, we shall fix δ to be greater than $\frac{1}{2}$ as required by this Lemma. But, we note that if the environment is so rich that there exists some $\tilde{a}^i$ with $v_i(\tilde{a}^i) = v_i(f)$ for all i (for instance, when utilities are transferable) then our main results below are true for any $\delta \in (0,1)$.

[10] If agent i has multiple maximal outcomes in some states, other agents' payoffs from such a regime depend on the actual outcomes chosen by i when he is the dictator and, therefore, are not necessarily unique.

[11] We later discuss an alternative requirement, which we call *non-exclusion*, that can serve the same purpose as condition ω in our analysis.

4.2.1 Three or more agents

We first consider the case of three or more agents. Our arguments are constructive. The discount factor is fixed (greater than one half) unless stated otherwise.

Consider an SCF f that satisfies condition ω. First, mechanism $g^* = (M, \psi)$ is defined as follows:

- For all i, $M_i = \Theta \times \mathbb{Z}_+$.
- Let ψ be:
  1. if $m_i = (\theta, 0)$ for all i, then $\psi(m) = f(\theta)$;
  2. if there exists some i such that $m_j = (\theta, 0)$ for all $j \neq i$ and $m_i \neq m_j$, then $\psi(m) = f(\theta)$;
  3. if $m = \left((\theta^i, z^i)\right)_{i \in I}$ is of any other type and i is lowest-indexed agent among those who announce the highest integer, then $\psi(m) = f(\theta^i)$.

Next, let R^* denote any regime satisfying the following transition rules:

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

The agents start by playing mechanism g^*. At any period in which this mechanism is played, the transition is as follows. If all agents play the same message, then the mechanism next period continues to be g^*. If all agents but one, say i, agree on a message, then the continuation regime at the next period is a history-independent regime in which the odd-man-out i gets a payoff exactly equal to the target level $v_i(f)$ (invoking Lemma 1).[12]

[12] If R^* represents a repeated contract, this feature can be interpreted as awarding the odd-man-out his *expectation damage*, possibly via a transfer from the other contracting parties.

Finally, if the message profile is of any other type, one of the agents who announce the highest integer becomes a dictator forever thereafter.

Notice that, unless all agents agree (when playing mechanism g^*), any strategic play in this regime effectively ends; for any other message profile, the continuation regime is history-independent and employs only dictatorial and/or trivial mechanisms.

Let us proceed by characterizing the set of Nash equilibria of this regime. First, it is straightforward to establish existence: R^* has a Nash equilibrium in Markov strategies which attains truth-telling and, hence, the desired social choice at every possible history.

Lemma 2. *Suppose that f satisfies condition ω. There exists $\sigma^* \in \Omega^\delta(R^*)$, which is Markov, such that, for any t, any $\theta(t)$ and any θ^t, (i) $g^{\theta(t)}(\sigma^*, R^*) = g^*$; and(ii) $a^{\theta(t),\theta^t}(\sigma^*, R^*) = f(\theta^t)$.*

Proof. Consider $\sigma^* \in \Sigma$ such that, for all i, $\sigma_i^*(\mathbf{h}, g^*, \theta) = \sigma_i^*(\mathbf{h}', g^*, \theta) = (\theta, 0)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ and any θ. Thus, at any t and any $\theta(t)$, we have $\pi_i^{\theta(t)}(\sigma^*, R^*) = v_i(f)$ for all i.

Now, consider any i making a unilateral deviation from σ^* by choosing some $\sigma_i' \neq \sigma_i^*$ which announces a different message at some $(\theta(t), \theta^t)$ following some history. But, by rule 2 of g^*, it follows that $a^{\theta(t),\theta^t}(\sigma_i', \sigma_{-i}^*, R^*) = a^{\theta(t),\theta^t}(\sigma^*, R^*) = f(\theta^t)$ while, by transition rule 2(b), we have $\pi_i^{\theta(t),\theta^t}(\sigma_i', \sigma_{-i}^*, R^*) = v_i(f)$. Thus, the deviation is not profitable. □

A critical feature of our regime construction is conveyed in our next Lemma: beyond the first period, as long as g^* is the mechanism played, each agent i's equilibrium continuation payoff is bounded below by the target payoff $v_i(f)$, regardless of the history.

Lemma 3. *Suppose that f satisfies condition ω. Fix any $\sigma \in \Omega^\delta(R^*)$. For any $t > 1$ and any $\theta(t)$, if $g^{\theta(t)}(\sigma, R^*) = g^*$, then $\pi_i^{\theta(t)}(\sigma, R^*) \geq v_i(f)$ for all i.*

Proof. Suppose not. Then, at some $t > 1$ and $\theta(t)$, $\pi_i^{\theta(t)}(\sigma, R^*) < v_i(f)$ for some i. Let $\theta(t) = (\theta(t-1), \theta^{t-1})$. By the transition rules of R^*, it must be that $g^{\theta(t-1)}(\sigma, R^*) = g^*$ and, for all i, $m_i^{\theta(t-1),\theta^{t-1}}(\sigma, R^*) = (\tilde{\theta}, 0)$ for some $\tilde{\theta}$.

Now, consider agent i deviating to another strategy σ_i' which is identical to the equilibrium strategy σ_i at every history, except in period $t-1$ and state θ^{t-1} after history $\mathbf{h}(\theta(t-1), \sigma, R^*)$, where it announces a positive integer. But, by rule 2 of g^*, it follows that

$$a^{\theta(t-1),\theta^{t-1}}(\sigma_i', \sigma_{-i}, R^*) = a^{\theta(t-1),\theta^{t-1}}(\sigma, R^*) = f(\tilde{\theta}) \ ,$$

while, by transition rule 2(b), we have

$$\pi_i^{\theta(t-1),\theta^{t-1}}(\sigma_i', \sigma_{-i}, R^*) = v_i(f) \ .$$

Thus, the deviation is profitable, contradicting the Nash equilibrium assumption. □

We next want to show that indeed mechanism g^* will always be played on the equilibrium path. To this end, let us introduce an additional minor condition.

Condition ν. If there exists some i such that $v_i(f) \leq v_i(a)$ for all $a \in A$ and $v_i(f) = v_i(a')$ for some $a' \in A$, then there exists some $j \neq i$ such that $v_j^j > v_j(a')$.

If f satisfies condition ν as well as condition ω, then, in any equilibrium, agents will always agree and, therefore, g^* will always be played, implementing outcomes belonging only to the *image* of f.

Lemma 4. *Suppose that f satisfies conditions ω and ν. Fix any $\sigma \in \Omega^\delta(R^*)$. For any t, any $\theta(t)$ and any θ^t, we have: (i) $g^{\theta(t)}(\sigma, R^*) = g^*$; (ii) $m_i^{\theta(t),\theta^t}(\sigma, R^*) = (\tilde{\theta}, 0)$ for all i; (iii) $a^{\theta(t),\theta^t}(\sigma, R^*) \in f(\Theta)$.*

Proof. Given the definitions of g^* and R^*, it suffices to establish the following: *For any t and any $\theta(t)$, if $g^{\theta(t)} = g^*$, then, for all i and any θ^t, $m_i^{\theta(t),\theta^t} = (\tilde{\theta}, 0)$ for some $\tilde{\theta}$.*

We first establish the following Fact.

<u>Fact 0</u>. Fix any i and any $a^i(\theta) \in A^i(\theta)$ for each θ. Assumption (A) implies that there exists $j \neq i$ such that $v_j^j > \sum_{\theta \in \Theta} p(\theta) u_j(a^i(\theta), \theta)$.

Proof of Fact 0. If there exists some $j \neq i$ such that $v_j^j > v_j^i$, then this is trivially true. So, suppose that $v_j^j = v_j^i$ for all $j \neq i$. We derive contradiction in the following way.

It suffices to establish the following: there exists some θ in which no $a^* \in A^i(\theta)$ is such that $a^* \in \{\arg\max_{a \in A} u_k(a, \theta)\} \cap \{\arg\max_{a \in A} u_l(a, \theta)\}$ for any $k, l \in I$.

Suppose not. So, suppose that, for some $k, l \in I$ and each θ, there exists some $a^*(\theta) \in A^i(\theta)$ such that $a^*(\theta) \in \{\arg\max_{a \in A} u_k(a, \theta)\} \cap \{\arg\max_{a \in A} u_l(a, \theta)\}$. Then, it must be that $v_k^k = v_k^l = \sum_{\theta \in \Theta} p(\theta) u_k(a^*(\theta), \theta)$ and, similarly, $v_l^l = v_l^k$. But, this contradicts (A).

We now proceed to prove the claim by way of contradiction. So, suppose that, at some t and $\theta(t)$, $g^{\theta(t)} = g^*$ but, for some θ^t, $m^{\theta(t),\theta^t}$ does not belong to rule 1 of g^*. There are two cases to consider.

<u>Case 1</u>: $m^{\theta(t),\theta^t}$ belongs to rule 3 of g^*.

In this case, by transition rule 2(c), the continuation regime at the next period is D^i for some i.

But then, by Fact 0, there must be some $j \neq i$ such that $\pi_j^{\theta(t),\theta^t} < v_j^j$. Consider j deviating to another strategy which yields the same outcome path as the original equilibrium strategy at every history, except in period t and state θ^t after history $\mathbf{h}(\theta(t))$, where it announces θ^i (the state that agent i announces at this history in equilibrium) and an integer higher than any integer that can be reported by the equilibrium profile at this history.

By rule 3 of g^*, such a deviation does not incur a one-period utility loss while strictly improving the continuation payoff as of the next period since the deviator becomes a dictator himself. This is a contradiction.

Case 2: $m^{\theta(t),\theta^t}$ belongs to rule 2 of g^*.

In this case, by transition rule 2(b), the continuation regime, S^i, at the next period belongs to $\mathcal{S}(i,a)$ for some i and some a. There are two possibilities to consider here.

First, suppose that $v_i(f) \leq v_i(a)$ for all a and $v_i(f) = v_i(a')$ for some a'. Thus, $S^i = \Phi^{a'}$. But then, by condition ν, there must exist some $j \neq i$ such that $v_j^j > v_j(a')$ and, hence, $v_j^j > \pi_j^{\theta(t),\theta^t}$. This agent j can profitably deviate by adopting a strategy which is identical to the equilibrium strategy at every history, except in period t and state θ^t after $\mathbf{h}(\theta(t))$ where it announces the same state as in the equilibrium strategy but an integer higher than one reported by i's equilibrium strategy at such a history. This is a contradiction.

Second, suppose otherwise. Thus, $v_i(f) > v_i(a)$ for some a, and S^i can be constructed such that $S^i \in \mathcal{S}(i,a) \backslash \Phi^a$, that is, involving the dictatorial regime $d(i)$. But then, by Fact 0 above, there must exist some $j \neq i$ such that $v_j^j > v_j^i$ and, hence, $v_j^j > \pi_j^{\theta(t),\theta^t}$. This agent j can profitable deviate in the same way as above, and we derive a contradiction. □

Given the previous two Lemmas, we can now tie down the equilibrium payoffs by invoking efficiency *on the image*.

Lemma 5. *Suppose that f is efficient on the image and satisfies conditions ω and ν. Then, for any $\sigma \in \Omega^\delta(R^*)$, $\pi_i^{\theta(t)}(\sigma, R^*) = v_i(f)$ for all i, any $t > 1$ and any $\theta(t)$.*

Proof. Suppose not. So, suppose that f is efficient on the image but there exist some $\sigma \in \Omega^\delta(R^*)$, $t > 1$ and $\theta(t)$ such that $\pi_i^{\theta(t)} \neq v_i(f)$ for some i. By Lemma 3, it must be that $\pi_i^{\theta(t)} > v_i(f)$. We know from part (iii) of Lemma 4 that, in this equilibrium, implemented outcomes always belong to $f(\Theta)$. Then, since f is efficient on the image, it must be that

$\left(\pi_j^{\theta(t)}\right)_{j\in I} \in co(V(f))$ and, therefore, there must exist some $i' \neq i$ such that $\pi_{i'}^{\theta(t)} < v_{i'}(f)$. But, this contradicts Lemma 3. □

We are now ready to present our next main result. With condition ν, we derive repeated implementation of an SCF that is efficient on the image. This minor detail can be avoided, however, with (global) efficiency.

Theorem 2. *Suppose that $I \geq 3$.*

1. *Consider an SCF f that satisfies conditions ω and ν. If f is also efficient on the image, then it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient on the image, then it is repeated-implementable in Nash equilibrium from period 2.*

2. *Consider an SCF f that satisfies condition ω. If f is also efficient, then it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient, then it is repeated-implementable in Nash equilibrium from period 2.*

Proof. 1. The first part of this claim follows immediately from Lemmas 2 and 5. To see the second part, suppose otherwise. But then, by part (iii) of Lemma 4, there must exist some $f' \in F(f)$ such that $v_i(f') = v_i(f)$ for all i. This contradicts strict efficiency of f.

2. Given part 1 above and the definition of condition ν, it suffices to consider the following case: for some i, we have $v_i(f) \leq v_i(a)$ for all a and, for any a' such that $v_i(f) = v_i(a')$, we have $v_j^j = v_j(a')$ for all $j \neq i$.

But then, if f is efficient, it must be that $v_j(f) = v_j(a')$ for all $j \neq i$; if f is strictly efficient in such a case, it must be constant, i.e. $f(\theta) = a'$ for all θ. Such an SCF can be trivially repeated-implemented via the regime $\Phi^{a'}$. □

Note that the above Theorem establishes repeated implementation from the second period and, therefore, unwanted outcomes may still be implemented in the first period. In Section 5 below, we shall demonstrate that the result can be sharpened by appealing to an equilibrium refinement based on complexity considerations.

4.2.2 Two agents

As in one-shot Nash implementation (Moore and Repullo [19] and Dutta and Sen [7]), the two-agent case brings non-trivial differences to the analysis. In particular, with three or more agents a unilateral deviation from "consensus" can be detected; with two agents, however, it is not possible to identify the misreport in the event of disagreement. In our repeated implementation setup, this creates a difficulty in establishing existence of an equilibrium in the canonical regime.

Let us introduce a *bad outcome* (with respect to an SCF); that is, consider an SCF f such that there exists $\tilde{a} \in A$ such that, for all i, we have $v_i(\tilde{a}) < v_i(f)$. Note, first, that we define a bad outcome in terms of expected utilities and, hence, it is a weaker requirement than the similar condition appearing in Moore and Repullo [19], which requires an outcome strictly dominated by $f(\theta)$ in each θ. Second, note that conditions ω and ν are vacuously satisfied in the presence of a bad outcome.

We can also make a stronger assumption about the bad outcome. we shall say that there exists a *strictly bad outcome* when considering an SCF f with the following property: there exists $\tilde{a} \in A$ such that, for each i, $u_i(\tilde{a}, \theta) \leq u_i(f(\theta), \theta)$ for all θ with the inequality being strict for at least one state.[13]

Theorem 3. *Suppose that $I = 2$.*

1. *Consider an SCF f such that there is a bad outcome. If f is efficient on the image, then it is payoff-repeated-implementable in Nash equilibrium from period 2 with sufficiently large discount factor; if f is strictly efficient on the image, then it is repeated-implementable in Nash equilibrium from period 2 with sufficiently large discount factor.*

2. *Consider an SCF f such that there is a strictly bad outcome. If f is efficient on the image, then it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient on the image, then it is repeated-implementable in Nash equilibrium from period 2.*

Proof. See Appendix. □

[13]Note that this definition is still weaker than the similar condition in Moore and Repullo [19], which requires the inequality to be *strict* for *every* state.

We note that the need for sufficiently large δ in the first part of Theorem 3 above arises from the definition of the bad outcome assumption. It serves to punish both players in the event of disagreement. But, each agent prefers it to the social choices only *on average* and, therefore, it is possible for an agent to make a one-period utility gain by deviating from agreement. For this reason, the same result can be obtained independently of the discount factor under the stronger definition of strict bad outcome.

5 Complexity-averse agents

Theorems 2 and 3 above do not guarantee period 1 implementation of the SCF. We next show that this can be achieved if the agents have a preference for less complex strategies at the margin.[14]

Consider *any* measure of complexity of a strategy under which a Markov strategy is simpler than a non-Markov strategy.[15] Then, refine Nash equilibrium as follows: a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ constitutes a Nash equilibrium with complexity cost (NEC) of regime R if, for all i, (i) σ_i is a best response to σ_{-i}; and (ii) there exists no σ_i' such that σ_i' is a best response to σ_{-i} at every history and σ_i' is simpler than σ_i.

In this refinement, complexity enters an agent's preferences *lexicographically*. By definition, the set of NECs is a subset of the set of Nash equilibria of a regime.[16] Also, it is important to note that the above notion of NEC is a much weaker refinement than other notions previously considered in the complexity literature (Abreu and Rubinstein [1], Kalai and Stanford [11], Lee and Sabourian [12], among others). Not only it allows for different measures of complexity, it asks for a simpler strategy that is a best response to the others' strategies *at every history*, not merely at the beginning of the game.

Let $\Omega^{\delta,c}(R)$ denote the set of NECs of regime R with discount factor δ. The following extends the notion of Nash repeated implementation to the case with complexity-averse

[14] We nonetheless assume that the social planner, if exists, does not care about the complexity of his chosen regime.

[15] There are many complexity notions that possess this property. One example is provided by Kalai and Stanford [11] who measure the number of *continuation strategies* that a strategy induces at different periods/histories of the game.

[16] Note that the complexity cost here concerns the cost associated with implementation, rather than computation, of a strategy.

agents.[17]

Definition 5. *An SCF f is repeated-implementable in Nash equilibrium with complexity cost if there exists a regime R such that (i) $\Omega^{\delta,c}(R)$ is non-empty; and (ii) every $\sigma \in \Omega^{\delta,c}(R)$ is such that $a^{\theta(t),\theta^t}(\sigma, R) = f(\theta)$ for any t, any $\theta(t)$ and any θ^t.*

We next report the following result for the case of $I \geq 3$ that sharpens Theorem 2. A corresponding result for the two-agent case can be similarly derived and, hence, omitted for expositional flow.

Theorem 4. *Suppose that $I \geq 3$. If an SCF f is strictly efficient and satisfies condition ω, then it is repeated-implementable in Nash equilibrium with complexity cost.*

Proof. See Appendix. □

In order to derive this result, we characterize the set of NECs of regime R^* defined above. Since, by definition, a NEC is also a Nash equilibrium, Lemmas 3-5 remain true for NEC. Moreover, since a Markov Nash equilibrium is itself a NEC, by Lemma 2, $\Omega^{\delta,c}(R^*)$ is non-empty. In the proof, we additionally establish that complexity considerations select only Markovian behavior in R^*. This, together with Theorem 2, delivers the result.

6 Concluding discussion

6.1 Non-exclusive SCF

In our analysis thus far, repeated implementation of an efficient SCF has been obtained with an auxiliary condition that, for each agent, the (one-period) expected utility from implementation of the SCF must be bounded below by that of some constant SCF. The role of this condition is to construct, for each agent, a history-independent and non-strategic continuation regime in which the agent derives a payoff equal to the target level. We next define another condition that can fulfil the same role.

Non-exclusion. For each i, there exists some $j \neq i$ such that $v_i(f) \geq v_i^j$.[18]

[17] We can similarly define payoff repeated implementation with complexity cost. It is straightforward to derive the corresponding results. We shall therefore omit their statements and proofs.

[18] Note that this property replaces both conditions ω and ν in the case of three or more agents.

The name of the property comes from the fact that, otherwise, there must exist some agent i such that $v_i(f) < v_i^j$ for all $j \neq i$; in other words, there exists an agent who strictly prefers a dictatorship by any other agent to the SCF itself.

Suppose now that each agent has a *unique* maximal outcome in each state. Then, if non-exclusion is satisfied, for any i, one can build a history-independent regime that appropriately alternates dictatorial mechanisms $d(i)$ and $d(j)$ for some $j \neq i$ (instead of $d(i)$ and some trivial mechanism $\phi(a)$) such that the agent's payoff is precisely the target level $v_i(f)$.

We cannot say that either condition ω (or ω') or non-exclusion is a weaker requirement than the other. Consider the following two examples.

First, consider $\mathcal{P}$ where $I = \{1, 2\}$, $A = \{a, b\}$, $\Theta = \{\theta', \theta''\}$, $p(\theta') = p(\theta'') = 1/2$ and the agents' state-contingent utilities are given below:

	θ'			θ''		
	$i = 1$	$i = 2$	$i = 3$	$i = 1$	$i = 2$	$i = 3$
a	1	3	2	3	2	1
b	3	2	1	1	3	2
c	2	1	3	2	1	3

Here, the SCF f such that $f(\theta') = a$ and $f(\theta'') = b$ is efficient but fails to satisfy condition ω because of agent 1. But, notice that f is non-exclusive.

Second, consider $\mathcal{P}$ where $I = \{1, 2, 3\}$, $A = \{a, b, c, d\}$, $\Theta = \{\theta', \theta''\}$ and the agents' state-contingent utilities are given below:

	θ'			θ''		
	$i = 1$	$i = 2$	$i = 3$	$i = 1$	$i = 2$	$i = 3$
a	3	2	0	1	2	1
b	2	1	1	2	3	0
c	1	3	1	3	1	1
d	0	0	0	0	0	0

Here, the SCF such that $f(\theta') = a$ and $f(\theta'') = b$ is efficient and also satisfies condition ω, but excludes agent 3.

6.2 Off the equilibrium

In one-shot implementation, it has been shown that one can improve the range of achievable objectives by employing extensive form mechanisms together with refinements of Nash equilibrium as solution concept (Moore and Repullo [18] and Abreu and Sen [2]).

Although this paper also considers a dynamic setup, the solution concept adopted is that of Nash equilibrium and our results do not rely on imposing a particular behavioral assumption off-the-equilibrium to "kill off" unwanted equilibria. Recall that, in each of the regimes constructed to prove our sufficiency results above, off-the-equilibrium continuation paths can involve either an indefinite repetition of a dictatorial or a trivial mechanism, or some combination of only the two. Neither case presents a strategic situation for the agents and, therefore, the issue of off-the-equilibrium credibility simply does not arise.

On the other hand, it is important to note that our existence results do not involve construction of Nash equilibria based on non-credible threats off-the-equilibrium. The equilibria identified for our main Theorems all satisfy subgame perfection. Thus, we can replicate the same set of results with subgame perfect equilibrium as the solution concept.

A related issue is that of efficiency of off-the-equilibrium paths. In one-shot implementation, it is often the case that off-the-equilibrium inefficiency is imposed in order to sustain desired outcomes on-the-equilibrium. Several authors have, therefore, investigated to what extent the possibility of *renegotiation* affects the scope of implementability (for example, Maskin and Moore [16]). In our repeated implementation setup with three or more agents, this needs not be a cause for concern since off-the-equilibrium outcomes in our regimes can actually be made efficient (the two-agent case involves an inefficient bad outcome). Recall that the requirement of condition ω is that, for each agent i, there exists some outcome $\tilde{a}^i$ which gives the agent an expected utility less than or equal to that of the SCF. If the environment is rich enough, such an outcome can indeed be found on the efficient frontier itself. Moreover, if the SCF is non-exclusive, the regimes can be constructed so that off-the-equilibrium is entirely associated with dictatorial outcomes, which are efficient.

7 Extensions

In this paper, we have restricted our attention to implementation in pure strategies only. It is in fact possible to extend our analysis to include mixed strategies. Conceptually, allowing for players' randomization has similar effects as assuming incomplete information. Therefore, the incomplete information analysis of the companion paper [13], can be modified to handle the case of mixed strategies under complete information. As can be seen in [13], however, this will generate somewhat weaker results than what is reported here.

Another important extension is to generalize the process with which individual preferences evolve. In this paper, and also in the companion paper, we consider the case in which preferences follow an i.i.d. process. If we were to introduce a history-dependent underlying distribution, the notion of repeated implementation might itself need a modification. Notice that our definition of efficiency actually depends on the distribution and, therefore, a general stochastic process opens door for time-varying social objectives. We shall leave this question for future research.

8 Appendix

Proof of Theorem 3 A proof of the second part of this Theorem is straightforward once the the first part has been established. We shall therefore present a proof of the first part only.

There exists a bad outcome $\tilde{a}$ such that, for each i, $v_i(\tilde{a}) < v_i(f) - \epsilon$ for some $\epsilon > 0$. From Lemma 1, let S^i be the sequence of dictatorial mechanism $d(i)$ and trivial mechanism $\phi(\tilde{a})$ such that $\pi_i(S^i) = v_i(f)$. Since S^i involves d^i, Assumption (A) above means that there exists some $\epsilon' > 0$ such that $v_1^1 > \pi_1(S^2) + \epsilon'$ and $v_2^2 > \pi_2(S^1) + \epsilon'$. As before, define $\rho \equiv \max_{i,\theta,a,a'} [u_i(a,\theta) - u_i(a',\theta)]$ and, also, $\bar{\delta} \equiv \frac{\rho}{\rho + \min\{\epsilon,\epsilon'\}}$.

Define mechanism $\tilde{g}^* = (M, \psi)$ as:

- For all i, $M_i = \Theta \times \mathbb{Z}_+$.
- ψ is such that

  1. if $m_i = (\theta, \cdot)$ and $m_j = (\theta, \cdot)$, then $\psi(m) = f(\theta)$;
  2. if $m_i = (\theta^i, z^i)$, $m_j = (\theta^j, 0)$ and $z^i \neq 0$, then $\psi(m) = f(\theta^j)$;
  3. for any other m, $\psi(m) = \tilde{a}$.

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

1. $\tilde{R}^*(\emptyset) = \tilde{g}^*$;

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

   (a) if $m_i^{t-1} = (\theta, 0)$ and $m_j^{t-1} = (\theta, 0)$, then $\tilde{R}^*(h) = \tilde{g}^*$;

   (b) if $m_i^{t-1} = (\theta^i, 0)$, $m_j^{t-1} = (\theta^j, 0)$ and $\theta^i \neq \theta^j$, then $\tilde{R}^*(h) = \Phi^{\tilde{a}}$;

   (c) if $m_i^{t-1} = (\theta^i, z^i)$, $m_j^{t-1} = (\theta^j, 0)$ and $z^i \neq 0$, then $\tilde{R}^*|h = S^i$;

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

Lemmas A1-A4 below complete the proof.

Lemma A1: Fix any $\sigma \in \Omega^\delta(\tilde{R}^*)$. For any $t > 1$ and any $\theta(t)$, if $g^{\theta(t)} = \tilde{g}^*$, then $\pi_i^{\theta(t)} \geq v_i(f)$.

Proof. Suppose not. So, suppose that at some $t > 1$ and $\theta(t)$, $g^{\theta(t)} = \tilde{g}^*$ but $\pi_i^{\theta(t)} < v_i(f)$ for some i. Let $\theta(t) = (\theta(t-1), \theta^{t-1})$. Given the transition rules, it must be that $g^{\theta(t-1)} = \tilde{g}^*$ and $m_i^{\theta(t-1),\theta^{t-1}} = m_j^{\theta(t-1),\theta^{t-1}} = (\cdot, 0)$.

Consider i deviating at $\left(\mathbf{h}(\theta(t-1)), \theta^{t-1}\right)$ such that the state reported is the same as in his original equilibrium strategy but the integer reported is positive. Given ψ, such a deviation does not alter the current period's implemented outcome but, by transition rule 2(c), it activates a continuation regime at the next period from which i obtains a payoff exactly $v_i(f)$. Hence, the deviation is profitable, and we have a contradiction. □

Lemma A2: Fix any $\delta \in \left(\bar{\delta}, 1\right)$ and any $\sigma \in \Omega^\delta(\tilde{R}^*)$. For any t and any $\theta(t)$, if $g^{\theta(t)} = \tilde{g}^*$, then $m_i^{\theta(t),\theta^t} = m_j^{\theta(t),\theta^t} = (\theta, 0)$ for any θ^t.

Proof. We prove the claim by contradiction. So, suppose that, for some t, $\theta(t)$ and θ^t, $g^{\theta(t)} = \tilde{g}^*$ but $m^{\theta(t),\theta^t}$ is not as in the claim. There are three cases to consider.

Case 1: $m_i^{\theta(t),\theta^t} = (\cdot, z^i)$ and $m_j^{\theta(t),\theta^t} = (\cdot, z^j)$ with $z^i, z^j > 0$.

In this case, by rule 3 of ψ, $\tilde{a}$ is implemented in the current period and, by transition rule 2(d), a dictatorship, by, say, i, follows forever thereafter. But then, by assumption (A) above, $v_j^j > v_j^i$ and, therefore, j can profitably deviate by announcing an integer higher than z^i at such a history; such a deviation does not alter the current period outcome from $\tilde{a}$ but switches dictatorship to himself as of the next period. This is a contradiction.

Case 2: $m_i^{\theta(t),\theta^t} = (\cdot, z^i)$ and $m_j^{\theta(t),\theta^t} = (\theta^j, 0)$ with $z^i > 0$.

In this case, by rule 2 of ψ, $f(\theta^j)$ is implemented in the current period and, by transition rule 2(c), continuation regime S^i follows thereafter.

Consider j deviating to another strategy identical to σ_j everywhere except, here, it announces an integer higher than z^i. Given rule 3 of ψ and transition rule 2(d), such a deviation yields a continuation payoff

$$(1-\delta)u_j(\tilde{a},\theta^t)+\delta v_j^j \ , \tag{4}$$

while the corresponding equilibrium payoff is

$$(1-\delta)u_j(f(\theta^j),\theta^t)+\delta\pi_j(S^i) \ . \tag{5}$$

But then, since $v_j^j > \pi_j(S^i)+\epsilon'$ and $\delta > \bar{\delta}$, (4) exceeds (5) and the deviation is profitable. Thus, we have a contradiction.

Case 3: $m_i^{\theta(t),\theta^t} = (\theta^i, 0)$ and $m_j^{\theta(t),\theta^t} = (\theta^j, 0)$ with $\theta^i \neq \theta^j$.

In this case, by rule 3 of ψ, $\tilde{a}$ is implemented in the current period and, by transition rule 2(b), in every period thereafter.

Consider any agent i deviating by announcing a positive integer. Given rule 2 of ψ and transition rule 2(c), such a deviation yields a continuation payoff

$$(1-\delta)u_i(f(\theta^j),\theta^t)+\delta v_i(f) \ , \tag{6}$$

while the corresponding equilibrium payoff is

$$(1-\delta)u_i(\tilde{a},\theta^t)+\delta v_i(\tilde{a}) \ , \tag{7}$$

But then, since $v_i(f) > v_i(\tilde{a})+\epsilon$ and $\delta > \bar{\delta}$, (6) exceeds (7) and the deviation is profitable. Thus, we have a contradiction. □

Lemma A3: Suppose that f is efficient on the image. Fix any $\delta \in (\bar{\delta},1)$. Then, for any $\sigma \in \Omega^\delta(\tilde{R}^*)$, $\pi_i^{\theta(t)} = v_i(f)$ for all i, any $t>1$ and any $\theta(t)$.

Proof. Given Lemmas A1-A2, we can directly apply the proof of Lemma 5 above. □

Lemma A4: Fix any $\delta \in (\bar{\delta},1)$. Then, $\Omega^\delta(\tilde{R}^*)$ is non-empty.

Proof. Consider a symmetric Markov strategy profile in which the true state and zero integer are always reported.

Each agent i can deviate in one of the following three ways:

(i) Announce the true state but a positive integer. Given rule 1 of ψ and transition rule 2(c), such a deviation is not profitable.

(ii) Announce a false state and a positive integer. Given rule 2 of ψ and transition rule 2(c), such a deviation is not profitable.

(iii) Announce zero integer but a false state. In this case, by rule 3 of ψ, $\tilde{a}$ is implemented in the current period and, by transition rule 2(b), in every period thereafter. The gain from such a deviation cannot exceed

$$(1-\delta)\left(\max_{\theta} u_i(\tilde{a},\theta) - \min_{a,\theta} u_i(a,\theta)\right) - \delta\epsilon < 0 \ ,$$

where the inequality holds since $\delta > \bar{\delta}$. Thus, the deviation is not profitable. □

Proof of Theorem 4 It suffices to show that every $\sigma \in \Omega^{\delta,c}(R^*)$ is Markov. Suppose that there exists some $\sigma \in \Omega^{\delta,c}(R^*)$ such that σ_i is non-Markov for some i. Then, consider i deviating to a Markov strategy, $\sigma_i' \neq \sigma_i$, which always announces a positive integer when playing g^*.

Fix any $\theta^1 \in \Theta$. By part 2 of Lemma 4 and rule 2 of g^*, we have $a^{\theta^1}(\sigma_i', \sigma_{-i}, R^*) = a^{\theta^1}(\sigma, R^*)$; by transition rule 2(b) of R^*, we have $\pi_i^{\theta^1}(\sigma_i', \sigma_{-i}, R^*) = v_i(f)$ for any θ^1. We know from Lemma 5 that $\pi_i^{\theta^1}(\sigma, R^*) = v_i(f)$ for any θ^1. Thus, the deviation does not alter i's payoff. But, since σ_i' is less complex than σ_i, such a deviation is worthwhile for i.

By similar arguments, σ_i' is the best response to σ_{-i} at every (off-the-outcome-path) history. This contradicts the NEC assumption.

References

[1] Abreu, D., and A. Rubinstein, "The Structure of Nash Equilibria in Repeated Games with Finite Automata," *Econometrica*, 56 (1988), 1259-1282.

[2] Abreu, D. and A. Sen, "Subgame Perfect Implementaton: A Necessary and Almost Sufficient Condition," *Journal of Economic Theory*, 50 (1990), 285-299.

[3] Abreu, D. and A. Sen, "Virtual Implementation in Nash Equilibrium," *Econometrica*, 59 (1991), 997-1021.

[4] Chambers, C. P., "Virtual Repeated Implementation," *Economics Letters*, 83 (2004), 263-268.

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

[6] Dutta, P. K., "A folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66 (1995), 1-32.

[7] Dutta, B. and A. Sen, "A Necessary and Sufficient Condition for Two-Person Nash Implementation," *Review of Economic Studies*, 58 (1991), 121-128.

[8] Jackson, M. O., "A Crash Course in Implementation Theory," *Social Choice and Welfare*, 18 (2001), 655-708.

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

[10] Kalai, E. and J. O. Ledyard, "Repeated Implementation," *Journal of Economic Theory*, 83 (1998), 308-317.

[11] Kalai, E. and W. Stanford, "Finite Rationality and Interpersonal Complexity in Repeated Games," *Econometrica*, 56 (1988), 397-410.

[12] Lee, J. and H. Sabourian, "Coase Theorem, Complexity and Transaction Costs," *Journal of Economic Theory*, 135 (2007), 214-235.

[13] Lee, J. and H. Sabourian, "Efficient Repeated Implementation II: Incomplete Information," *mimeo* (2008).

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

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

[16] Maskin, E. and J. Moore, "Implementation and Renegotiation," *Review of Economic Stidies*, 66 (1999), 39-56.

[17] Matsushima, H., "A New Approach to the Implementation Problem," *Journal of Economic Theory*, 45 (1988), 128-144.

[18] Moore, J. and R. Repullo, "Subgame Perfect Implementation," *Econometrica*, 56 (1988), 1191-1220.

[19] Moore, J. and R. Repullo, "Nash Implementation: A Full Characterization," *Econometrica*, 58 (1990), 1083-1099.

[20] Mueller, E. and M. Satterthwaite, "The Equivalence of Strong Positive Association and Strategy-proofness," *Journal of Economic Theory*, 14 (1977), 412-418.

[21] Saijo, T., "On Constant Maskin Monotonic Social Choice Functions," *Journal of Economic Theory*, 42 (1987), 382-386.

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