Efficient Repeated Implementation*

Jihong Lee†
Seoul National University

Hamid Sabourian‡
University of Cambridge

December 2010

Abstract

This paper examines repeated implementation of a social choice function (SCF) with infinitely-lived agents whose preferences are determined randomly in each period. An SCF is repeated-implementable in Nash equilibrium if there exists a sequence of (possibly history-dependent) mechanisms such that (i) its Nash equilibrium set is non-empty and (ii) every Nash equilibrium outcome corresponds to the desired social choice at every possible history of past play and realizations of uncertainty. We show, with minor qualifications, that in the complete information environment an SCF is repeated-implementable in Nash equilibrium if and only if it is efficient. We also discuss several extensions of our analysis, including constructions employing only finite mechanisms.

JEL Classification: A13, C72, C73, D02, D70

Keywords: Repeated implementation, Nash implementation, Efficiency, Mixed strategies, Finite mechanisms, Complexity

---

*The authors are grateful to Stephen Morris and three anonymous referees for their helpful comments and suggestions that have led to the present version of the paper. We have also benefited from conversations with Bhaskar Dutta, Matt Jackson, Eric Maskin and Roberto Serrano. Jihong Lee acknowledges financial support from the Korea Research Foundation Grant funded by the Korean Government (KRF-2008-327-B00103).

†Department of Economics, Seoul National University, Seoul 151-746, Korea, jihonglee@snu.ac.kr

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

1 Introduction

Implementation theory, sometimes referred to as the theory of *full* implementation, has been concerned with designing mechanisms, or game forms, that implement desired social choices in *every* equilibrium of the mechanism. Numerous characterizations of implementable social choice rules have been obtained in one-shot settings in which agents interact only once. However, many real world institutions, from voting and markets to contracts, are used repeatedly by their participants. Despite its relevance, implementation theory has yet to offer much to the question of what is generally implementable in repeated contexts (see, for example, the surveys of Jackson [14], Maskin and Sjöström [23] and Serrano [32]).[1]

In many repeated settings, the agent's preferences change over time in an uncertain manner and the planner's objective is to repeatedly implement the same social choice for each possible preference profile. A number of applications naturally fit this description. In repeated voting or auctions, the voters' preferences over candidates or the bidders' valuations over the objects could follow a stochastic process, with the planner's goal being for instance to always enact an outcome that is Condorcet-consistent or to sell each object to the bidder with highest valuation. Similarly, a community that collectively owns a technology could repeatedly face the problem of efficiently allocating resources under changing circumstances.

This paper examines such a repeated implementation problem in complete information environments. In our setup, the agents are infinitely-lived and their preferences are represented by state-dependent utilities with the state being drawn randomly in each period from an identical prior distribution. Utilities are not necessarily transferable, and the realizations of states are complete information among the agents.[2]

In the one-shot implementation problem with complete information, the critical condition for implementing a social choice rule is the well-known (Maskin) monotonicity. This condition is necessary and, together with some minor qualification, also sufficient.[3] As is the case between one-shot and repeated games, however, a repeated implementation

[1]The literature on dynamic mechanism design does not address the issue of full implementation since it is concerned only with establishing a single equilibrium of some mechanism with desired properties.

[2]A companion paper [19] explores the case of incomplete information.

[3]Monotonicity can be a strong requirement. Some formal results showing its restrictiveness can be found in Mueller and Satterthwaite [27], Dasgupta, Hammond and Maskin [9] and Saijo [29].

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 social choice function that satisfies sufficiency conditions for Nash implementation in the one-shot complete information setup (e.g. monotonicity and no veto power) and a mechanism that implements it (e.g. Maskin [21]). Suppose now that the agents play this mechanism repeatedly and in each period a state is drawn independently from a fixed distribution, with its realization being complete information.[4] 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 history. However, we also know from the study of repeated games (e.g. Mailaith and Samuelson [20]) that unless minmax payoff profile of the stage game lies on the efficient payoff frontier of the repeated game, by the Folk theorem, there will be many equilibrium paths along which unwanted outcomes are implemented. Thus, the conditions that guarantee one-shot implementation are not sufficient for repeated implementation. Our results below also show that they are not necessary either.

Given the multiple equilibria and collusion possibilities in repeated environments, at first glance, implementation in such settings seems a daunting task. But our understanding of repeated interactions also provides us with several clues as to how it may be achieved. 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 need 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 by the agents, thereby creating difficulties towards full repeated implementation, 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.

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

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. In fact, we observe institutions with similar features. For instance, many constitutions involve explicit provisions for amendment,[5] while a designer of repeated auctions or other repeated allocation mechanisms often commits to excluding collusive bidders or free-riders from future participation.

Formally, we consider repeated implementation of a social choice function (henceforth, called SCF) in the following sense: there exists a regime such that (i) its equilibrium set is non-empty and (ii) in any equilibrium of the regime, the desired social choice is implemented at every possible history of past play of the regime and realizations of states. A weaker notion of repeated implementation seeks the equilibrium continuation payoff (discounted average expected utility) of each agent at every possible history to correspond precisely to the one-shot payoff (expected utility) of the social choices. Our main analysis adopts Nash equilibrium as the solution concept.[6]

We first demonstrate the following necessity result: if the agents are sufficiently patient and an SCF is repeated-implementable, it cannot be strictly Pareto dominated (in terms of expected utilities) by any convex combination of SCFs whose *ranges* belongs to that of the desired SCF. Just 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.

It is then shown that, under some minor conditions, any SCF that is efficient in the range can be repeatedly implemented. This sufficiency result is obtained by constructing for each SCF a canonical regime in which, at any history along an equilibrium path, each agent's continuation payoff has a lower bound equal to his payoff from the SCF, thereby ensuring the individually rational payoff profile in any continuation game to be no less than the desired profile. It then follows that if the desired payoff profile is located on the efficient frontier the agents cannot sustain any collusion away from it; moreover, if there is a unique SCF associated with such payoffs than repeated implementation of the desired outcomes is achieved.

The construction of the canonical regime involves two steps. We first show, for each player i, that there exists a regime S^i in which the player obtains a payoff exactly equal

[5] Barbera and Jackson [6] explore the issue of "stability" of constitutions (voting rules).

[6] Our results do not rely on imposing credibility *off*-the-equilibrium to sharpen predictions, as done in Moore and Repullo [25], Abreu and Sen [3] and others.

to that from the SCF and, then, embed this into the canonical regime such that each agent i can always induce S^i in the continuation game by an appropriate deviation from his equilibrium strategy. The first step is obtained by applying Sorin's [31] observation that with infinite horizon any payoff can be generated exactly by the discounted average payoff from some sequence of outcomes, as long as the discount factor is sufficiently large.[7] The second step is obtained by allowing each agent the possibility of making himself the "odd-one-out" in any equilibrium.

Our main analysis is extended in several directions. In particular, we also explore what can be achieved with regimes employing only *finite* mechanisms. While this does not alter the results when we restrict attention to pure strategies, the use of finite mechanisms creates the problem of unwanted mixed strategy equilibria. The problem is particularly severe in one-shot settings because Jackson [13] shows that even when an SCF is Nash implementable there could still be a mixed equilibrium that *strictly Pareto dominates* it. In our repeated setup, Jackson's criticism works in our favor. We construct a regime with finite mechanisms in which, with minor qualifications, every mixed Nash equilibrium is itself dominated by other (pure) equilibria that achieves desired efficient implementation.

Furthermore, using the same construction, and focusing on subgame perfect equilibria, we eliminate randomization altogether by introducing a mild equilibrium refinement based on complexity considerations. Specifically, in our construction, any equilibrium involving mixing can only be sustained by strategies that are *unnecessarily too complex*; that is, there are simpler strategies that would generate the same payoff at any on- or off-the-equilibrium history of the game. Therefore, such mixed equilibria will not be chosen when the agents have, at least at the margin, a preference for less complex strategies.

To this date, only few papers address the problem of repeated implementation. Kalai and Ledyard [17] and Chambers [7] ask the question of implementing an infinite sequence of outcomes when the agents' preferences are fixed. Kalai and Ledyard [17] find that, if the 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 [7] 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

[7] In our setup, the threshold on discount factor required for the main sufficiency results is one half and, therefore, an arbitrarily large discount factor is not needed.

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 [16] consider "budgeted" mechanisms in a finitely linked, or repeated, incomplete information implementation problem with *independent private values.* They find that for any *ex ante* Pareto efficient SCF all equilibrium payoffs of such a budgeted mechanism must approximate the target payoffs corresponding to the SCF, as long as the agents are sufficiently patient and the horizon is sufficiently long. In contrast to [16], our setup deals with infinitely-lived agents and the case of complete information (see [19] for our incomplete information analysis). In terms of results, we derive a necessary condition as well as precise, rather than approximate, repeated implementation of an efficient SCF at every possible history of the regime, not just the payoffs computed at the outset. The sufficiency results do not require the discount factor to be arbitrarily large and are obtained with arguments that are very much distinct from those of [16].

The paper is organized as follows. Section 2 first introduces the implementation problem in the one-shot setup with complete information which will lay out the basic definitions and notation throughout the paper. Section 3 then describes the problem of infinitely repeated implementation. Our main results are presented and discussed in Section 4. We consider the key extensions of our analysis in Section 5 before concluding in Section 6. Some proofs are relegated to an Appendix. Also, we provide a Supplementary Material to present some results and proofs whose details are left out for expositional reasons.

2 Preliminaries

Let I be a finite, non-singleton set of agents; with some abuse of notation, I also denotes the cardinality of this set. Let A be a finite set of outcomes and 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 $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$.

An SCF f in an implementation problem $\mathcal{P}$ is a mapping $f : \Theta \to A$ such that $f(\theta) \in A$ for any $\theta \in \Theta$. The *range* of f is the set $f(\Theta) = \{a \in A : a = f(\theta) \text{ for some } \theta \in \Theta\}$. 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 range belongs to $f(\Theta)$.

For an outcome $a \in A$, define $v_i(a) = \sum_{\theta \in \Theta} p(\theta) u_i(a, \theta)$ as its expected utility, or (one-shot) payoff, to agent i. Similarly, though with some abuse of notation, for an SCF f define $v_i(f) = \sum_{\theta \in \Theta} p(\theta) u_i(f(\theta), \theta)$. Denote the profile of payoffs associated with f by $v(f) = (v_i(f))_{i \in I}$. Let $V = \{v(f) \in \mathbb{R}^I : f \in F\}$ be the set of expected utility profiles of all possible SCFs. Also, for a given $f \in F$, let $V(f) = \{v(f') \in \mathbb{R}^I : f' \in F(f)\}$ be the set of payoff profiles of all SCFs whose ranges belong to the range of f. We write $co(V)$ and $co(V(f))$ for the convex hulls of the two sets, respectively.

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. Furthermore, v' *strictly* Pareto dominates v if the inequality is strict for *all* i. An *efficient* SCF is defined as follows.

Definition 1 *An SCF f is efficient if there exists no $v' \in co(V)$ that Pareto dominates $v(f)$; f is strictly efficient if it is efficient and there exists no $f' \in F$, $f' \neq f$, such that $v(f') = v(f)$.*

Our notion of efficiency is similar to *ex ante* Pareto efficiency used by Jackson and Sonnenschein [16]. 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 (discounted average) payoffs that can be obtained in an infinitely repeated implementation problem.[8]

We also define efficiency *in the range* as follows.

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

As a benchmark, we next specify Nash implementation in the one-shot context. 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 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.

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. We then say that an SCF f is *Nash implementable* if

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

there exists a mechanism g such that, for all $\theta \in \Theta$, $\psi^g(m) = f(\theta)$ for all $m \in \mathcal{N}_g(\theta)$. The seminal result on (one-shot) Nash implementation is due to Maskin [21]: (i) If an SCF f is Nash implementable, f satisfies monotonicity; (ii) If $I \geq 3$, and if f satisfies monotonicity and no veto power, f is Nash implementable.[9] As mentioned before, monotonicity can be a restrictive condition, and one can easily find cases in standard problems such as voting or auction where, for example, efficient SCFs are not monotonic and hence not (one-shot) Nash implementable.[10]

3 Repeated Implementation

3.1 An Illustrative Example

We begin our analysis of repeated implementation by discussing an example that will illustrate the key issues. 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 efficient, monotonic and satisfies no veto power. The Maskin mechanism, $\mathcal{M} = (M, \psi)$, for f is defined as follows: $M_i = \Theta \times A \times \mathbb{Z}_+$ (where $\mathbb{Z}_+$ is the set of non-negative integers) for all i and ψ satisfies

1. if $m_i = (\theta, f(\theta), 0)$ for all i, $\psi(m) = f(\theta)$;
2. if there exists some i such that $m_j = (\theta, f(\theta), 0)$ for all $j \neq i$ and $m_i = (\cdot, \tilde{a}, \cdot) \neq m_j$, $\psi(m) = \tilde{a}$ if $u_i(f(\theta), \theta) \geq u_i(\tilde{a}, \theta)$ and $\psi(m) = f(\theta)$ if $u_i(f(\theta), \theta) < u_i(\tilde{a}, \theta)$;

[9] 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 $i \in I$ and $b \in A$ 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 $a = f(\theta)$.

[10] An efficient SCF may not even satisfy *ordinality*, which allows for virtual implementation (Matsushima [24] and Abreu and Sen [4]).

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

By monotonicity and no veto power of f, for each θ, the unique Nash equilibrium of $\mathcal{M}$ consists of each agent announcing $(\theta, f(\theta), 0)$, thereby inducing outcome $f(\theta)$.

Next, consider the infinitely repeated version of Maskin mechanism, where in each period state θ is drawn randomly and the agents play the same Maskin mechanism. Clearly, this repeated game with random states admits an equilibrium in which the agents play the unique Nash equilibrium of the stage game 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 and the SCF cannot be (uniquely) implemented.

For instance, consider the following repeated game strategies which implement outcome b in both states of each period. 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 each agent plays 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 sufficiently patient, agent 3 does not want to deviate either. This player can deviate in state θ' and obtain c instead of b but this would be met by punishment in which his continuation payoff is a convex combination of 2 (in θ') and 4 (in θ''), which is less than the equilibrium payoff.

In the above example, we have deliberately chosen an SCF that is efficient (as well as monotonic and satisfying no veto power) so that the Maskin mechanism in the one-shot framework induces unique Nash equilibrium payoffs on its efficient frontier. Despite this, we cannot repeatedly implement the SCF via a repeated Maskin mechanism. The reason is that in this example the Nash equilibrium payoffs differ from the minmax payoffs of the stage game. For instance, agent 1's minmax utility in θ' is equal to 0, resulting from $m_2 = m_3 = (\theta'', f(\theta''), 0)$, which is less than his utility from $f(\theta') = a$; in θ'', minmax utilities of agents 2 and 3, which both equal 2, are below their respective utilities from $f(\theta'') = b$. As a result, the set of individually rational payoffs in the repeated game is not singleton, and 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 irrelevant 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 (repeated game) payoffs are given by a mapping $\pi_i : A^\infty \to \mathbb{R}$ such that

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

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

We want to repeatedly implement an SCF in each period by devising a mechanism for each period. A *regime* specifies a sequence 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.

To formally define a regime, we need some notation. 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 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 t and any $h, h' \in H^t$, $R(h) = R(h')$. 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 $h, h' \in H^\infty$, $R(h) = R(h')$.[11]

Given a regime, a (pure) strategy for an agent depends on the sequence of realized states as well as the history of mechanisms and message profiles played.[12] 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 $(\mathbf{h}, g, \theta) \in \mathbf{H}^\infty \times G \times \Theta$. 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 (history-independent) strategy if and only if $\sigma_i(\mathbf{h}, g, \theta) = \sigma_i(\mathbf{h}', g, \theta)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$, $g \in G$ and $\theta \in \Theta$. 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 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

[11] A constitution (over voting rules) can therefore be thought of as a regime in the following sense. In each period, each agent reports his preference over the candidate outcomes and also chooses a voting rule to be enforced in the next period. The current voting rule aggregates the agents' first reports, while the amendment rule dictates the transition according to the second reports.

[12] We later extend the analysis to allow for mixed (behavioral) strategies. See Section 5.1.

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 current 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 current state is θ^t.

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

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

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

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 any i, $t \geq \tau$ and $\theta(t)$. An SCF f is repeated-implementable in Nash equilibrium from period τ if, in addition, every $\sigma \in \Omega^\delta(R)$ is such that $a^{\theta(t),\theta^t}(\sigma, R) = f(\theta^t)$ for any $t \geq \tau$, $\theta(t)$ and 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 from some period τ requires the existence of a regime in which every Nash equilibrium delivers the correct continuation payoff profile or the correct outcomes from period τ onwards for every possible sequence of state realizations.

4 Main Results

4.1 Necessity

As illustrated by the example in Section 3.1, our understanding of repeated games suggests that some form of efficiency ought to play a necessary role towards repeated implementation. However, note that any constant SCF is trivially repeated-implementable, implying that an SCF needs not be efficient over the entire set of possible SCFs. Our first result establishes the following: if the agents are sufficiently patient and an SCF f is repeated-implementable from any period, then there cannot be a payoff vector v' belonging to the convex hull of all feasible payoffs that can be constructed from the *range* of f such that all agents strictly prefer v' to $v(f)$.

We demonstrate this result by showing that, if this were not the case, there would be a "collusive" equilibrium in which the agents obtain the higher payoff vector v'. To construct this collusive equilibrium, we first invoke the result by Fudenberg and Maskin [11] on convexifying the set of payoffs without public randomization in repeated games to show that, with sufficiently large δ, there exists a sequence of non-truthful announcements and corresponding outcomes in the range of f such that the payoff profile v' is obtained. Then, we show that these announcements can be supported in equilibrium by constructing strategies in which any unilateral deviation triggers the original equilibrium in the continuation game (that repeated-implements f).

Theorem 1 *Consider any SCF f such that $v(f)$ is strictly Pareto dominated by another payoff profile $v' \in co\,(V(f))$. Then there exists $\bar{\delta} \in (0,1)$ such that, for any $\delta \in (\bar{\delta}, 1)$ and period τ, f is not repeated-implementable in Nash equilibrium from period τ.*[13]

Proof. By assumption, there exists $\epsilon > 0$ such that $v'_i > v_i(f) + 2\epsilon$ for all i. Let $\delta^1 = \frac{2\rho}{2\rho+\epsilon}$, where $\rho \equiv \max_{i \in I, \theta \in \Theta, a, a' \in A}\left[u_i(a,\theta) - u_i(a',\theta)\right]$.

Since $v' \in co\,(V(f))$, there exists $\delta^2 > 0$ such that, for all $\delta \in (\delta^2, 1)$, there exists an infinite sequence of SCFs $F' = \{f^1, f^2, \ldots\}$ such that

$$f^t \in F(f) \text{ for all integer } t \tag{1}$$

[13]The necessary condition here requires the payoff profile of the SCF f to lie on the frontier of $co(V(f))$. Thus, it will correspond to efficiency in the range when $co(V(f))$ is strictly convex.

and, for any t',

$$\left| v' - (1-\delta) \sum_{t \geq t'} \delta^{t-t'} v(f^t) \right| < \epsilon. \tag{2}$$

The proof of this claim is analogous to the standard result by Fudenberg and Maskin [11] on convexifying the set of payoffs without public randomization in repeated games (see Lemma 3.7.2 of Mailath and Samuelson [20]).

Next, let $\overline{\delta} = \max\{\delta^1, \delta^2\}$. Fix any $\delta \in (\overline{\delta}, 1)$ and any sequence $F' = \{f^1, f^2, \ldots\}$ that satisfies (1) and (2) for any date t'. Also, fix any date τ. We want to show that f cannot be repeatedly implemented from period τ. Suppose not; then there exists a regime R^* that repeated-implements f from period τ.

For any strategy profile σ in regime R^*, any player i, any date t and any sequence of states $\theta(t)$, let $M_i(\theta(t), \sigma, R^*)$ and $\psi^{\theta(t)}(\sigma, R)$ denote, respectively, the set of messages that i can play and the corresponding outcome function at history $\mathbf{h}(\theta(t), \sigma, R^*)$. Also, with some abuse of notation, for any $m_i \in M_i(\theta(t), \sigma, R^*)$ and any $\theta^t \in \Theta$, let $\pi_i^{\theta(t),\theta^t}(\sigma)|m_i$ represent i's continuation payoff from period $t+1$ if the sequence of states $(\theta(t), \theta^t)$ is observed, i deviates, by playing m_i, from σ_i for only one period at $\mathbf{h}(\theta(t), \sigma, R^*)$ after observing θ^t and every other agent plays the regime according to σ_{-i}.

Consider any $\sigma^* \in \Omega^\delta(R^*)$. Since σ^* is a Nash equilibrium that repeated-implements f from period τ, the following must be true about the equilibrium path: for any i, $t \geq \tau$, $\theta(t)$, θ^t and $m_i' \in M_i(\theta(t), \sigma^*, R^*)$,

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

where $a \equiv \psi^{\theta(t)}(\sigma^*, R^*)(m_i', m_{-i}^{\theta(t),\theta^t}(\sigma^*, R^*))$. This implies that, for any i, $t \geq \tau$, $\theta(t), \theta^t$ and $m_i' \in M_i(\theta(t), \sigma^*, R^*)$,

$$\delta \pi_i^{\theta(t),\theta^t}(\sigma^*)|m_i' \leq (1-\delta)\rho + \delta v_i(f). \tag{3}$$

Next, note that, since $f^t \in F(f)$, there must exist a mapping $\lambda^t : \Theta \to \Theta$ such that $f^t(\theta) = f(\lambda^t(\theta))$ for all θ. Consider the following strategy profile σ': for any i, g, and θ, (i) $\sigma_i'(\mathbf{h}, g, \theta) = \sigma_i^*(\mathbf{h}, g, \theta)$ for any $\mathbf{h} \in \mathbf{H}^t$, $t < \tau$; (ii) for any $\mathbf{h} \in \mathbf{H}^t$, $t \geq \tau$, $\sigma_i'(\mathbf{h}, g, \theta) = \sigma_i^*(\mathbf{h}, g, \lambda^t(\theta))$ if $\mathbf{h}$ is such that there has been no deviation from σ', while $\sigma_i'(\mathbf{h}, g, \theta) = \sigma_i^*(\mathbf{h}, g, \theta)$ otherwise.

Then, by (2), we have

$$\pi_i^{\theta(t)}(\sigma', R) = (1-\delta) \sum_{t \geq \tau} \delta^{t-\tau} v(f^t) > v_i' - \epsilon \text{ for all } i, t \geq \tau \text{ and } \theta(t) \tag{4}$$

Given the definitions of σ' and $\sigma^* \in \Omega^\delta(R^*)$, and since $v_i' - \epsilon > v_i(f)$, (4) implies that it pays no agent to deviate from σ' at any history before period τ.

Next, fix any player i, any date $t \geq \tau$, any sequence of states $\theta(t)$ and any state θ^t. By (4), we have that agent i's continuation payoff from σ' at $\mathbf{h}(\theta(t), \sigma', R^*)$ after observing θ^t is no less than

$$(1-\delta)u_i\left(a^{\theta(t),\theta^t}(\sigma', R^*), \theta^t\right) + \delta(v_i' - \epsilon). \tag{5}$$

On the other hand, the continuation payoff of i from any unilateral one-period deviation $m_i' \in M_i(\theta(t), \sigma', R^*)$ from σ' at $\theta(t)$ and θ^t is given by

$$(1-\delta)u_i\left(a', \theta^t\right) + \delta\pi_i^{\theta(t),\theta^t}(\sigma')|m_i', \tag{6}$$

where $a' = \psi^{\theta(t)}(\sigma', R^*)(m_i', m_{-i}^{\theta(t),\theta^t}(\sigma', R^*))$.

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

$$\pi_i^{\theta(t),\theta^t}(\sigma')|m_i' = \pi_i^{\tilde{\theta}(t),\lambda^t(\theta^t)}(\sigma^*)|m_i'.$$

Then, by (3) above, the deviation payoff (6) is less than or equal to

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

This, together with $v_i' > v_i(f) + 2\epsilon$, $\delta > \bar{\delta} = \max(\delta^1, \delta^2)$ and the definition of δ^1, implies that (5) exceeds (6). But, this means that it does not pay any agent i to deviate from σ' at any date $t \geq \tau$. Therefore, σ' must also be a Nash equilibrium of regime R^*.

Since, by (4), $\pi_i^{\theta(t)}(\sigma', R^*) > v_i' - \epsilon > v_i(f) = \pi_i^{\theta(t)}(\sigma^*, R^*)$ for any i, $t \geq \tau$ and $\theta(t)$, we then have a contradiction against the assumption that R^* repeated-implements f from period τ. ■

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.

First, we call a *constant rule* mechanism one that enforces a single outcome (constant SCF). Formally, $\phi(a) = (M, \psi)$ is such that $M_i = \{\emptyset\}$ for all i and $\psi(m) = a \in A$ for all

$m \in M$. Also, let $d(i)$ denote a *dictatorial* mechanism in which agent i is the dictator, or simply i-dictatorship; 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 $m \in M$.

Next, let $v_i^i = \sum_{\theta \in \Theta} p(\theta) \max_{a \in A} u_i(a, \theta)$ denote agent i's maximal one-period payoff. Clearly, v_i^i is i's payoff when i is the dictator and he acts rationally. Also, let $A^i(\theta) \equiv \{\arg\max_{a \in A} u_i(a, \theta)\}$ represent the set of i's best outcomes in state θ. Define the maximum payoff i can obtain when agent $j \neq i$ is the dictator by $v_i^j = \sum_{\theta \in \Theta} p(\theta) \max_{a \in A^j(\theta)} u_i(a, \theta)$.

We make the following assumption throughout the paper.

(A) There exist some i and j such that $A^i(\theta) \cap A^j(\theta)$ is empty for some θ.

This assumption is equivalent to assuming that $v_i^i \neq v_i^j$ for some i and j. It implies that in some state there is a conflict between some agents on the best outcome. Since we are concerned with repeated implementation of efficient SCFs, Assumption (A) incurs no loss of generality when each agent has a unique best outcome for each state: if Assumption (A) were not to hold, we could simply let any agent choose the outcome in each period to obtain repeated implementation of an efficient SCF.

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. One could compare it to the *bad outcome* condition appearing in Moore and Repullo [26] (which requires existence of an outcome *strictly* worse than the desired social choice for *all* agents in *every* state). Condition ω is weaker for three reasons. First, condition ω does not require that there be a single constant SCF to provide the lower bound for all agents; second, for each i, outcome $\tilde{a}^i$ is worse than the SCF only on average; third, the inequality is weak. In many applications, condition ω is naturally satisfied (e.g. zero consumption in the group allocation problem mentioned in the Introduction). Furthermore, there are other properties that can serve the same role, which we discuss in Section 4.3 below.

Now, let Φ^a denote a stationary regime in which the constant rule mechanism $\phi(a)$ is repeated forever and let D^i denote a stationary regime in which the dictatorial mechanism $d(i)$ is repeated forever. Also, 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. For any $i, j \in I$, $a \in A$ and $S^i \in \mathcal{S}(i, a)$, we denote by $\pi_j(S^i)$ the maximum payoff j can obtain when S^i is enforced and agent i always chooses a best outcome under $d(i)$.

Our first Lemma applies the result of Sorin [31] 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 constant rule mechanisms, as long as the discount factor is greater than a half.

Lemma 1 *Consider an SCF f and any i. Suppose that there exists $\tilde{a}^i \in A$ such that $v_i(f) \geq v_i(\tilde{a}^i)$. Then, for any $\delta > \frac{1}{2}$, there exists $S^i \in \mathcal{S}(i, \tilde{a}^i)$ such that $\pi_i(S^i) = v_i(f)$.*

Proof. By assumption there exists some outcome $\tilde{a}^i$ such that $v_i(f) \in [v_i(\tilde{a}^i), v_i^i]$. Since $v_i(\tilde{a}^i)$ is the one-period payoff of i when $\phi(\tilde{a}^i)$ is the mechanism played and v_i^i is i's payoff when $d(i)$ is played and i behaves rationally, it follows from the algorithm of Sorin [31] (see Lemma 3.7.1 of Mailath and Samuelson [20]) that there exists a regime $S^i \in \mathcal{S}(i, \tilde{a}^i)$ that alternates between $\phi(a)$ and $d(i)$, and generates the payoff $v_i(f)$ exactly. ■

The above statement assumes that $\delta > \frac{1}{2}$ because $v_i(f)$ is a convex combination of exactly *two* payoffs $v_i(\tilde{a}^i)$ and v_i^i. For the remainder of the paper, unless otherwise stated, δ will be fixed to be greater than $\frac{1}{2}$ as required by this Lemma. But, note that if the environment is sufficiently rich that, for each i, one can find some $\tilde{a}^i$ with $v_i(\tilde{a}^i) = v_i(f)$ (for instance, when utilities are quasi-linear and monetary transfers can be arranged) then our results below are true for any $\delta \in (0, 1)$.

Three or more agents The analysis with three or more agents is somewhat different from that with two players. We begin with the former case and assume that $I \geq 3$.

Our arguments are constructive. First, fix any SCF f that satisfies condition ω and define mechanism $g^* = (M, \psi)$ as follows: $M_i = \Theta \times \mathbb{Z}_+$ for all i, and ψ is such that (i) if $m_i = (\theta, \cdot)$ for at least $I - 1$ agents, $\psi(m) = f(\theta)$ and (ii) if $m = ((\theta^i, z^i))_{i \in I}$ is of any other type, $\psi(m) = f(\tilde{\theta})$ for some arbitrary but fixed state $\tilde{\theta} \in \Theta$.

Next, we define our canonical regime. Let R^* denote any regime in which $R^*(\emptyset) = g^*$ and, for any $h = ((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 1$ and $g^{t-1} = g^*$, the following transition rules hold:

Rule 1: If $m_i^{t-1} = (\cdot, 0)$ for all i, $R^*(h) = g^*$.

Rule 2: If there exists some i such that $m_j^{t-1} = (\cdot, 0)$ for all $j \neq i$ and $m_i^{t-1} = (\cdot, z^i)$ with $z^i > 0$, $R^*|h = S^i$, where $S^i \in \mathcal{S}(i, \tilde{a}^i)$ such that $v_i(\tilde{a}^i) \leq v_i(f)$ and $\pi_i(S^i) = v_i(f)$ (by condition ω and Lemma 1, regime S^i exists).

Rule 3: If m^{t-1} is of any other type and i is lowest-indexed agent among those who announce the highest integer, $R^*|h = D^i$.

Regime R^* starts with mechanism g^*. At any period in which this mechanism is played, the transition is as follows. If all agents announce zero, then the mechanism next period continues to be g^*. If all agents but one, say i, announce zero and i does not, then the continuation regime at the next period is a history-independent regime in which the "odd-one-out" i can guarantee himself a payoff exactly equal to the target level $v_i(f)$ (invoking Lemma 1). Finally, if the message profile is of any other type, one of the agents who announce the highest integer becomes a dictator forever thereafter.

Note that, unless all agents "agree" on zero when playing mechanism g^*, the game effectively ends; for any other message profile, the continuation regime is history-independent and employs only dictatorial and/or constant rule mechanisms.

We now characterize the set of Nash equilibria of regime R^*. 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 always bounded below by the target payoff $v_i(f)$. Otherwise, the agent whose continuation payoff falls below the target level could profitably deviate by announcing a positive integer in the previous period, thereby making himself the "odd-one-out" and hence guaranteeing the target payoff.

Lemma 2 *Suppose that f satisfies condition ω. Fix any $\sigma \in \Omega^\delta(R^*)$. For any $t > 1$ and $\theta(t)$, 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, $m_i^{\theta(t-1),\theta^{t-1}}(\sigma, R^*) = (\theta, 0)$ for some θ.

Consider agent i deviating to another strategy σ_i' identical to the equilibrium strategy σ_i at every history, except at $\mathbf{h}(\theta(t-1), \sigma, R^*)$ and period $t-1$ state θ^{t-1} where it announces the state announced by σ_i, θ, and a positive integer. Note that the outcome function ψ of mechanism g^* is independent of integers and, therefore, the outcome at

$(\mathbf{h}(\theta(t-1), \sigma, R^*), \theta^{t-1})$ does not change, i.e. $a^{\theta(t-1),\theta^{t-1}}(\sigma_i', \sigma_{-i}, R^*) = a^{\theta(t-1),\theta^{t-1}}(\sigma, R^*)$. But, by Rule 2, S^i will be the continuation regime at the next period and i can obtain continuation payoff $v_i(f)$. Thus, the deviation is profitable, contradicting the Nash equilibrium assumption. ■

We next show that indeed mechanism g^* will always be played on the equilibrium path. Note that, in our dynamic construction, the agents play an "integer game" over the identity of dictator in the continuation game. Given Assumption (A), when two or more agents announcing a positive integer, there must be another agent who can profitably deviate to a higher integer. In order to ensure that there cannot be an equilibrium with an "odd-one-out," say i, and hence the continuation regime S^i, we also assume the following: for each i and $\tilde{a}^i \in A$ used for constructing S^i,

$$\text{if } v_i(f) = v_i(\tilde{a}^i) \text{ then } v_j^j > v_j(\tilde{a}^i) \text{ for some } j. \tag{7}$$

Lemma 3 *Suppose that f satisfies ω. Also, suppose that, for each i, outcome $\tilde{a}^i \in A$ used in the construction of S^i above satisfies (7). Then, for any $\sigma \in \Omega^\delta(R^*)$, t, $\theta(t)$ and θ^t, we have: (i) $g^{\theta(t)}(\sigma, R^*) = g^*$; (ii) $m_i^{\theta(t),\theta^t}(\sigma, R^*) = (\cdot, 0)$ for all i; (iii) $a^{\theta(t),\theta^t}(\sigma, R^*) \in f(\Theta)$.*

Proof. Note that $R^*(\emptyset) = g^*$. Thus, by Rule 1 and induction, and by ψ, it suffices to show the following: *For any t and $\theta(t)$, if $g^{\theta(t)} = g^*$ then $m_i^{\theta(t),\theta^t} = (\cdot, 0)$ for all i and θ^t.*

We shall use proof by contradiction. To do so, we first establish two claims that will ensure that, if the statement were not true, Assumption (A) and the extra condition (7) would imply existence of an agent who could profitably deviate.

Claim 1: *Fix any i and any $a^i(\theta) \in A^i(\theta)$ for every θ. There exists $j \neq i$ such that $v_j^j > \sum_\theta p(\theta) u_j(a^i(\theta), \theta)$.*

To prove this claim, suppose otherwise; then $v_j^j = \sum_\theta p(\theta) u_j(a^i(\theta), \theta)$ for all $j \neq i$. But this means that $a^i(\theta) \in A^j(\theta)$ for all $j \neq i$ and θ. Since by assumption $a^i(\theta) \in A^i(\theta)$, this contradicts Assumption (A).

Claim 2: *Fix any $\sigma \in \Omega^\delta(R^*)$, t, $\theta(t)$ and θ^t. If $g^{\theta(t)} = g^*$ and $m_i^{\theta(t),\theta^t} = (\cdot, z^i)$ with $z^i > 0$ for some i then there must exist some $j \neq i$ such that $\pi_j^{\theta(t),\theta^t} < v_j^j$.*

To prove this claim note that, given the definition of R^*, the continuation regime at the next period is either D^i or S^i for some i. There are two cases to consider.

Case 1: The continuation regime is $S^i = \Phi^{\tilde{a}^i}$ (S^i enforces $\tilde{a}^i \in A$ at every period).

In this case $\pi_i^{\theta(t),\theta^t} = v_i(f) = v_i(\tilde{a}^i)$. Then the claim follows from $\pi_j^{\theta(t),\theta^t} = v_j(\tilde{a}^i)$ and condition (7).

Case 2: The continuation regime is either D^i or $S^i \neq \Phi^{\tilde{a}^i}$.

By assumption under $d(i)$ every agent j receives at most $v_j^i \leq v_j^j$. Also, when the constant rule mechanism $\phi(\tilde{a}^i)$ is played every agent j receives a payoff $v_j(\tilde{a}^i) \leq v_j^j$. Since in this case the continuation regime involves playing either $d(i)$ or $\phi(\tilde{a}^i)$, it follows that, for every j, $\pi_j^{\theta(t),\theta^t} \leq v_j^j$. Furthermore, by Claim 1, it must be that this inequality is strict for some $j \neq i$. This is because in this case there exists some $t' > t$ and some sequence of states $\theta(t') = (\theta(t), \theta^{t+1}, ..\theta^{t'-1})$ such that the continuation regime enforces $d(i)$ at history $\mathbf{h}(\theta(t'))$; but then $a^{\theta(t'),\theta} \in A^i(\theta)$ for all θ and therefore, by Claim 1, there exists an agent $j \neq i$ such that $v_j^j > \sum_\theta p(\theta) u_j(a^{\theta(t'),\theta}, \theta)$.

Now, suppose that, at some t and $\theta(t)$, $g^{\theta(t)} = b^*$ but $m_i^{\theta(t),\theta^t} = (\cdot, z^i)$ with $z^i > 0$ for some i and θ^t. Then, by Claim 2, there exists $j \neq i$ such that $\pi_j^{\theta(t),\theta^t} < v_j^j$. Next consider j deviating to another strategy identical to σ_j at every history, except at $(\mathbf{h}(\theta(t)), \theta^t)$ where it announces the same state as σ_j but an integer higher than any integer that can be reported by σ at this history. Given ψ, such a deviation does not incur a one-period utility loss while strictly improving the continuation payoff as of the next period since, by Rule 3, the deviator j becomes a dictator himself and by Claim 2 $\pi_j^{\theta(t),\theta^t} < v_j^j$. This is a contradiction. ■

Given the previous two lemmas, we can now pin down the equilibrium payoffs by invoking efficiency *in the range*.

Lemma 4 *Suppose that f is efficient in the range and satisfies condition ω. Also, suppose that, for each i, outcome $\tilde{a}^i \in A$ used in the construction of S^i above satisfies (7). Then, for any $\sigma \in \Omega^\delta(R^*)$, $\pi_i^{\theta(t)}(\sigma, R^*) = v_i(f)$ for any i, $t > 1$ and $\theta(t)$.*

Proof. Suppose not; then f is efficient in the range 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 2, it must be that $\pi_i^{\theta(t)} > v_i(f)$. Also, by part (iii) of Lemma 3, $\left(\pi_j^{\theta(t)}\right)_{j\in I} \in co(V(f))$. Since f is efficient in the range, it then follows that there must exist some $j \neq i$ such that $\pi_j^{\theta(t)} < v_j(f)$. But, this contradicts Lemma 2. ■

It is straightforward to show that R^* has a Nash equilibrium in Markov strategies which attains truth-telling and, hence, the desired social choice at every possible history.

Lemma 5 *Suppose that f satisfies condition ω. There exists $\sigma^* \in \Omega^\delta(R^*)$, which is Markov, such that, for any t, $\theta(t)$ and θ^t, (i) $g^{\theta(t)}(\sigma^*, R^*) = g^*$; (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 θ. Thus, at any t and $\theta(t)$, we have $\pi_i^{\theta(t)}(\sigma^*, R^*) = v_i(f)$ for all i. Consider any i making a unilateral deviation from σ^* by choosing some $\sigma_i' \neq \sigma_i^*$ which announces a different message at some $(\theta(t), \theta^t)$. But, given ψ, 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 Rule 2, $\pi_i^{\theta(t),\theta^t}(\sigma_i', \sigma_{-i}^*, R^*) = v_i(f)$. Thus, the deviation is not profitable.[14] ∎

We are now ready to present our main results. The first result requires a slight strengthening of condition ω in order to ensure implementation of SCFs that are efficient *in the range.*

Condition ω'. For each i, there exists some $\tilde{a}^i \in A$ such that

(a) $v_i(f) \geq v_i(\tilde{a}^i)$ and

(b) if $v_i(f) = v_i(\tilde{a}^i)$ then either (i) there exists j such that $v_j^j > v_j(\tilde{a}^i)$ or (ii) the payoff profile $v(\tilde{a}^i)$ does not Pareto dominate $v(f)$.

Note that the additional requirement (b) in ω' that separates it from ω applies only in the special case when $v_i(f) = v_i(\tilde{a}^i)$. Even in such a non-generic case, (b) amounts to a minimal addition. For example, part (i) of (b) is satisfied if, instead of Assumption (A), at least three agents have distinct best outcomes in some state.

Theorem 2 *Suppose that $I \geq 3$, and consider an SCF f satisfying condition ω'. If f is efficient in the range, it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient in the range, it is repeated-implementable in Nash equilibrium from period 2.*

Proof. Consider any profile of outcomes $(\tilde{a}^1, \ldots, \tilde{a}^I)$ satisfying condition ω'. There are two cases to consider.

[14] In this Nash equilibrium, each agent is indifferent between the equilibrium and any unilateral deviation. The following modification to regime R^* will admit a *strict* Nash equilibrium with the same properties: for each i, construct S^i such that i obtains a payoff $v_i(f) - \epsilon$ for some arbitrarily small $\epsilon > 0$ (to do so, we will need the inequality in condition ω to be strict). This will, however, result in the equilibrium payoffs of our canonical regime to *approximate* the efficient target payoffs.

Case 1: For all i, $\tilde{a}^i \in A$ satisfies condition (7).

In this case the first part of the theorem follows immediately from Lemmas 4 and 5. To prove the second part, fix any $\sigma \in \Omega^\delta(R^*)$, i, $t > 1$ and $\theta(t)$. Then,

$$\pi_i^{\theta(t)} = \sum_{\theta^t \in \Theta} p(\theta^t) \left[(1-\delta) u_i(a^{\theta(t),\theta^t}, \theta^t) + \delta \pi_i^{\theta(t),\theta^t} \right]. \quad (8)$$

Also, by Lemma 4 we have $\pi_i^{\theta(t)} = v_i(f)$ and $\pi_i^{\theta(t),\theta^t} = v_i(f)$ for any θ^t. But then, by (8), we have $\sum_{\theta^t} p(\theta^t) u_i(a^{\theta(t),\theta^t}, \theta^t) = v_i(f)$. Since, by part (iii) of Lemma 3, $a^{\theta(t),\theta^t} \in f(\Theta)$, and since f is strictly efficient in the range, the claim follows.

Case 2: For some i, condition (7) does not hold.

In this case, $v_i(f) = v_i(\tilde{a}^i)$ and $v_j^j = v_j(\tilde{a}^i)$ for all $j \neq i$. Then, by b(ii) of condition ω', $v(\tilde{a}^i)$ does not Pareto dominate $v(f)$. Since $v_j^j \geq v_j(f)$, it must then be that $v_j(f) = v_j(\tilde{a}^i)$ for all j. Such an SCF can be trivially payoff-repeated-implemented via $\Phi^{\tilde{a}^i}$. Furthermore, since $v_j(f) = v_j(\tilde{a}^i) = v_j^j$ for all j, f is efficient. Thus, if f is strictly efficient (in the range), it must be constant, i.e. $f(\theta) = \tilde{a}^i$ for all θ, and hence can also be repeated-implemented via $\Phi^{\tilde{a}^i}$. ■

Note that when f is efficient (over the entire set of SCFs) part (b)(ii) of condition ω' is vacuously satisfied. Therefore, we can use condition ω instead of ω' to establish repeated implementation with efficiency.

Corollary 1 *Suppose that $I \geq 3$, and consider an SCF f satisfying condition ω. If f is efficient, it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient, it is repeated-implementable in Nash equilibrium from period 2.*

Note that Theorem 2 and its Corollary establish repeated implementation from the second period and, therefore, unwanted outcomes may still be implemented in the first period. This point will be discussed in more detail below.

Two agents As in one-shot Nash implementation (Moore and Repullo [26] and Dutta and Sen [10]), 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 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.

As identified by Dutta and Sen [10], a necessary condition for existence of an equilibrium in the one-shot setup is a *self-selection* requirement that ensures the availability of a punishment whenever the two players disagree on their announcements of the state but one of them is telling the truth. We show below that, with two agents, such a condition together with condition ω (ω'), delivers repeated implementation of an SCF that is efficient (efficient in the range). Formally, for any f, i and θ, let $L_i(\theta) = \{a \in A | u_i(a, \theta) \leq u_i(f(\theta), \theta)\}$ be the set of outcomes that are no better than f for agent i. We say that f satisfies *self-selection* if $L_1(\theta) \cap L_2(\theta') \neq \emptyset$ for any $\theta, \theta' \in \Theta$.[15]

Theorem 3 *Suppose that $I = 2$, and consider an SCF f satisfying condition ω (ω') and self-selection. If f is efficient (in the range), it is payoff-repeated-implementable in Nash equilibrium from period 2; if f is strictly efficient (in the range), it is repeated-implementable in Nash equilibrium from period 2.*

For the proof, which appear in the Supplementary Material (Section A), we construct a new regime $\widehat{R}$ that is identical to the canonical regime R^* with three or more agents, except that at any history the immediate outcome following announcement of different states is chosen according to the self-selection condition to support truth-telling in equilibrium. Formally, we replace mechanism g^* in the construction of R^* by a new mechanism $\hat{g} = (M, \psi)$ defined as follows: $M_i = \Theta \times \mathbb{Z}_+$ for all i and ψ is such that

1. if $m_1 = (\theta, \cdot)$ and $m_2 = (\theta, \cdot)$, then $\psi(m) = f(\theta)$;
2. if $m_1 = (\theta^1, \cdot)$ and $m_2 = (\theta^2, \cdot)$, and $\theta^1 \neq \theta^2$, then $\psi(m) \in L_1(\theta^2) \cap L_2(\theta^1)$ (by self-selection, this is well defined).

Thus, regime $\widehat{R}$ is such that $\widehat{R}(\emptyset) = \hat{g}$ and, for any $h = ((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})) \in H^t$ such that $t > 1$ and $g^{t-1} = \hat{g}$, the following transition rules hold:

Rule 1: If $m_1^{t-1} = (\cdot, 0)$ and $m_j^{t-1} = (\cdot, 0)$, then $\widehat{R}(h) = \hat{g}$.

Rule 2: If $m_i^{t-1} = (\cdot, z^i)$, $m_j^{t-1} = (\cdot, 0)$ and $z^i \neq 0$, then $\widehat{R}|h = S^i$ (Lemma 1).

Rule 3: If m^{t-1} is of any other type and i is lowest-indexed agent among those who announce the highest integer, then $\widehat{R}|h = D^i$.

[15] Self-selection is clearly weaker than the bad outcome condition in [26].

The replacement of g^* by $\hat{g}$ ensures that with two players the regime has a Nash equilibrium in which each player announces the true state and zero integer at every history. By self-selection, any unilateral deviation results in a current period outcome that is no better for the deviator; as with the three-or-more-agent construction, by making himself the "odd-one-out," the deviator obtains the same (target level) continuation payoff at the next period. Showing that every equilibrium of $\widehat{R}$ repeatedly implements the SCF from period 2 (in terms of payoffs or outcomes) proceeds analogously to the corresponding characterization for R^* with $I \geq 3$.

The purpose of self-selection here is to ensure existence of an equilibrium by appealing to one-shot incentives. In our repeated setup, there are alternative ways to obtain a similar result if the agents are sufficiently patient. For instance, we show in the Supplementary Material that with large enough δ the two requirements of self-selection and condition ω' in Theorem 3 above can be replaced by assuming an outcome $\tilde{a}$ that is strictly worse than f for both players *on average*, i.e. $v_i(\tilde{a}) < v_i(f)$ for all $i = 1, 2$.

4.3 Discussion

We next offer some discussion of our main results above.

More on condition ω In our analysis, repeated implementation of an efficient SCF has been obtained with an auxiliary condition ω (or its variant ω') which assumes that, for each agent, the payoff 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 i, a history-independent and non-strategic continuation regime S^i in which the agent derives a payoff equal to the target level $v_i(f)$.

While condition ω is satisfied in many applications, it is by no means necessary. An obvious alternative to construct such a regime S^i is to sequence dictatorship of i with dictatorship of another player j if j-dictatorship generates a unique payoff to i no greater than $v_i(f)$. Denoting the set of players whose dictatorships induce a unique payoff to i by $\Gamma_i = \{j \neq i \mid v_i^j = \sum_{\theta\in\Theta} p(\theta)u_i(a(\theta),\theta); \forall a(\theta) \in A^j(\theta), \forall\theta\}$, we can define another condition that can fulfill the same role as condition ω: an SCF f is *non-exclusive* if for each i, there exists some $j \in \Gamma_i$ such that $v_i^j \leq v_i(f)$.[16]

[16]The name of this 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

It is important to note that, with $I = 2$, the inequality part of non-exclusion is vacuously satisfied if the SCF is efficient and, therefore, we only need to assume that $\Gamma_i = \{j\}$ for each $i = 1, 2$. This is true, for instance, if $A^i(\theta)$ is a singleton set for all θ, i.e. each player's best response when dictator is always unique.[17]

More generally, constructing regime S^i could also be achieved with dictatorial mechanisms over restricted sets of outcomes. Specifically, for each agent j and any $N \subseteq A$, let $A^j(N,\theta) = \{a \in N \mid u_j(a,\theta) \geq u_j(a',\theta)\ \forall a' \in N\}$ be the outcomes that j would choose from the restricted outcome set N in state θ when he is dictator, and let $v_i^j(N) = \sum_{\theta\in\Theta} p(\theta) \max_{a\in A^j(N,\theta)} u_i(a,\theta)$ be i's maximum payoff from j-dictatorship over N, with $v^j(N)$ denoting the corresponding payoff profile. Also, define $\Gamma_i(N)$ as the set of all agents other than i such that i has a unique payoff from their dictatorships over N.[18] Then, for each i, S^i can be constructed if there exist a set N and a player $j \in \Gamma_i(N)$ such that $v_i^j(N) \leq v_i(f)$. Note that both condition ω and non-exclusion are equivalent to the above condition when N is a singleton set or the entire set A, respectively. Thus, for repeated-implementing efficient SCFs the two conditions can be subsumed by the following: for each i, there exists some $v = (v_1, \ldots, v_I) \in \{v^j(N)\}_{j\in\Gamma_i(N), N\in 2^A}$ such that $v_i \leq v_i(f)$.

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 (e.g. Moore and Repullo [25] and Abreu and Sen [3]). Although this paper also considers a dynamic setup, the solution concept adopted is that of Nash equilibrium and our characterization results do not rely on imposing off-the-equilibrium credibility to eliminate unwanted equilibria.[19] At the same time, our existence results do not involve construction of Nash equilibria based on non-credible threats off-the-equilibrium. 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 extensive

any other agent to the SCF itself (i.e. "excluded" by the SCF).

[17]While non-exclusion can replace condition ω for repeated implementation of efficient SCFs, we still need an extra condition similar to (7) to achieve the results with efficiency in the range. Specifically, for each i, if $v_i(f) = v_i^j$ for all $j \in \Gamma_i$ then there must be some $k \neq i$ and some $j \in \Gamma_i$ such that $v_k^k > v_k^j$.

[18]Formally, $\Gamma_i(N) = \Big\{ j \neq i \mid v_i^j(N) = \sum_{\theta\in\Theta} p(\theta) u_i(a(\theta),\theta); \forall a(\theta) \in A^j(N,\theta), \forall\theta \Big\}$.

[19]In particular, note that we do not require each player i to behave rationally when he is dictator at some off-the-equilibrium history. Lemmas 2-3 only appeal to the *possibility* that dictatorial payoffs could be obtained by the deviator.

form 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 implementability (e.g. Maskin and Moore [22]). For many of our repeated implementation results, this needs not be a cause for concern since off-the-equilibrium outcomes in our regimes can actually be made efficient. If the environment is rich enough, the outcomes needed for condition ω could be found on the efficient frontier itself. Moreover, if the SCF is non-exclusive, the regimes can also be constructed so that off-the-equilibrium is entirely associated with dictatorships, which are efficient.

Period 1 The critical aspect of our constructions behind Theorems 2-3 is that if any player expects a payoff below his target level from the continuation play then this player could deviate in the previous period and make himself the "odd-one-out." This argument ensures that from period 2 desired outcomes are implemented. Our results however do not guarantee period 1 implementation of the SCF; in fact one can easily find an equilibrium of regime R^* or $\widehat{R}$ where the players report false states and integer zero in period 1 (at every other history they follow truth-telling and announce zero). If the SCF satisfies the standard conditions required for one-shot implementation, nonetheless, our constructions can be altered to achieve period 1 implementation. For example, with monotonicity and no veto power we could just modify mechanism for period 1 as in Maskin [21].

We could also deal with period 1 implementation if there were a pre-play round that takes place before the first state is realized. In such a case, prior to playing the canonical regime one could let the players simply announce a non-negative integer with the same transition rules such that equilibrium payoffs at the beginning of the game correspond exactly to the target levels.

Alternatively, we could consider an equilibrium refinement. In the Supplementary Material (Section B), we formally introduce agents who possess, at least at the margin, a preference for simpler strategies, in a similar way that complexity-based equilibrium refinements have yielded sharper predictions in various dynamic game settings (e.g. Abreu and Rubinstein [2], Chatterjee and Sabourian [8], Gale and Sabourian [12]). By adopting a natural measure of complexity and a refinement based on very mild criteria in terms of complexity, we show that every equilibrium in the canonical regimes above must be Markov and hence the main sufficiency results extend to implementation from outset. Similar refinements will later be used to analyze constructions employing only finite mechanisms; we refer the reader to Section 5.2 below.

Social choice correspondence Our analysis could be extended to repeated implementation of a social choice *correspondence* (SCC) as follows. For any mapping $\mathbb{F} : \Theta \to 2^A \backslash \{\emptyset\}$, let $F(\mathbb{F}) = \{f \in F : f(\theta) \in \mathbb{F}(\theta)\ \forall\theta\}$. Then an SCC $\mathbb{F}$ is repeated-implementable if we can find a a regime such that for any $f \in F(\mathbb{F})$ there exists a Nash equilibrium that repeated-implements it, in the sense of Definition 3, and every Nash equilibrium repeated-implements some $f \in F(\mathbb{F})$. With this definition, it is trivially the case that our necessary condition for repeated implementation in Theorem 1 also holds for each $f \in F(\mathbb{F})$.

We can also obtain an equivalent set of sufficiency results to Theorems 2-3 (and Corollary 1) for repeated-implementing $\mathbb{F}$ by modifying the canonical regime as follows. In period 1, each agent first announces an SCF from the set $F(\mathbb{F})$; if all announce the same SCF, say, f, then they play the canonical regime, R^* when $I \geq 3$ or $\widehat{R}$ when $I = 2$, defined for f, while otherwise they play the canonical regime that corresponds to some arbitrary $\tilde{f} \in F(\mathbb{F})$. If every $f \in F(\mathbb{F})$ satisfies efficiency and the other auxiliary conditions, such a regime would repeated-implement $\mathbb{F}$. Thus, when indifferent among several (efficient) SCFs, the planner can let the agents themselves choose a particular SCF and payoff profile in the first period.

Learning by the planner In a dynamic environment, one may ask what would happen if the planner could also observe the state at the end of a period with some probability, say, ϵ. Depending on the interpretation of the state, this could be an important issue. While our sufficiency results clearly remain true, the necessity result is robust to such learning by the planner in the following sense. Suppose that an SCF f is repeated-implementable but strictly dominated by another SCF (in its range). Then, for sufficiently small values of ϵ, the regime must admit another equilibrium in which the agents collude to achieve the superior payoffs by similar arguments to those behind Theorem 1 above.

5 Mixed Strategies and Finite Mechanisms

In this section, we broaden the scope of our sufficiency results by considering mixed strategies as well as constructions involving only finite mechanisms.

5.1 Mixed strategies

In our analysis thus far, repeated implementation of an efficient SCF has been obtained under restriction to pure strategies. In the static Nash implementation literature, it is well known that the canonical mechanism can be modified to deal with mixed strategies (Maskin [21]). The unbounded nature of the integer game ensures that there cannot be an equilibrium in pure or mixed strategies in which positive integers are announced.

It is similarly possible to incorporate mixed (behavioral) strategies into our repeated implementation setup. In the Supplementary Material (Section C), we establish an analogous sufficiency result to the results of Section 4.2 for the case of $I \geq 3$ (the two-agent case can be dealt with similarly and hence omitted). Specifically, we show that an SCF that satisfies efficiency (strict efficiency) and condition ω can be payoff-repeated-implemented (repeated-implemented) in pure or mixed strategy Nash equilibrium from period 2.[20]

We obtain these results with the same canonical regime R^*. With mixed strategies, each player i faces uncertainty about the others' messages and, therefore, the "odd-one-out" argument first obtains a lower bound for each player's *expected* continuation payoffs at each history (in contrast to Lemma 2). If the SCF is efficient these expected continuation payoffs are equal to the target levels. Given this, integer arguments can be extended to show that, whether playing pure or mixed strategies, the agents must always announce zero at every history and hence mechanism g^* must always be played. Although the players may still mix over their reports on state, we can then once again apply the previous arguments to reach the results.

5.2 Finite mechanisms

Our sufficiency results appeal to integer games to determine the continuation play at each history. In the one-shot implementation literature, integer-type arguments have been at times criticized for its lack of realism or for technical reasons (e.g. being unbounded or not having undominated best responses). Such criticisms may also be applied to our constructions. One response, both in static and our repeated setups, is that integers are used to demonstrate what can possibly be implemented in most general environments; in specific examples more appealing constructions may also work. Furthermore, given Theorem 1, our sufficiency results show that indeed efficiency is a relatively tight (necessary)

[20]With mixed strategies, our necessity result (Theorem 1) holds via identical arguments.

condition for repeated implementation.

Another response in the static implementation literature to the criticism of integer games has been to restrict attention to *finite* mechanisms, such as the modulo game. Using a finite mechanism to achieve Nash implementation, however, brings an important drawback: unwanted *mixed* strategy equilibria. This could be particularly problematic in one-shot settings because as Jackson [13] has shown a finite mechanism Nash-implementing an SCF could invite unwanted mixed equilibria that strictly Pareto dominate the SCF.[21]

If we exclude mixed strategies, it is also straightforward to replace the integer games in our repeated game constructions with a finite alternative like the modulo game. More challenging is the issue of unwanted mixed strategy equilibria in a regime that employs only finite mechanisms. Regarding this issue, note that we are implementing an efficient SCF and, hence, there cannot be another mixed equilibrium that dominates it. In fact, below we go further and construct a regime with finite mechanisms (involving at most three integers) that, under minor qualifications, possesses the following two features.

First, every non-pure Nash equilibrium of the regime is strictly Pareto dominated by the pure equilibria which obtain implementation of the efficient SCF. Thus, we turn Jackson's criticism of one-shot Nash implementation into our favor: non-pure equilibria in our repeated settings are less plausible from the same efficiency perspective.

Second, and more importantly, we can eliminate randomization altogether by considering Nash equilibrium strategies that are credible (subgame perfect) and by invoking an additional equilibrium refinement, based on complexity considerations, that is appealing and very marginal.

The basic idea that we introduce to obtain these twin findings is that, even with simple finite mechanisms, the freedom to choose different mechanisms at different histories enables the planner to design a regime with the following property: if the players were to randomize in equilibrium, the strategies would prescribe (i) inefficient outcomes and (ii) a complex pattern of behavior (i.e. choosing different mixing probabilities at different histories) that could not be justified by payoff considerations, as simpler strategies could induce the same payoff as the equilibrium strategy at every history.

To save space, here we present the formal analysis only for the two-agent case. The

[21] In order to address mixed strategies with finite mechanisms, the static implementation literature has explored the role of refinements and/or virtual implementation in specific environments (e.g. Jackson, Palfrey and Srivastava [15], Sjöström [30] and Abreu and Matsushima [1]).

analysis for the case of $I \geq 3$ is more lengthy and complicated. We offer a brief remark on this case at the end of this section and leave the details to the Supplementary Material. Unless otherwise mentioned, all formal proofs in this section appear in the Appendix.

Construction Suppose that $I = 2$, and fix any SCF f that satisfies efficiency and self-selection, as required for Theorem 3 above. To obtain our results, we modify the two-agent canonical regime $\widehat{R}$ in Section 4.2 as follows. First, we replace mechanism $\hat{g}$ in regime $\widehat{R}$ by the following *extensive form* mechanism, referred to as g^e:

> Stage 1 - Each agent $i = 1, 2$ announces a state, θ^i, from Θ. If $\theta^1 = \theta^2 = \theta$, $f(\theta)$ is implemented; otherwise, an outcome from the set $L_1(\theta^2) \cap L_2(\theta^1)$ is implemented.
>
> Stage 2 - Each agent announces an integer, z^i, from the set $\mathcal{Z} \equiv \{0, 1, 2\}$.

In this mechanism, the outcome implemented as a function of the states announced is essentially the same as that of mechanism $\hat{g}$. However, it differs from $\hat{g}$ in that the agents only choose integers 0,1 or 2, and also, the mechanism has a two-stage sequential structure. The latter change, as will be clarified shortly, enables us to define the notion of complexity of a strategy in a natural way.

Our second modification to $\widehat{R}$ involves changing the continuation regimes after a non-zero integer announcement. As we argued in Section 4.3, with two agents, S^i can be constructed by alternating dictatorships of the two players, as long as the payoff profile from each dictatorship is unique (e.g. each dictator has a unique best response in each state). For the rest of this section, we assume that this is indeed the case. Thus, for each $i = 1, 2$, there exists a regime S^i that yields a unique (discounted average) payoff profile $w^i = (w^i_i, w^i_j)$ such that $w^i_i = v_i(f)$. Furthermore, since f is efficient it must be that, for $j \neq i$, $w^i_j \leq w^j_j$. If the latter inequality binds then f can be payoff-implemented from period 1 by simply adopting S^i.[22] Thus, with (almost) no loss of generality, assume that

$$w^i_j < w^j_j \text{ for every } j \neq i. \tag{9}$$

Given (9), there must then exist regimes $X(t)$ for each $t = 1, 2, \ldots$ and Y that respectively induce unique payoff profiles $x(t)$ and y satisfying the following condition:

$$w^2_1 < y_1 < x_1(t) < w^1_1 \quad \text{and} \quad w^1_2 < x_2(t) < y_2 < w^2_2. \tag{10}$$

[22]Note that, in this case, we can invoke Fudenberg and Maskin [11] and construct S^i in such a way that the continuation payoffs at any date is also arbitrarily close to $v(f)$.

To construct these regimes, let $x(t) = \lambda(t)w^1 + (1-\lambda(t))w^2$ and $y = \mu w^1 + (1-\mu)w^2$ for some $0 < \mu < \lambda(t) < 1$. By (9), these payoffs satisfy (10). Furthermore, since w^i for each i is a convex combination of the two dictatorial payoffs v^1 and v^2, such payoffs can be obtained by regimes that appropriately alternate between the two dictatorships.

Now, we define new regime R^e inductively as follows: (i) mechanism g^e is implemented at $t = 1$ and (ii) if, at some date t, g^e is the mechanism played with a pair of states $\underset{\sim}{\theta}$ and a pair of integers $\underset{\sim}{z} = (z^1, z^2)$ announced over the two stages, the continuation mechanism/regime at the next period is as follows:

Rule 1: If $z^1 = z^2 = 0$, the mechanism next period is g^e.

Rule 2: If $z^1 > 0$ and $z^2 = 0$ ($z^1 = 0$, $z^2 > 0$), the continuation regime is S^1 (S^2).

Rule 3: Suppose that $z^1, z^2 > 0$. Then, we have the following.

Rule 3.1: If $z^1 = z^2 = 1$, the continuation regime is $X \equiv X(\tilde{t})$ for some arbitrary $\tilde{t}$, with the payoffs denoted by x.

Rule 3.2: If $z^1 = z^2 = 2$, the continuation regime is $X(t)$.

Rule 3.3: If $z^1 \neq z^2$, the continuation regime is Y.

As in $\widehat{R}$, announcement of any non-zero integer in regime R^e ends the game, and if only one agent announces zero, the other agent obtains his target payoff in the continuation regime. The rest of transitions are designed to achieve our new objectives. In particular, when both agents report 2 (Rule 3.2) the continuation regimes could actually be different across periods. This feature will later be used to facilitate our refinement arguments.

Next we define strategies in R^e. In this regime the histories that matter are only those at which the agents engage in mechanism g^e. Using the same notation as before, we denote by $\mathbf{H}^t$ the set of all such finite histories observed by the agents at the beginning of period t; let $\mathbf{H}^\infty = \cup_{t=1}^{\infty} \mathbf{H}^t$. Also, since g^e has a two-stage sequential structure, we need to additionally describe information available within a period. We refer to such information as *partial history* and denote it by $d \in \Theta \cup (\Theta \times \Theta^2) \equiv D$; thus $d = \theta$ represents the beginning of stage 1 of g^e after state θ has been realized and $d = (\theta, \underset{\sim}{\theta})$ refers to the beginning of stage 2 after realization of θ followed by profile $\underset{\sim}{\theta}$ announced in stage 1.

Then, using Δ before a set to denote the probability distributions over the set, we represent a mixed (behavioral) strategy of agent $i = 1, 2$ in regime R^e as the mapping

$b_i : \mathbf{H}^\infty \times D \rightarrow (\triangle\Theta) \cup (\triangle\mathcal{Z})$ such that, for any $\mathbf{h} \in \mathbf{H}^\infty$, $b_i(\mathbf{h}, d) \in \triangle\Theta$ if $d \in \Theta$ and $b_i(\mathbf{h}, d) \in \triangle\mathcal{Z}$ if $d \in \Theta \times \Theta^2$. Let B_i be the set of i's strategies in R^e. With slight abuse of notation, we write $\pi_i^{\mathbf{h}}(b, R^e)$ as player i's continuation payoff at history $\mathbf{h} \in \mathbf{H}^\infty$ under strategy profile b.

Properties of Nash equilibria We begin by obtaining an important property of Nash equilibrium of the above regime. At any information set on the equilibrium path where the players face stage 2 (the integer part) of mechanism g^e, they must all be either playing 0 for sure and obtaining the target payoffs $v(f)$ in the continuation game next period, or mixing between 1 and 2 for sure and obtaining less than $v(f)$.

Lemma 6 *Consider any Nash equilibrium of regime R^e. Fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d = (\theta, \underset{\sim}{\theta}) \in \Theta \times \Theta^2$ on the equilibrium path. Then, one of the following must hold at $(\mathbf{h}, d)$:*

1. Each i announces 0 for sure and his continuation payoff at the next period is $v_i(f)$.

2. Each i announces 1 or 2 for sure, with the probability of choosing 1 equal to $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i} \in (0,1)$, and his continuation payoff at the next period is less than $v_i(f)$.

Let us sketch the steps of the proof. First, assume no randomization over integers at the relevant history. Then, the inequalities of (10) ensure that in equilibrium the agents must announce zero and hence, by reasoning similar to the "odd-one-out" argument, the continuation payoff of each i is $v_i(f)$.

Second, we show that if the players are mixing over integers then zero cannot be chosen. Since $x_i(t) > w_i^j$ and $y_i > w_i^j$ for $i, j = 1, 2$, the transition rules imply that each agent prefers to announce 1 than to announce 0 if the other player is announcing a positive integer for sure. It then follows that if agent i attaches a positive weight to 0 then the other agent j must also do the same, and i's continuation payoff is at least (strictly larger than) $v_i(f)$ when j announces 0 for sure (attaches a positive weight to a positive integer). Applying this argument to both agents leads to a contradiction against the assumption that the SCF is efficient.

Finally, i's continuation payoff at the next period when both choose a positive integer is x_i, $x_i(t)$ or y. The precise probability of choosing integer 1 by i in the case of mixing is determined trivially by these payoffs as in the lemma. Also, since these payoffs are all by assumption less than $v_i(f)$, we have that mixing results in continuation payoffs strictly below the target levels.

Given Lemma 6 we can also show that if the players mix over integers at any on-the-equilibrium history it must be in period 1; otherwise, both players must be playing 0 in the previous period where either player i could profitably deviate by announcing a positive integer and activating continuation regime S^i. Combining this with Lemma 6, we state the following.

Proposition 1 *Fix any Nash equilibrium b of regime R^e.*

1. If any player mixes over integers on the equilibrium path, then both players randomize at some partial history in stage 2 of period 1; furthermore, $\pi_i^{\mathbf{h}}(b, R^e) \leq v_i(f)$ for all i and any on-the-equilibrium history $\mathbf{h} \in \mathbf{H}^2$ with the inequality being strict at every such history that involves randomization in stage 2 of period 1.

2. Otherwise, $\pi_i^{\mathbf{h}}(b, R^e) = v_i(f)$ for any i, any $t > 1$ and any (on-the-equilibrium) history $\mathbf{h} \in \mathbf{H}^t$.

Note also that R^e trivially admits a Nash equilibrium in which each agent always announces the true state and integer zero. This, together with part 2 of Proposition 1, implies that R^e payoff-repeated-implements f from period 2 if players do not randomize over integers. Part 1 demonstrates that such randomization results in an expected continuation payoff from period 2 that is less than the target payoff for every player. This however does not rule out the possibility that in period 1 some player obtains a one-period expected payoff greater than the target level. With sufficiently large δ one-period payoffs have small weights and therefore we can claim that any such randomized Nash equilibrium is dominated by equilibria that do not involve the randomization.

Refinement We now introduce our refinement arguments. Note first that, if we apply subgame perfection, the statements of Lemma 6 above can be readily extended to hold for any on- or *off*-the-equilibrium history after which the agents find themselves in the integer part of mechanism g^e; that is, in a subgame perfect equilibrium (SPE) of regime R^e, at *any* $(\mathbf{h}, (\theta, \underline{\theta}))$ they must either choose 0 for sure or mix between 1 and 2.

Next, we add to the construction of R^e the following property: the sequence of regimes $\{X(t)\}_t$ is such that, in addition to (10) above, the corresponding payoffs $\{x(t)\}_t$ satisfy

$$x_1(t) \neq x_1(t') \text{ and } x_2(t) \neq x_2(t') \text{ for any } t, t', t \neq t'. \tag{11}$$

Note that this can be done simply by ensuring that the sequence $\{\lambda(t) : \lambda(t) \in (\mu, 1) \ \forall (t)\}$ used before to construct these regimes is such that $\lambda(t) \neq \lambda(t')$ for any $t \neq t'$.

Clearly, this additional feature does not alter Lemma 6, or its extension to SPE. However, it implies for any SPE that, if the agents mix over integers at some period t and partial history within t on or off the equilibrium path, each i's mixing probability, given by $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i}$, is determined *uniquely* by t.

We next introduce a "small" cost associated with implementing a more complex strategy. Complexity of a strategy can be measured in a number of ways. For our analysis, it is sufficient to have a notion of complexity that captures the idea that stationary behavior (always making the same choice) at every stage in mechanism g^e is simpler than taking different actions in g^e at different histories. We adopt the following.

Definition 4 *For any i and any pair of strategies $b_i, b_i' \in B_i$, we say that b_i is more complex than b_i' if the strategies are identical everywhere except, after some partial history in mechanism g^e, b_i' always behaves (randomizes) the same way while b_i does not. Formally, there exists some $d' \in D \equiv \Theta \cup (\Theta \times \Theta^2)$ with the following properties:*

1. $b_i'(\mathbf{h}, d) = b_i(\mathbf{h}, d)$ for all $\mathbf{h} \in \mathbf{H}^\infty$ and all $d \in D$, $d \neq d'$.

2. $b_i'(\mathbf{h}, d') = b_i'(\mathbf{h}', d')$ for all $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$.

3. $b_i(\mathbf{h}, d') \neq b_i(\mathbf{h}', d')$ for some $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$.

Notice that this definition imposes a very weak and intuitive partial order over the strategies. It has a similar flavor to the complexity notions used by Chatterjee and Sabourian [8], Sabourian [28] and Gale and Sabourian [12] who consider bargaining and market models. Our results also hold with other similar complexity measures, which we discuss in further detail at the end of this section.[23]

Using Definition 4, we refine the set of SPEs as follows.[24]

Definition 5 *A strategy profile b is a weak perfect equilibrium with complexity cost (WPEC) of regime R^e if b is an SPE and for each i no other strategy $b_i' \in B_i$ is such that*

1. b_i' is less complex than b_i; and

2. b_i' is a best response to b_{-i} at every information set for i (on or off the equilibrium).

[23] We could also adopt measures such as counting the number of "states of the automaton" implementing the strategy (e.g. Abreu and Rubinstein [2]) or the "collapsing state condition" (Binmore, Piccione and Samuelson [5]).

[24] In the Supplementary Material (Section D), we discuss another equivalent way of introducing credibility and complexity cost into the equilibrium notion by explicitly considering the possibility of errors (off-the-equilibrium moves).

WPEC is a very mild refinement of SPE since it requires players to adopt minimally complex strategies among the set of strategies that are best responses *at every information set*. This means that complexity appears lexicographically after *both* equilibrium and off-equilibrium payoffs in each player's preferences. This contrasts with the more standard equilibrium notion in the literature on complexity in dynamic games that requires strategies to be minimally complex among those that are best responses only on the equilibrium path.[25] This latter approach, however, has been criticized for prioritizing complexity costs ahead of off-equilibrium payoffs in preferences. Our notion of WPEC avoids this issue since it only excludes strategies that are unnecessarily complex without any payoff benefit on or off the equilibrium.

Note that strategies that always report the true state and zero continue to be an equilibrium with this notion. We now show that there cannot be a WPEC in which mixing over integers occurs.

Lemma 7 *Fix any WPEC of regime R^e. Also, fix any $\mathbf{h} \in H^\infty$ and $d \in \Theta \times \Theta^2$ (on or off the equilibrium path). Then, every agent announces zero for sure at this history.*

To obtain this lemma we show that, otherwise, either agent could deviate to another less complex strategy identical to the equilibrium strategy everywhere except that it always responds to partial history $d = (\theta, \underset{\sim}{\theta})$ by announcing 1, and obtain the same payoff at every history. Three crucial features of our regime construction deliver this argument. First, the deviation is less complex because the mixing probabilities are uniquely determined by the date t and, hence, the equilibrium strategy must prescribe different behaviors at different histories. Second, since the players can only randomize between 1 and 2, the deviation would not affect payoffs at histories where the equilibrium strategies randomize. Finally, since at histories where the equilibrium strategies do not mix they report 0 for sure with continuation payoffs equal to $v(f)$, by reporting 1 the deviator becomes the "odd-one-out" and ensures the same target payoff.

Recall that for any SPE in which no randomization over integers occurs the continuation payoffs at any history beyond period 1 are equal to $v(f)$. Then, by Lemma 7, we can immediately establish the following.

[25]The two exceptions in the existing literature are Kalai and Neme [18] and Sabourian [28]. The notion of WPEC was first introduced by [28].

Proposition 2 *1. If f is efficient, every WPEC, b, of regime R^e payoff-repeated-implements f from period 2 at every history; i.e.* $\pi_i^{\mathbf{h}}(b, R^e) = v_i(f)$ *for all i, $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t$.*

2. If f is strictly efficient, every WPEC, b, of regime R^e repeated-implements f from period 2 at every history.

The complexity of our regime construction driving these WPEC results is captured by the non-stationarity of continuation regimes $\{X(t)\}$. This feature generates Nash equilibrium strategies involving randomization to follow a complex pattern. The arguments behind Lemma 7 would in fact hold as long as these continuation regimes differ in just two periods. In general, the planner could also write the continuation regime $X(\cdot)$ as a function of the entire (publicly observable) history instead of just its date, thereby reinforcing the complexity of such mixing behavior and making our arguments more compelling.

One may, however, question why we consider a preference for less complex behavior only by the agents and not by the planner. We note here that our complexity notion only calls for any additional complexity of a strategy to be justified by payoffs. Similarly, for the planner the complexity of the above regime is warranted in the sense that it allows for better implementation results.

Alternative complexity measure The basic idea behind our complexity measure is that stationary behavior is simple. Definition 4 captures this by saying that a strategy that at every date t responds identically to some partial history d, independently of the previous history of play before t, is less complex than one that responds differently to the same partial history d. Another measure with a similar flavor would be to say that a strategy that announces the same integer (or state) regardless of *both* the history before the date and the partial history within the date is less complex than one that announces different integers (or states) while being identical everywhere else. In the Supplementary Material we formally explain how our results can be extended to such a partial order.[26]

Three or more agents In the Supplementary Materia (Section D), we extend the above

[26] With this alternative complexity measure, our characterization of WPECs of regime R^e above remains true via exactly the same reasoning. Also, when the self-selection condition is strict (so that there exists an outcome strictly worse than what the SCF prescribes in each state), R^e admits a truth-telling equilibrium that is robust to this refinement. When strict self-selection is not satisfied, however, truth-telling will not constitute an equilibrium of R^e with this alternative definition of complexity, as the players may be able to economize on complexity cost of reporting different (true) states without affecting payoffs. In the Supplementary Material (Section D) we show that this case can be handled by slightly modifying R^e.

analysis to the case of $I \geq 3$. Two points are worth making. First, analogous results are obtained if, as in the two-agent case, we can construct for each agent i a history-independent and non-strategic regime S^i that yields a unique payoff profile $(w_i^i, \ldots, w_I^i)$ such that $w_i^i = v_i(f)$ and, as in (9), $w_j^i < w_j^j$ for all $j \neq i$. While this condition is innocuous when $I = 2$ (as we only needed to assume that dictatorial payoffs are unique), it is less so when $I \geq 3$.

Second, in the regime constructed to handle the case of $I \geq 3$ the size of the mechanim does not increase with the number of players (involves at most three integers). This contrasts with standard finite constructions like modulo game where the number of integers needed is at least the number of players. Specifically, we extend R^e as follows: only two agents are given the option to choose from $\{0, 1, 2\}$ in the integer part of mechanism g^e while all remaining agents choose from $\{0, 1\}$. The transition rules are set in a way that integer reports of the first two players are given priority over the rest and, therefore, the key qualitative features of the two-agent regime also hold with more than two players.

6 Conclusion

This paper sets up a problem of infinitely repeated implementation with stochastic preferences and establishes that, with minor qualifications, a social choice function is repeated-implementable in Nash equilibrium in complete information environments if and only if it is efficient (in the range). We also argue that our results are robust to various refinements and extend them to incorporate mixed strategies and finite mechanisms.

Our findings contrast with those obtained in the literature on static Nash implementation in which monotonicty occupies a critical position. The reason for this fundanmental difference is that in our repeated implementation setup the agents learn the infinite sequence of states gradually rather than all at once.[27]

In the one-shot implementation problem with *in*complete information, full implementation requires incentive compatibility in addition to Bayesian Monotonicity (an extension of Maskin monotonicity). The main arguments developed in this paper can be extended to show that neither is necessary for repeated implementation. A companion paper (Lee and Sabourian [19]) establishes the following results.

[27] If the agents learned the states at once and the SCF were a mapping from the set of such sequences Θ^∞ to the set of infinite outcomes A^∞ the problem would be analogous to one-shot implementation.

First, in a general incomplete information setup, we show that an SCF satisfying efficiency and incentive compatibility can be repeated-implemented in Bayesian Nash equilibrium. In a regime similar to the canonical regimes in this paper, efficiency pins down continuation payoffs of every equilibrium; incentive compatibility ensures existence.[28] Second, restricting attention to the case of *interdependent values*, repeated implementation of an efficient SCF is obtained when the agents are sufficiently patient by replacing incentive compatibility with an intuitive condition that we call *identifiability*. This condition stipulates that a unilateral deviation from truth-telling can be detected by another player after the outcome is implemented in the period. Given this, 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.

There are several important questions still outstanding. In particular, it remains to be seen whether efficiency is also necessary in incomplete information settings. The sufficiency results in [19] also assume either incentive compatibility or identifiability in the case of interdependent values, and leaves open the issue of how important these assumptions are in general.

Another interesting direction for future research is to generalize the process with which individual preferences evolve. However, allowing for such non-stationarity makes it difficult to define efficiency of social choices. Also, this extension will introduce the additional issue of learning.

Appendix

Proof of Lemma 6 Fix any $\mathbf{h} \in \mathbf{H}^\infty$ and $d \in \Theta \times \Theta^2$. For any $i = 1, 2$, let Π_i denote i's continuation payoff at the next period if both agents announce zero. Also, let z^i denote the integer that i ends up choosing at $(\mathbf{h}, d)$. At this history the players either randomize (over integers) or do not randomize. We consider each case separately.

<u>*Case 1*</u>: No player randomizes.

[28]With incomplete information, we evaluate repeated implementation in terms of *expected* continuation payoffs computed at the beginning of a regime. This is because continuation payoffs in general depend on an agent's ex post beliefs about the others' past private information at different histories but we do not want our solution concept to depend on such beliefs.

In this case we show that each player must play 0 for sure. Suppose otherwise; then some i plays $z^i \neq 0$ for sure and the other announces z^j for sure. We derive contradiction by considering the following subcases.

Subcase 1A: $z^i > 0$ and $z^j = 0$.

The continuation regime at the next period is S^i (Rule 2). But then, since $y_j > w_j^i$ by construction, j can profitably deviate by choosing a strategy identical to the equilibrium strategy except that it announces the positive integer other than z^i at this history, which activates the continuation regime Y instead of S^i (Rule 3.3). This is a contradiction.

Subcase 1B: $z^i > 0$ and $z^j > 0$.

The continuation regime is either X, $X(t)$ or Y (Rule 3). Since $y_2 > x_2(t)$ for any t, it follows that if the continuation regime is X or $X(t)$ then player 2 can profitably deviate just as in Subcase 1A, a contradiction. Since $x_1 > y_1$, if the continuation regime is Y player 1 can profitably deviate and we obtain a similar contradiction.

Given that both players choose 0 for sure, the players then face mechanism g^e at the next period. Therefore, we can apply the "odd-one-out" argument of Lemma 2 and efficiency of f to show that the continuation payoffs must equal $v(f)$.

Case 2: Some player randomizes.

We proceed by first establishing the following two claims.

Claim 1: For each i, the continuation payoff from announcing 1 is greater than that from announcing 0, if $z^j > 0$ for sure, $j \neq i$.

Proof of Claim 1. If i announces zero, by Rule 2, his continuation payoff is w_i^j. If he announces 1, by Rules 3.1 and 3.3, the continuation payoff is $x_i > w_i^j$ or $y_i > w_i^j$.

Claim 2: Suppose that agent i announces 0 with a positive probability. Then the other agent j must also announce 0 with a positive probability and $\Pi_i \geq v_i(f)$. Furthermore, $\Pi_i > v_i(f)$ if j does not choose 0 for sure.

Proof of Claim 2. By Claim 1, playing 1 must always yield a higher continuation payoff for player i than playing 0, except when j plays 0. Since i plays 0 with a positive probability, it must then be that (i) j chooses 0 with a positive probability and (ii) if j also attaches a positive weight to a positive integer, i's continuation payoff is greater from choosing 0 than from choosing 1, i.e. $\Pi_i > v_i(f)$. Finally, if j chooses 0 for sure then i obtains Π_i from 0 and $v_i(f)$ from 1 or 2; hence in this case we must have $\Pi_i \geq v_i(f)$.

We now show that, in this Case 2, both players choose a postive integer for sure. To show this suppose otherwise; then some player chooses 0 with a positive probability. By

Claim 2, the other player must also play 0 with a positive probability and, also, $\Pi_i \geq v_i(f)$ for any $i = 1, 2$. Moreover, since this case assumes that some player is choosing 0 with a probability less than one, by appealing to Claim 2 once again, it must be that at least one of the inequalities $\Pi_1 \geq v_1(f)$ or $\Pi_2 \geq v_2(f)$ is strict. But, since f is efficient, this is a contradiction.

In this case, therefore, both players mix between 1 and 2 for sure and, by simple computation, it must be that each i plays 1 with probability $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i} \in (0,1)$. Furthermore, since for each i, $v_i(f)$ exceeds $x_i, x_i(t)$ or y, it follows that the continuation payoff at the next period must be less than $v_i(f)$.

Proof of Proposition 1 Given Lemma 6, it suffices to show that if the players mix over integers on the equilibrium path it must happen in period 1. Suppose not; so, there exists a Nash equilibrium b such that, for some $t > 1$, there exist $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in \Theta \times \Theta^2$ that occur on the equilibrium path at which the players are mixing.

First, note that by Lemma 6 the players must have all announced 0 for sure in the previous period. Thus, we can deduce that $\pi_i^{\mathbf{h}^t}(b, R^e) = v_i(f)$ for all $i = 1, 2$ by invoking the same arguments as in Lemma 2 and efficiency of f. Second, for any (on-the-equilibrium) partial history $d' \in \Theta \times \Theta^2$, we can also apply similar reasoning to show that $\pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}}(b, R^e) = v_i(f)$ for all i if $\underset{\sim}{z} = (0, 0)$.

Next, let $r(d', \underset{\sim}{z})$ denote the probability of $(d', \underset{\sim}{z})$ occurring at $\mathbf{h}^t$ under b, and let $a^{\mathbf{h}^t, d'}$ denote the outcome implemented at $(\mathbf{h}^t, d')$. Then, with slight abuse of notation, i's continuation payoff at $\mathbf{h}^t$ can be written as

$$\pi_i^{\mathbf{h}^t}(b, R^e) = \sum_{(d', \underset{\sim}{z}) \in \Theta \times \Theta^2 \times \mathcal{Z}} r(d', \underset{\sim}{z}) \left[(1-\delta) u_i(a^{\mathbf{h}^t, d'}, d') + \delta \pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}} \right] = v_i(f). \qquad (12)$$

Lemma 6 implies that, for any i and any d', it must be either that $\underset{\sim}{z} = (0, 0)$ and hence, by the argument above, $\pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}} = v_i(f)$, or that both players announce a positive integer and hence $\pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}} < v_i(f)$. Thus, since we assume that mixing occurs after d, it follows from (12) that $\sum_{(d', \underset{\sim}{z})} r(d', \underset{\sim}{z}) u_i(a^{\mathbf{h}^t, d'}, d') > v_i(f)$ for all i. But this contradicts that f is efficient.

Proof of Lemma 7 Suppose not. Then, there exists a WPEC, b, such that, by Lemma 6 applied to SPE, at some t, $\mathbf{h}^t \in \mathbf{H}^t$ and $d = (\theta, \underset{\sim}{\theta}) \in \Theta \times \Theta^2$, the two agents play 1 or 2 for sure; each i plays 1 with probability $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i}$. Furthermore, by construction,

$x_1(t')$ and $x_2(t')$ are distinct for each t' and, therefore, it follows that, for some $t' \neq t$ and $\mathbf{h}^{t'} \in H^{t'}$, and for each i, we have $b_i(\mathbf{h}^t, d) \neq b_i(\mathbf{h}^{t'}, d)$.

Now, consider any $i = 1, 2$ deviating to another strategy b_i' that is identical to the equilibrium strategy b_i except that, for all $\mathbf{h} \in \mathbf{H}^\infty$, $b_i'(\mathbf{h}, d)$ prescribes announcing 1 with probability 1. Since b_i' is less complex than b_i, we obtain a contradiction by showing that $\pi_i^{\mathbf{h}}(b_i', b_{-i}, R^e) = \pi_i^{\mathbf{h}}(b, R^e)$ for all $\mathbf{h} \in \mathbf{H}^\infty$. To do so, it suffices to fix any history $\mathbf{h}$ and consider continuation payoffs after the given partial history d. Given Lemma 6, there are two cases to consider at $(\mathbf{h}, d)$.

First, if the other agent j mixes between 1 and 2, by part 2 of Lemma 6, i is indifferent between choosing 1 and 2. Second, suppose that j plays 0 for sure. Then, by Lemma 6, i also plays 0 for sure and obtains a continuation payoff equal to $v_i(f)$ in equilibrium. Deviation also induces the same continuation payoff $v_i(f)$ as it makes i the "odd-one-out."

References

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

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

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

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

[5] Binmore, K. M. Piccione and L. Samuelson (1998): "Evolutionary Stability in Alternating-Offers Bargaining Games," *Journal of Econic Theory*, 80, 257-291.

[6] Barbera, S. and M. O. Jackson (2004): "Choosing How to Choose: Self-Stable Majority Rules and Constitutions," *Quarterly Journal of Economics*, 119, 1011-1048.

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

[8] Chatterjee, K. and H. Sabourian (2000): "Multiperson Bargaining and Strategic Complexity," *Econometrica*, 68, 1491-1509.

[9] 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.

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

[11] Fudenberg, D. and E. Maskin (1991): "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, 53, 428-438.

[12] Gale, D. and H. Sabourian (2005): "Complexity and Competition," *Econometrica*, 73, 739-770.

[13] Jackson, M. O. (1992): "Implementation in Undominated Strategies: A Look at Bounded Mechanisms," *Review of Economic Studies*, 59, 757-775.

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

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

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

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

[18] Kalai, E. and A. Neme, "The Strength of a Little Perfection," *International Journal of Game Theory*, 20 (1992), 335-355.

[19] Lee, J. and H. Sabourian (2010): "Efficient Repeated Implementation with Incomplete Information," *in preparation.*

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

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

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

[23] Maskin, E. and T. Sjöström (2002): "Implementation Theory," in *Handbook of Social Choice and Welfare, Vol. 1*, ed. by K. Arrow *et al*, North-Holland.

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

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

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

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

[28] Sabourian, H. (2004): "Bargaining and Markets: Complexity and the Competitive Outcome ," *Journal of Economic Theory*, 116, 189-228.

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

[30] Sjöström, T. (1994): "Implementation in Undominated Nash Equilibria without Integer Games," *Games and Economic Behavior*, 6, 502-511.

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

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