Complexity and Repeated Implementation*

Jihong Lee† Hamid Sabourian‡
Seoul National University University of Cambridge

January 2015

Abstract

This paper examines the problem of repeatedly implementing an efficient social choice function when the agents' preferences evolve randomly. We show that the freedom to set different mechanisms at different histories can give the planner an additional leverage to deter undesirable behavior even if the mechanisms are restricted to be simple and finite. Specifically, we construct a history-dependent sequence of simple mechanisms such that, with minor qualifications, every pure subgame perfect equilibrium delivers the correct social choice at every history, while every mixed equilibrium is strictly Pareto-dominated. More importantly, when faced with agents with a preference for less complex strategies at the margin, the (efficient) social choice function can be repeatedly implemented in subgame perfect equilibrium in pure or mixed strategies. Our results demonstrate a positive role for complexity considerations in mechanism design.

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

Keywords: Complexity, Repeated implementation, Efficiency, Finite mechanisms, Mixed strategies, Subgame perfect equilibrium

*This paper was previously circulated under the title "Repeated Implementation with Finite Mechanisms and Complexity." The authors wish to thank Ariel Rubinstein and seminar participants at CUHK, Hitotsubashi, HKU, Iowa, Michigan, Northwestern, Princeton, SMU, UCL, Malaga Economic Theory Workshop, SAET Meeting Faroe, SING Meeting Budapest and Warwick Theory Workshop for helpful comments. Jihong Lee's research was supported by the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2014S1A3A2043505) and by the Institute of Economic Research of Seoul National University.

†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

Contents

1 Introduction

The success of a society often hinges on the design of its institutions, from markets to voting. From a game-theoretic perspective, the basic requirement of an institution is that it admits an equilibrium satisfying properties that the society deems desirable, as forwarded by the literature on mechanism design. A more satisfactory way of designing an institution is to have all of its equilibria to be desirable, or to achieve *full* implementation.

In a recent paper, Lee and Sabourian [21] (henceforth, LS) extend the scope of implementation to infinitely repeated environments in which the agents' preferences evolve stochastically, and demonstrate a fundamental difference between the problems of one-shot and repeated implementation. In particular, they establish, with minor qualifications, that in complete information environments a social choice function is repeatedly implementable in Nash equilibrium if and only if it is efficient, thereby dispensing with Maskin monotonicity [25] that occupies the critical position in one-shot implementation and yet often amounts to a very restrictive requirement, incompatible with many desirable normative properties including efficiency (e.g. Mueller and Satterthwaite [34], Saijo [36]). The notion of efficiency represents the basic goal of an economic system and therefore the sufficiency results in LS offer strong implications.

Despite the appeal of its results, the full implementation approach has often been criticized for employing abstract institutions that neither square up to the demands of real world mechanism design, nor are theoretically appealing. The implementation literature has therefore engaged in multiple debates as to whether it can maintain the high standards of its theoretical objective without exposing its key results to hinge on these issues (see, for instance, the surveys of Moore [30], Jackson [16], Maskin and Sjöström [26], and Serrano [38]). The purpose of this paper is to bring the repeated analysis of LS to the realm of these debates. We adopt a novel approach that appeals to bounded rationality of agents and seek also to gain insights into a broader motivating enquiry: can a small departure from fully rational behavior on the part of individuals work in the favor of the society to broaden the scope of implementability? Specifically, we pursue the implications of agents who have a preference for less complex strategies (at the margin) on the mechanism designer's ability to discourage undesired equilibrium outcomes.[1]

Many strong implementation results (including those of LS) have been obtained through

[1] The complexity cost in our analysis is concerned with implementation of a strategy. The players are assumed to have full computational capacity to derive best responses.

the usage of unbounded integer games which rule out certain undesired outcomes via an infinite chain of dominated actions. One response in the implementation literature, both in one-shot and repeated setups, to the criticism of its constructive arguments is that the point of using abstract mechanisms is to demonstrate what can possibly be implemented in most general environments; in specific situations, more appealing constructions may also work. According to this view, the constructions allow us to show how tight the necessary conditions for implementation are. Another response in the one-shot literature has been to restrict attention to more realistic, *finite* mechanisms. However, using a finite mechanism such as the modulo game to achieve Nash implementation brings an important drawback: unwanted *mixed* strategy equilibria. This could be particularly problematic in one-shot settings since, as Jackson [15] has shown, a finite mechanism that Nash implements a social choice function could invite unwanted mixed equilibria that strictly Pareto dominate the desired outcomes.

In this paper, we apply our bounded rationality approach to the issue of implementing efficient social choice functions in a repeated environment with only simple mechanisms. In order to achieve implementation under changing preferences, a mechanism has to be devised in each period to elicit the agents' information. A key insight in LS is that the mechanisms can themselves be made contingent on past histories in a way that, roughly put, each agent's individually rational equilibrium payoff at every history is equal to the target payoff that he derives from implementation of the desired social choices. Part of the arguments for this result involves an extension of the integer game.

Here, we show that it is possible to construct a sequence of simple and finite mechanisms that has, under minor qualifications, the following equilibrium features:

- Every pure strategy subgame perfect equilibrium repeatedly implements the efficient social choice function, while every mixed strategy subgame perfect equilibrium is strictly Pareto-dominated by the pure equilibria.

- Randomization can be eliminated altogether by making the sequence of mechanisms non-stationary or history-dependent (different mechanisms are enforced at different public histories) and by invoking an additional equilibrium refinement, based on a "small" cost associated with implementing a more complex strategy.

Thus, even with simple finite mechanisms, the freedom to choose different mechanisms at different histories enables the planner to design a sequence of mechanisms such that every pure equilibrium attains the desired outcomes; at the same time, if the players were to randomize in equilibrium, the strategies would prescribe:

(i) inefficient outcomes, which therefore make non-pure equilibria in our repeated settings less plausible from the efficiency perspective (as alluded to by Jackson [15]); and, moreover,

(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.

We emphasize that, although the evolution of mechanisms follows a non-stationary path, each mechanism that we employ has a simple two-stage sequential structure and a finite number of actions that is independent of the number of players (unlike the modulo game, for instance).

Our complexity refinement is particularly appealing and marginal for two reasons. On the one hand, the notion of complexity needed to obtain the result stipulates only a partial order over strategies such that stationary behavior (i.e., always making the same choice) is simpler than taking different actions at different histories (any measure of complexity that satisfies this will suffice). On the other hand, the equilibrium refinement requires players to adopt minimally complex strategies among the set of strategies that are best responses at every history. This is a significantly weaker refinement of equilibrium than the one often adopted in the literature on complexity in dynamic games that asks strategies to be minimally complex among those that are best responses only on the equilibrium path (see Abreu and Rubinstein [2] and the survey of Chatterjee and Sabourian [6], among others).

The paper is organized as follows. In Section 2, we describe and discuss the problem of repeated implementation. Section 3 presents our main analysis and results for the case of two agents. The analysis for the case of three of more agents, appearing in Section 4, builds on from the material on the two-agent case. Section 5 presents several extensions of our results, and Section 6 offers a detailed discussion of how our work relates to the existing implementation literature, including previous studies on bounded mechanisms and mixed strategies in one-shot implementation. Section 7 concludes. Appendices and a Supplementary Material are provided to present some proofs and additional results omitted from the main text for expositional reasons.

2 The Setup

The following describe the repeated implementation setup introduced by LS.

2.1 Basic Definitions and Notation

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

An SCF f in an implementation problem $\mathcal{P}$ is a mapping $f : \Theta \to A$, and the *range* of f is the set $f(\Theta) = \{a \in A : a = f(\theta) \text{ for some } \theta \in \Theta\}$. 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 ranges belong 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 (one-shot) expected utility, or payoff, to agent i with $v(a) = (v_i(a))_{i\in I}$. Similarly, for an SCF f, define $v_i(f) = \sum_{\theta\in\Theta} p(\theta)u_i(f(\theta), \theta)$. Denoting the profile of payoffs associated with f by $v(f) = (v_i(f))_{i\in I}$, let $V = \left\{v(f) \in \mathbb{R}^I : f \in F\right\}$ be the set of expected utility profiles of all possible SCFs. Also, for a given $f \in F$, let $V(f) = \left\{v(f') \in \mathbb{R}^I : f' \in F(f)\right\}$ be the set of payoff profiles of all SCFs whose ranges belong to the range of f. We refer to $co(V)$ and $co(V(f))$ as the convex hulls of the two sets, respectively.

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

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

(b) 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 $v(f') = v(f)$; f is strongly efficient in the range if it is strictly efficient in the range and $v(f)$ is an extreme point of $co(V(f))$.

2.2 Repeated Implementation

We refer to $\mathcal{P}^\infty$ as the infinite repetitions of the implementation problem $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$. Periods are indexed by $t \in \mathbb{Z}_{++}$ and the agents' common discount factor is $\delta \in (0, 1)$. In each period, the state is drawn from Θ from an independent and identical probability distribution p. For an (uncertain) infinite sequence of outcomes $a^\infty = \left(a^{t,\theta}\right)_{t \in \mathbb{Z}_{++}, \theta \in \Theta}$, where $a^{t,\theta} \in A$ is the outcome implemented in period t and state θ, agent i's (repeated game) payoff is given by

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

We assume that the structure of $\mathcal{P}^\infty$ (including the discount factor) is common knowledge among the agents and, if there is one, the planner. The realized state in each period is complete information among the agents but unobservable to a third party.

Next, we define mechanisms and regimes. A (multi-stage) mechanism is defined as $g = ((M^g(k))_{k=1}^K, \psi)$, where K is the number of stages of the mechanism, $M^g(k) = M_1^g(k) \times \cdots \times M_I^g(k)$ is a cross product of message spaces at stage $k = 1, \ldots, K$, and, letting $M^g \equiv M^g(1) \times \cdots \times M^g(K)$, $\psi^g : M^g \to A$ is an outcome function such that $\psi^g(m) \in A$ for any K-stage history of message profiles $m = (m^1, \ldots, m^K) \in M^g$. We say that mechanism g is *finite* if K is finite and $|M_i^g(k)| < \infty$ for every agent i and stage k. Let G be the set of all *feasible* mechanisms.

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

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

We define a regime, R, as a mapping $R: H^\infty \to G$.[2] Let $R|h$ refer to the *continuation regime* that regime R induces at history $h \in H^\infty$ (thus, $R|h(h') = R(h,h')$ for any $h, h' \in H^\infty$). We say that a regime R is *history-independent* if and only if, for any t and any $h, h' \in H^t$, $R(h) = R(h')$, and that a regime R is *stationary* if and only if, for any $h, h' \in H^\infty$, $R(h) = R(h')$.

Given a regime, an agent can condition his actions on the past history of realized states as well as that of mechanisms and message profiles played. Define $\mathbf{H}^t = (\mathcal{E} \times \Theta)^{t-1}$ 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 $\mathbf{h}$. Also, since we allow for mechanisms with multi-stage sequential structure, we additionally describe information available within a period, or *partial history*. For any K-stage mechanism g and for any $1 \leq k \leq K$, let $D_k^g = \Theta \times M^g(1) \times \cdots \times M^g(k-1)$ denote the set of partial histories that can occur within the first $k-1$ stages of g. Here, we take $D_1^g = \Theta$; that is, the play of the first stage starts with the arrival of a random state. Let $D^g = \cup_k D_k^g$ and $D = \cup_g D^g$ with its typical element denoted by d.

Then, we can write each agent i's mixed (behavioral) strategy as a mapping $\sigma_i : \mathbf{H}^\infty \times G \times D \to \cup_{g,k} \mathcal{M}_i^g(k)$ such that $\sigma_i(\mathbf{h}, g, d) \in \mathcal{M}_i^g(k)$ for any $\mathbf{h} \in \mathbf{H}^\infty$, $g \in G$ and $d \in D_k^g$. 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 $\sigma_i(\mathbf{h}, g, d) = \sigma_i(\mathbf{h}', g, d)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$, $g \in G$ and $d \in D$. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is Markov if and only if σ_i is Markov for each i.

Suppose that R is the regime and σ the strategy profile chosen by the agents. Then, for any date t and history $\mathbf{h} \in \mathbf{H}^t$, we define the following:

- $g^{\mathbf{h}}(\sigma, R) \equiv (M^{\mathbf{h}}(\sigma, R), \psi^{\mathbf{h}}(\sigma, R))$ refers to the mechanism played at $\mathbf{h}$.
- $\pi_i^{\mathbf{h}}(\sigma, R)$, with slight abuse of notation, denotes agent i's expected continuation payoff at $\mathbf{h}$. For notational simplicity, let $\pi_i(\sigma, R) \equiv \pi_i^{\mathbf{h}}(\sigma, R)$ for $\mathbf{h} \in \mathbf{H}^1$.
- $A^{\mathbf{h},\theta}(\sigma, R) \subset A$ denotes the set of outcomes implemented with positive probability at $\mathbf{h}$ when the current state is θ.

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

[2]Therefore, we restrict attention to deterministic transitions of mechanisms. We below discuss how our constructive arguments can be made simpler if one allows for random transitions.

Let $\mathcal{S}$-equilibrium be a game theoretic solution concept, and given regime R with discount factor δ, let $\Omega^\delta(R) \subseteq \Sigma$ denote the set of (pure or mixed) $\mathcal{S}$-equilibrium strategy profiles. LS propose the following two notions of repeated implementation.[3]

Definition 2 *(a) An SCF f is payoff-repeatedly implementable in $\mathcal{S}$-equilibrium from period τ if there exists a regime R such that $\Omega^\delta(R)$ is non-empty and every $\sigma \in \Omega^\delta(R)$ is such that $\pi_i^{\mathbf{h}}(\sigma, R) = v_i(f)$ for any $i \in I$, $t \geq \tau$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$ on the equilibrium path.*

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

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 equilibrium delivers the correct continuation payoff profile or the correct outcomes from period τ onwards for every possible sequence of state realizations.

With no restrictions on the set of feasible mechanisms and regimes, LS established that, with some minor qualifications, an SCF satisfying efficiency in the range (strict efficiency in the range) is payoff-repeatedly implementable (repeatedly implementable) in Nash equilibrium.[4] In this paper, we pursue repeated implementation of efficient social choice functions using only simple finite mechanisms. Our approach involves adopting equilibrium refinements that incorporate credibility (subgame perfection) and complexity. While the corresponding sufficiency results in LS were based on one-shot mechanisms, the constructive arguments in the main analysis below make use of multi-stage mechanisms as it better facilitates our complexity treatment. We consider one-shot mechanisms in Section 5.1.[5]

[3]In the analysis of LS, the solution concept is Nash equilibrium and only single-stage mechanisms are considered.

[4]LS also showed that weak efficiency in the range is a necessary condition for Nash repeated implementation when the agents are sufficiently patient. In Section 6 below, we offer a detailed comparison between our results and the sufficiency results of LS.

[5]Note that the individuals play multi-stage mechanisms repeatedly in our setup, and therefore, the requirement of subgame perfection itself does not have the same bite as in one-shot implementation with extensive form mechanisms.

2.3 Complexity and Equilibrium

We next introduce a solution concept that incorporates a "small" cost associated with implementing a more complex strategy.

Complexity of a strategy in a given regime 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 within a mechanism is simpler than taking different actions in the mechanism at different histories.

Definition 3 *For any $i \in I$ and any pair of strategies $\sigma_i, \sigma_i' \in \Sigma_i$, we say that σ_i is more complex than σ_i' if the strategies are identical everywhere except, after some partial history within some mechanism, σ_i' always behaves (randomizes) the same way while σ_i does not. Formally, there exist some $g' \in G$ and $d' \in D$ with the following properties:*

(i) $\sigma_i'(\mathbf{h}, g, d) = \sigma_i(\mathbf{h}, g, d)$ *for all* $\mathbf{h} \in \mathbf{H}^\infty$ *and all* $(g, d) \neq (g', d') \in G \times D$.

(ii) $\sigma_i'(\mathbf{h}, g', d') = \sigma_i'(\mathbf{h}', g', d')$ *for all* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$.

(iii) $\sigma_i(\mathbf{h}, g', d') \neq \sigma_i(\mathbf{h}', g', d')$ *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 [5], Sabourian [35] and Gale and Sabourian [12] who consider bargaining and market games. Our results also hold with other similar complexity measures, which we discuss in further detail in Section 5 below.

Using Definition 3, we refine the set of subgame perfect equilibria as follows.

Definition 4 *A strategy profile σ is a weak perfect equilibrium with complexity cost (WPEC) of regime R if σ is a subgame perfect equilibrium (SPE) and for each $i \in I$ no other strategy $\sigma_i' \in \Sigma_i$ is such that*

(i) σ_i' *is less complex than* σ_i*; and*

(ii) σ_i' *is a best response to* σ_{-i} *at every information set for* i *(on or off the equilibrium).*

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 repeated and bargaining games that requires strategies to be minimally complex among those that are best responses only on the equilibrium path (see Section 3.4 below for a formal definition).[6] 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.[7]

2.4 Obtaining Target Payoffs

An important feature of the constructive arguments behind LS's sufficiency results is that in the repeated implementation setup one can obtain the target expected payoff associated with the desired SCF for any single player precisely as the (discounted average) expected payoff of some history-independent and non-strategic regime, as long as the discount factor is sufficiently large. Such regimes involve enforcement of some constant outcomes and/or dictatorships.

A *constant rule* mechanism refers to a mechanism that enforces a single outcome (constant SCF). Formally, $\phi(a) = (M, \psi)$ is a one-stage mechanism such that $M_i = \{\emptyset\}$ for all $i \in I$ and $\psi(m) = a \in A$ for all $m \in M$. Also, for any SCF $f \in F$, let $d(i)$ denote a (one-stage) *dictatorial* mechanism (or simply i-dictatorship) in which agent i is the dictator and can choose any outcome in the *range* of f; thus, $d(i) = (M, \psi)$ is such that $M_i = f(\Theta)$, $M_j = \{\emptyset\}$ for all $j \neq i$ and $\psi(m) = m_i$ for all $m \in M$.

Note that, for any i, the dictatorship $d(i)$ must yield a unique expected utility $v_i^i = \sum_{\theta\in\Theta} p(\theta) \max_{a\in f(\Theta)} u_i(a, \theta)$ to dictator i if he acts rationally in the mechanism; but since there may not a be a unique optimal outcome for the dictator, multiple payoffs can arise for the other players. For simplicity, we assume throughout that, for each i, $d(i)$ also yields a unique payoff to every $j \neq i$ if i acts rationally. Our results below are not affected by relaxing this assumption (see Section 1 of the Supplementary Material). Let v_j^i denote the unique payoff of player j from $d(i)$ if i acts rationally, and let $v^i = (v_i^1, .., v_i^I)$. Clearly,

[6] The two exceptions in the existing literature are Kalai and Neme [18] and Sabourian [35]. The notion of WPEC was first introduced by [35].

[7] Note also that complexity cost enters the agents' preferences lexicographically. All our results below hold when the decision maker admits a positive trade-off between complexity cost and (on- or off-path) payoff.

$v_i^i \geq v_i(f)$.

Next, let $\mathcal{W} = \{v^i\}_{i \in I} \cup \{v(a)\}_{a \in A}$ denote the set of all (one-period) payoff profiles from dictatorial and constant rule mechanisms. Then, by constructing history-independent regimes that alternate between those mechanisms, we have the following.

Lemma 1 *(a) Suppose that $\delta \in \left(\frac{1}{2}, 1\right)$. Fix any $i \in I$, and suppose that there exists an outcome $\tilde{a}^i \in A$ such that $v_i(\tilde{a}^i) \leq v_i(f)$. Then, there exists a history-independent regime that generates a unique (discounted average) payoff to agent i equal to $v_i(f)$ in any Nash equilibrium.*

(b) Fix any $\mathcal{W}^* \subset \mathcal{W}$ and suppose that $\delta \in \left(1 - \frac{1}{|\mathcal{W}^*|}, 1\right)$. Then, for any payoff profile $w \in co(\mathcal{W}^*)$ and any $\epsilon > 0$, there exists a history-independent regime that generates a unique Nash equilibrium payoff w' such that $\mid w - w' \mid < \epsilon$.

Proof. (a) Fix any i. Note that $v_i^i \geq v_i(f) \geq v_i(\tilde{a}^i)$. Since $\delta > \frac{1}{2}$, by the algorithm of Sorin [40] (see also Lemma 3.7.1 of Mailath and Samuelson [24]), we can construct a history-independent regime, by appropriately alternating $d(i)$ and $\phi(\tilde{a}^i)$, such that player i obtains a unique equilibrium payoff equal to $v_i(f)$.

(b) Fix any $w \in co(\mathcal{W}^*)$ and any $\epsilon > 0$. Since $\delta \in \left(1 - \frac{1}{|\mathcal{W}^*|}, 1\right)$, by Sorin [40], there exists a history-independent regime that appropriately alternates between the dictatorial and constant rule mechanisms whose one-period payoffs belong to $\mathcal{W}^*$ such that if each dictator always picks his best one-shot alternative, the resulting discounted average payoff profile is w. Let us denote this regime by W.

To obtain a regime that induces a unique payoff w' such that $\mid w - w' \mid < \epsilon$ in any Nash equilibrium, consider another regime that is identical to W for some finite T periods, followed by non-contingent implementation of some arbitrary but fixed outcome in every period thereafter. The resulting Nash equilibrium payoffs are unique since each player must then pick his best one-shot outcome whenever he is dictator; moreover, if T is sufficiently large, the payoffs satisfy $\mid w - w' \mid < \epsilon$. ■

Part (a) of Lemma 1 also appeared in LS; part (b) is new and will be exploited in our analysis. We raise two remarks on the constructions in this lemma. First, while for any $w \in co(\mathcal{W}^*)$, the algorithm of Sorin [40] implies that there must exist a regime that induces w as a Nash equilibrium payoff profile by appropriately alternating between dictatorial and constant rule mechanisms, we cannot ensure that this is the case in *any* Nash equilibrium of the regime. The reason is that if the regime involves serial dictatorships by different players then it may be possible for collusion to occur intertemporally and induce outcome paths that are different from those in which each dictator always follows his unique myopic best action. Hence, a regime with infinitely many dictatorships given to multiple players can generate multiple equilibrium payoffs. To avoid such collusive outcomes, our construction in the proof of part (b) activates a permanent implementation of some constant outcome beyond a certain finite period T. The uniqueness of equilibrium continuation payoff at T ensures, via backward induction, that each player behaves myopically whenever he is dictator; at the same time, if T is large, the impact on the average payoffs is small.

Second, in this paper we restrict ourselves to deterministic regimes. One consequence of this restriction is that the results described in the above lemma require the discount factor to be above a certain positive lower bound. Another consequence is that the regimes constructed in the proof of the above results involve a precise sequencing of dictatorial and constant rule mechanisms. If the planner were able to condition the choice of mechanism on some random public signals, each of the results in the above lemma could be obtained by constructing an alternative regime that initially chooses among the set of dictatorial and constant rule mechanisms according to the appropriate probability distribution and then repeats the realized mechanism forever thereafter. Such regimes establish the above lemma for any δ and do not involve changing the mechanisms at different dates. Thus, with public randomization, we could do away from imposing any restriction on the discount factor and construct in some sense simpler mechanisms that do not require precise tracking of the time.[8] A public randomization device may not be available, however.

3 Two Agents

In this paper, we first report our results for $I = 2$. Our approach to the case of $I \geq 3$ involves more complicated constructions that will build on the material of this section.

[8]In part (b) of Lemma 1 the payoff vector w is obtained approximately because we want to ensure that collusion among different dictators does not occur. With public randomization, the collusion possibility no longer poses an issue and the result in (b) can in fact be obtained exactly.

3.1 Regime Construction

Our objective is to obtain a repeated implementation result for an SCF that is efficient in the range. To do so, we introduce two additional properties of the SCF.

First, as in the one-shot implementation problem, there is a difference between the two-agent and three-or-more-agent cases in our setup for ensuring the existence of truth-telling equilibrium. This is due to the fact that, with two agents, it is not possible to identify the misreport in the event of disagreement. One way to deter deviations from truth-telling in our regime construction with $I = 2$ is to invoke an additional requirement known as self-selection, as adopted in the one-shot literature.[9] Formally, for any f, i and θ, let $L_i(\theta) = \{a \in f(\Theta) | u_i(a, \theta) \leq u_i(f(\theta), \theta)\}$ be the set of outcomes among the range of f that make agent i worse off than f. We say that f satisfies *self-selection in the range* if $L_1(\theta) \cap L_2(\theta') \neq \emptyset$ for any $\theta, \theta' \in \Theta$. We assume this condition here for ease of exposition. It can be dropped when the agents are sufficiently patient, since intertemporal incentives can then be designed to support truth-telling; see Section 2 of the Supplementary Material for a formal analysis.

Second, as in LS, we consider SCFs induce payoffs that are, for each player, bounded below by the payoff of some constant outcome. This enables us to appeal to Lemma 1 above to build continuation regimes that provide correct intertemporal incentives for full repeated implementation.

Condition ϕ. (i) For all $i \in I$, there exists $\tilde{a}^i \in f(\Theta)$ such that $v_i(\tilde{a}^i) \leq v_i(f)$.

(ii) For all $i \in I$ and $\gamma \in [0, 1]$, $v(f) \neq \gamma v^i + (1 - \gamma) v(\tilde{a}^i)$.

Part (i) strengthens condition ω appearing in LS by requiring the outcome $\tilde{a}^i$ to be found in the range of the SCF, while the inequality here is allowed to be weak. The restriction to $f(\Theta)$ is imposed to obtain repeated implementation of an SCF that is efficient in the range and hence can be relaxed if one deals with efficiency instead. Part (ii) is almost without loss of generality. If this were not to hold, the history-independent regime described in the proof of part (a) of Lemma 1 would induce a unique payoff profile equal $v(f)$ if $\delta > 1/2$.[10]

[9] This condition is originally from Dutta and Sen [8] and is weaker than the "bad outcome" condition in Moore and Repullo [32].

[10] In fact, by Fudenberg and Maskin [11], one could also build a regime such that the agents' continuation payoffs at every date approximate $v(f)$ with δ sufficiently close to 1.

Using condition ϕ, we can construct some history-independent regimes that yield unique average payoffs with the following specific properties.

Lemma 2 *Suppose that $I = 2$, and fix an SCF f that satisfies efficiency in the range and condition ϕ. Suppose also that $\delta \in \left(\frac{3}{4}, 1\right)$. Then, we obtain the following:*

(a) For each $i \in I$, there exists a history-independent regime, referred to as S^i, that yields a unique Nash equilibrium (discounted average) payoff profile $w^i = (w^i_1, w^i_2)$ such that

$$w^i_i = v_i(f) \quad \text{and} \quad w^i_j < v_j(f),\ j \neq i.$$

(b) There exist history-independent regimes $\{X(t)\}_{t=1,2,\ldots}$ and Y that respectively induce unique Nash equilibrium payoff profiles $x(t) = (x_1(t), x_2(t))$ and $y = (y_1, y_2)$ satisfying the following condition:

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

Proof. (a) Fix any i. Then, by part (a) of Lemma 1, there exists a regime S^i that induces a unique payoff profile $w^i = (w^i_1, w^i_2)$ such that $w^i_i = v_i(f)$. Efficiency in the range of f and part (ii) of condition ϕ imply that, for $j \neq i$, $w^i_j < v_j(f)$.

(b) To construct regimes $\{X(t)\}_{t=1,2,\ldots}$ and Y, we first set, for each date t,

$$x'(t) = \lambda(t)w^1 + (1 - \lambda(t))w^2 \quad \text{and} \quad y' = \mu w^1 + (1 - \mu)w^2$$

for some $0 < \mu < \lambda(t) < 1$. Since $w^j_i < w^i_i$ for all i, j, $i \neq j$, the resulting payoffs satisfy, for any t,

$$w^2_1 < y'_1 < x'_1(t) < w^1_1 \text{ and } w^1_2 < x'_2(t) < y'_2 < w^2_2. \tag{2}$$

Since w^i is itself generated by a convex combination of v^i and $v(\tilde{a}^i)$, it follows that $x'(t)$ and y' can be written as convex combinations of $v^1, v^2, v(\tilde{a}^1)$ and $v(\tilde{a}^2)$. Then, since $\delta > 3/4$, and by part (b) of Lemma 1, there must exist history-independent regimes that induce unique equilibrium payoffs that are arbitrarily close to $x'(t)$ and y'. Since $x'(t)$ and y' satisfy the strict inequalities described in (2), we also have regimes $X(t)$, for each t, and Y that induce unique payoffs $x(t)$ and y, respectively, satisfying (1). ■

We assume throughout in this section that $\delta > 3/4$ as required by Lemma 2. As mentioned before, this assumption is not needed with a public randomization device as the results described in Lemma 2 could then be obtained for any δ (see footnote 14 below).

These constructions are illustrated in Figure 1 below (where, with slight abuse of notation, π_i refers to i's average repeated game payoff).

Figure 1: Regime construction

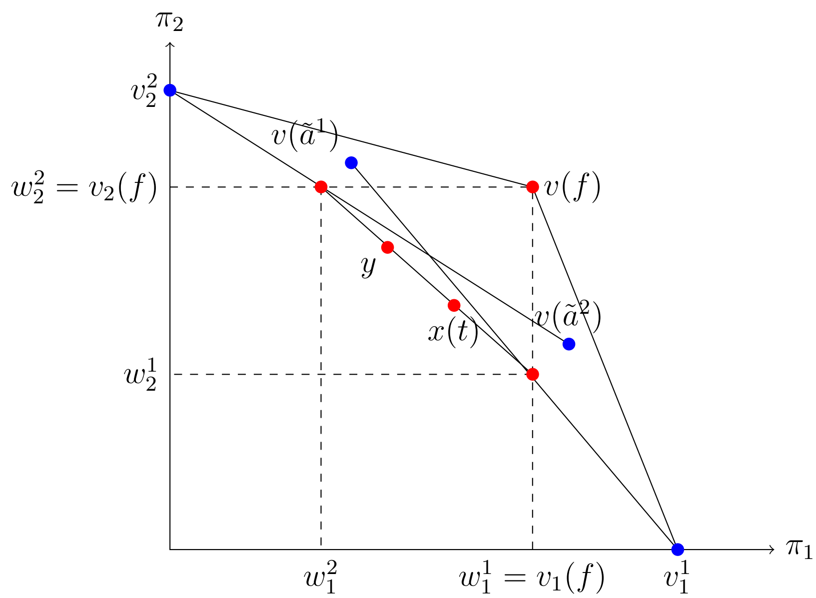

Now, for an SCF f that satisfies efficiency in the range, self-selection in the range and condition ϕ, we define the following multi-stage mechanism, referred to as g^e:

- Stage 1 - Each agent $i = 1, 2$ announces a state, θ_i, from Θ.
- Stage 2 - Each agent announces an integer, z_i, from the set $\mathcal{Z} \equiv \{0, 1, 2\}$.

The outcome function of this mechanism depends solely on the agents' announcement of states in Stage 1 and is given below:

(i) If $\theta_1 = \theta_2 = \theta$, $f(\theta)$ is implemented.

(ii) Otherwise, an outcome from the set $L_1(\theta_2) \cap L_2(\theta_1)$ is implemented.

Using this mechanism together with the history-independent regimes $X(t)$ and Y constructed above, we define regime R^e inductively as follows.

First, mechanism g^e is played in $t = 1$. Second, if, at some date $t \geq 1$, g^e is the mechanism played with a pair of states $\underset{\sim}{\theta} = (\theta_1, \theta_2)$ announced in Stage 1 followed by integers $\underset{\sim}{z} = (z_1, z_2)$ in Stage 2, the continuation mechanism or regime at the next period is given by the transition rules below:

- Rule A.1: If $z_1 = z_2 = 0$, then the mechanism next period is g^e.
- Rule A.2: If $z_i > 0$ and $z_j = 0$ for some $i, j = 1, 2$, then the continuation regime is S^i.
- Rule A.3: If $z_1, z_2 > 0$, then we have the following:
  - Rule A.3(i): If $z_1 = z_2 = 1$, the continuation regime is $X \equiv X(\tilde{t})$ for some arbitrary but fixed $\tilde{t}$, with the payoffs henceforth denoted by x.
  - Rule A.3(ii): If $z_1 = z_2 = 2$, the continuation regime is $X(t)$.
  - Rule A.3(iii): If $z_1 \neq z_2$, the continuation regime is Y.

This regime thus employs only the outcomes in the range of the SCF, $f(\Theta)$. Let us summarize other key features of this regime construction. First, in mechanism g^e, which deploys only two stages, the implemented outcome depends solely on the announcement of states, while the integers dictate the continuation mechanism. The set of integers contains only three elements. Second, announcement of any non-zero integer effectively ends the strategic part of the game. When only one agent, say i, announces a positive integer this agent obtains his target payoff $v_i(f)$ in the continuation regime S^i (Rule A.2). The rest of transitions are designed to rule out unwanted randomization behavior. In particular, when both agents report positive integers, by (1), the continuation regimes are such that the corresponding continuation payoffs, $x(t)$ or y, are strictly Pareto-dominated by the target payoffs $v(f)$. Furthermore, when both agents report integer 2 (Rule A.3(ii)) the continuation regimes could actually be different across periods. This feature will later be used to facilitate our complexity refinement arguments.

Note that, in this regime, the histories that involve strategic play are only those at which the agents engage in mechanism g^e. At any other public history h, the continuation regime is S^i, $X(t)$ for some t or Y involving only dictatorial and constant rule mechanisms. Furthermore, by Lemma 2, in any subgame perfect equilibrium, the continuation payoff at any such h is unique and given by w^i, $x(t)$ for some t or y, respectively, satisfying the conditions described in Lemma 2 if $\delta > 3/4$.[11]

As mentioned before, the only histories that matter for strategic play here are those at which the next mechanism is g^e. Therefore, with some abuse of notation, we simplify the definitions of relevant histories and strategies for the above regime as follows.

We denote by $\mathbf{H}^t$ the set of all finite histories observed by the agents at the beginning of period t given that the mechanism to play in t is g^e; let $\mathbf{H}^\infty = \cup_{t=1}^{\infty} \mathbf{H}^t$. Let $D_1 = \Theta$ and $D_2 = \Theta \times \Theta^I$ denote the set of partial histories at Stage 1 and at Stage 2 of the two-stage mechanism g^e, respectively. Thus, $d = \theta$ is a partial history that represents the beginning of Stage 1 after state θ has been realized, and $d = (\theta, \underset{\sim}{\theta}) \in D_2$ refers to the beginning of Stage 2 after realization of θ followed by profile $\underset{\sim}{\theta} \in \Theta^2$ announced in Stage 1.

Then, a mixed (behavioral) strategy of agent $i = 1, 2$ in regime R^e is written simply as the mapping $\sigma_i : \mathbf{H}^\infty \times D \to (\triangle\Theta) \cup (\triangle\mathcal{Z})$ such that, for any $\mathbf{h} \in \mathbf{H}^\infty$, $\sigma_i(\mathbf{h}, d) \in \triangle\Theta$ if $d \in D_1$ and $\sigma_i(\mathbf{h}, d) \in \triangle\mathcal{Z}$ if $d \in D_2$. Let Σ_i be the set of i's strategies in R^e. We write $\pi_i^{\mathbf{h}}(\sigma, R^e)$ as player i's continuation payoff under strategy profile σ at history $\mathbf{h} \in \mathbf{H}^\infty$.

3.2 Subgame Perfect Equilibria

We first consider the set of subgame perfect equilibria of the above regime. Let us begin by establishing existence of an equilibrium in which the desired social choice is always implemented. In this equilibrium, both players adopt Markov strategies, always announcing the true state followed by integer zero.

Lemma 3 *Regime R^e admits a subgame perfect equilibrium (SPE), σ^*, in Markov strategies such that, for any t, $\mathbf{h} \in \mathbf{H}^t$ and $\theta \in \Theta$, (i) $g^{\mathbf{h}}(\sigma^*, R^e) = g^e$ and (ii) $A^{\mathbf{h},\theta}(\sigma^*, R^e) = \{f(\theta)\}$.*

Proof. Consider $\sigma^* \in \Sigma$ such that, for all i, $\sigma_i^*(\mathbf{h}, \theta) = \theta$ for all $\mathbf{h} \in \mathbf{H}^\infty$ and $\theta \in D_1$, and $\sigma_i^*(\mathbf{h}, (\theta, \underset{\sim}{\theta})) = 0$ for all $\mathbf{h} \in \mathbf{H}^\infty$ and $\left(\theta, \underset{\sim}{\theta}\right) \in D_2$.[12] Clearly, this profile satisfies (i) and (ii) in the claim. Thus, at any $\mathbf{h} \in \mathbf{H}^\infty$, $\pi_i^{\mathbf{h}}(\sigma^*, R^e) = v_i(f)$ for all i.

[11] If regimes with random transitions are feasible then we can obtain the same results with no restriction on δ by constructing an alternative regime that is otherwise identical to R^e except that whenever at least one player announces a positive integer, the counterpart regime randomizes, with appropriate probability distributions, between four stationary continuation regimes, each of which repeatedly enforces $d(1)$, $d(2)$, $\phi(\tilde{a}^1)$ or $\phi(\tilde{a}^2)$.

[12] Here we have abused the notation slightly to describe pure strategies.

To show that σ^* is an SPE, consider a unilateral one-step deviation by any agent i. Fix any $\mathbf{h} \in \mathbf{H}^\infty$. There are two cases to consider. First, fix any partial history θ. By the outcome function of g^e and self-selection in the range, one-step deviation to a non-truthful state does not improve one-period payoff; also, since the other player's strategy is Markov and the transition rules do not depend on Stage 1 actions, the continuation payoff at the next period is unaffected. Second, fix any partial history $\left(\theta, \underset{\sim}{\theta}\right)$. In this case, by Rule A.2, the continuation payoff from deviating to any positive integer is identical to the equilibrium payoff, which is equal to $v_i(f)$. ■

We now turn to characterizing the properties of the set of SPEs. Our next Lemma is concerned with the players' equilibrium behavior whenever they face Stage 2 (the integer part) of mechanism g^e. It shows that at any such history both players must 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 $v(f)$. Thus, in terms of continuation payoffs, mixing is strictly Pareto-dominated by the pure strategy equilibrium.

Lemma 4 *Consider any SPE of regime R^e. Fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d = (\theta, \underset{\sim}{\theta}) \in D_2$. Then, one of the following must hold at $(\mathbf{h}, d)$:*

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

(b) Each agent 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 $v_i(f)$.

Proof. See Appendix A.1. ■

To gain intuition for the above result, consider the matrix below that contains the corresponding continuation payoffs when at least one player announces a positive integer.

First, from Figure 2, the inequalities of (1) imply that any equilibrium with pure strategy at the relevant history must play 0. Since the continuation regime S^i gives i his target payoff $w_i^i = v_i(f)$ and a payoff, w_j^i, that is strictly lower than y_j for the other player j, a strictly profitable deviation opportunity arises whenever there is an "odd-one-out" announcing a positive integer; if both players announce positive integers, the fact that $x_1(t) > y_1$ and $y_2 > x_2(t)$ imply that a deviation opportunity exists for one of the two players.

Figure 2: Continuation payoffs

		Player 2		
		0	1	2
	0	$\cdot$	w^2	w^2
Player 1	1	w^1	x	y
	2	w^1	y	$x(t)$

Second, if all the players announce zero, it then follows that each player i's continuation payoff must be bounded below by $v_i(f)$ since, otherwise, the player could deviate by reporting a positive integer and obtain $v_i(f)$ from the continuation regime S^i (Rule A.2). Since this is true for all i, the efficiency in the range of the SCF then implies that the continuation payoffs are equal to the target payoffs for all agents.

Next, 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 $v_i(f)$, with it being strictly greater than $v_i(f)$ when j plays a positive integer with positive probability. Applying this argument to both agents leads to a contradiction against the assumption that the SCF is *efficient in the range*.

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 4, we can also show that if the players were to mix over integers at any history on the equilibrium path, it must occur 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. The properties of subgame perfect equilibria of our regime can then be summarized as follows.

Proposition 1 *Consider any SPE σ of regime R^e. Then, one of the following must hold:*

(a) Each agent i announces 0 for sure at any $(\mathbf{h}, d) \in \mathbf{H}^\infty \times D_2$ on the equilibrium path, and $\pi_i^{\mathbf{h}}(\sigma, R^e) = v_i(f)$ for any $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t$ on the equilibrium path.

(b) Each agent i mixes between 1 and 2 at some $d \in D_2$ in period 1 on the equilibrium path, and his continuation payoff at the next period is less than $v_i(f)$; hence, $\pi_i(\sigma, R^e) < v_i(f)$ if δ is sufficiently large.

Proof. See Appendix A.1. ■

Thus, if we restrict attention to pure strategies, the first part of this Proposition and Lemma 3 imply that we obtain payoff-repeated implementation in subgame perfect equilibrium from period 2.[13] Furthermore, any mixed strategy equilibrium of our regime is strictly Pareto-dominated by any pure strategy equilibrium in terms of continuation payoffs from period 2.

3.3 WPEC

Our characterization of SPEs of regime R^e demonstrates that in any equilibrium the players must either continue along the desired path of play or fall into coordination failure early on in the game by mixing over the positive integers in period 1 which leads to strictly inefficient continuation payoffs. We now introduce our refinement (WPEC) arguments based on complexity considerations to select the former.

In order to obtain our selection results, we add to the construction of R^e the following property: the sequence of regimes $\{X(t)\}_{t=1}^\infty$ is such that, in addition to (1) above, the corresponding payoffs $\{x(t)\}_{t=1}^\infty$ satisfy

$$x_1(t') \neq x_1(t'') \text{ and } x_2(t') \neq x_2(t'') \text{ for some } t', t''. \tag{3}$$

Note that one way to achieve this involves taking the sequence $\{\lambda(t) : \lambda(t) \in (\mu, 1)\ \forall(t)\}$ used to construct these regimes in the proof of part (b) of Lemma 2 and set it such that $\lambda(t') \neq \lambda(t'')$ for at least two distinct dates t' and t''.

Clearly, this additional feature does not alter Lemmas 3 and 4. However, it implies for any SPE that, if at some period t on or off the equilibrium path, an agent mixes

[13] Note that, in part (a) of Proposition 1, there might still exist equilibria where players randomize over different state announcements at the first stage of the extensive form mechanism at any period. Even so, efficiency implies that they obtain the continuation payoff profile $v(f)$.

over integers, by choosing integer 1 with probability $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i}$, then his behavior in the integer part of mechanism g^e is not stationary across periods.[14]

We next show that if the players face a small cost of implementing a more complex strategy, mixing over integers can no longer be part of equilibrium behavior in our regime.

Lemma 5 *Fix any WPEC of regime R^e. Also, fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d \in D_2$ (on or off the equilibrium path). Then, each agent announces zero for sure at this history.*

Proof. See Appendix A.1. ■

To obtain this lemma we suppose otherwise. Then, some agent must respond differently to some partial history $d \in D_2$ depending on what happened in the past. But then, this agent could deviate to another less complex strategy identical to the equilibrium strategy everywhere except that it always responds to d 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.

Note that, since Markov strategies are simplest strategies according to Definition 3, Lemma 3 continues to hold with WPEC. Thus, combining the previous lemmas, we establish the following main result.

Theorem 1 *Suppose that $I = 2$ and $\delta \in \left(\frac{3}{4}, 1\right)$. If an SCF f is efficient in the range, and satisfies self-selection in the range and condition ϕ, f is payoff-repeatedly implementable in WPEC from period 2.*

Proof. This follows immediately from Lemmas 3-5. ■

[14] Our results are unaffected by making $X(\cdot)$ dependent on the entire history and not just its date. See Section 7 for further discussion on this issue.

Notice that the extent of implementation achieved here actually goes beyond that of Definition 2 and Theorem 1 since we obtain the desired payoffs at every on- and *off*-the-equilibrium history after period 1 at which mechanism g^e is played. To see this, combine Lemma 5 with part (a) of Lemma 4.

To obtain repeated implementation in terms of outcomes, as in LS, we need to go beyond efficiency in the range. LS assume pure strategies and hence invoke strict efficiency; here, we use strong efficiency.

Corollary 1 *Suppose that, in addition to the conditions in Theorem 1, f is strongly efficient in the range. Then, f is repeatedly implementable in WPEC from period 2.*

Proof. It suffices to show that every WPEC σ of R^e is such that $A^{\mathbf{h},\theta}(\sigma, R^e) = \{f(\theta)\}$ for any $t \geq 2$, $\mathbf{h} \in \mathbf{H}^t(\sigma, R^e)$ and $\theta \in \Theta$.

Fix any WPEC σ of regime R^e. Also, fix any $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t$. For each θ and $a \in f(\Theta)$, let $r(a, \theta)$ denote the probability that outcome a is implemented in equilibrium at $(\mathbf{h}, \theta)$. By Lemmas 4 and 5, we know that, for any i,

$$\pi_i^{\mathbf{h}}(\sigma, R^e) = (1-\delta) \sum_{\theta \in \Theta, a \in f(\Theta)} p(\theta) r(a,\theta) u_i(a,\theta) + \delta v_i(f) = v_i(f),$$

which implies that $\sum_{\theta \in \Theta, a \in f(\Theta)} p(\theta) r(a,\theta) u_i(a,\theta) = v_i(f)$.

Strong efficiency in the range implies that there does not exist a random SCF $\xi : \Theta \to \triangle(f(\Theta))$ such that $v(\xi) = v(f)$. Therefore, the claim follows. ∎

3.4 Further Refinement and Period 1

Our results do not ensure implementation of the desired outcomes in period 1. One way to sharpen our results in this direction is to consider a stronger equilibrium refinement in line with the standard literature on strategic complexity in dynamic games (e.g. Abreu and Rubinstein [2], Sabourian [35], Lee and Sabourian [20]) and to require the strategies to be minimally complex mutual best responses only *on* the equilibrium path.

Definition 5 *A strategy profile σ is a perfect equilibrium with complexity cost (PEC) of regime R^e if σ is an SPE and for each $i \in I$ no other strategy $\sigma_i' \in \Sigma_i$ is such that (i) σ_i' is less complex than σ_i and (ii) σ_i' is a best response to σ_{-i}.*

Compared with WPEC, this concept prioritizes the complexity of their strategies over off-the-equilibrium payoffs and hence selects minimally complex strategies over a larger set.

An alternative way of thinking about the issue of credibility of strategies and complexity considerations is to introduce two kinds of perturbations and find the limiting Nash equilibrium behavior as these perturbations become arbitrarily small (e.g. Chatterjee and Sabourian [5], Sabourian [35] and Gale and Sabourian [12]). One perturbation allows for a small but positive cost of choosing a more complex strategy; another perturbation represents a small but positive and independent probability of making an error (off-the-equilibrium-path move). The notions of WPEC and PEC can then be interpreted as the limiting Nash behavior as the two types of perturbation go to zero. The difference is that the WPEC results hold for such limiting equilibria *independently* of the order of the limiting arguments, while with PEC the order of limit is the complexity cost first and then the tremble.

Clearly, the Markov equilibrium constructed for Lemma 3 is a PEC, and since WPEC is itself a PEC, all our previous results on the properties of WPEC above remain valid under PEC. Additionally, we can show that every PEC in regime R^e must be Markov. Therefore, under the notion of PEC, the same repeated implementation results as in Theorem 1 above are obtained from period 1.[15]Appendix A.2 presents the formal details.[16]

4 Three or More Agents

In this section, we extend the arguments developed for the two-agent case to deal with the case of three or more agents.

[15]Similar complexity considerations were introduced by LS (see the Supplemental Material [22]) to their constructions with integer games to achieve efficient repeated implementation from period 1. Mezzetti and Renou [29] show repeated Nash implementation of dynamically monotonic SCFs from $t = 1$.

[16]Our WPEC and PEC results in this paper assume lexicographic preferences with respect to complexity cost. The same set of results also hold when the agents make a positive trade-off between complexity cost and payoffs. This is because the set of equilibria then is contained in the set of equilibria with lexicographic preferences, while our regime constructions admit an equilibrium that is Markov and hence cannot be improved upon in terms of complexity.

4.1 Regime Construction

An important part of our constructive arguments with two agents above was to construct for each i a history-independent and non-strategic regime S^i that generates a payoff profile $w^i = (w^i_i, w^i_j)$ such that $w^i_i = v_i(f)$ and $w^i_j < v_j(f)$. This was facilitated by condition ϕ. To address the problem of three or more agents, we make use of a slightly strengthened version below. Similarly to before, assume for simplicity that every i-dictatorship yields unique payoffs under rational dictator.

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

(ii) $v_i^j(f(\Theta)) < v_i(f)$ for all $i, j \in I$, $i \neq j$.

Part (i) states that there exists some outcome that is strictly worse for all agents than the given SCF in expectation. It is similar to but weaker than the "bad outcome" condition in Moore and Repullo (1990) as the outcome $\tilde{a}$ is compared to the SCF only on average (and not in every state). This amounts to a weak requirement on the SCF and is naturally satisfied in many environments (e.g. zero consumption in allocation problems). Part (ii) here would be satisfied with weak inequality when there are two agents and f is efficient in the range.

Now, consider any SCF that is efficient in the range and satisfies condition ϕ^*. Then, as in the two-player case, by part (i) of condition ϕ^* and by part (a) of Lemma 1, one can construct for each agent i a regime S^i that induces a unique Nash equilibrium payoff profile $w^i = (w^i_1, \ldots, w^i_I)$ such that $w^i_i = v_i(f)$. Efficiency in the range of f and condition ϕ^* imply that, for $j \neq i$, $w^i_j < v_j(f)$.

Fix any two agents k and l, and the corresponding regimes S^k and S^l with unique payoffs w^k and w^l. We can further extend part (b) of Lemma 2 for the two-agent case to show the existence of the following (history-independent) regimes.

Lemma 6 *Suppose that f is efficient in the range and satisfies condition ϕ^*. Also, suppose that $\delta \in \left(1 - \frac{1}{|I|+1}, 1\right)$. Then, for any subset of agents $C \subseteq I$ and each date $t = 1, 2, \ldots$, there exist regimes $S^C, X(t), Y$ that respectively induce unique Nash equilib-*

rium payoff profiles $w^C, x(t), y \in co(V(f))$ satisfying the following conditions:[17]

$$w_k^l < y_k < x_k(t) < w_k^k \quad and \quad w_l^k < x_l(t) < y_l < w_l^l \tag{4}$$

$$x_k(t') \neq x_k(t'') \quad and \quad x_l(t') \neq x_l(t'') \quad for\ some\ t', t'' \tag{5}$$

$$w_k^C < w_k^k \quad if\ C \neq \{k\} \quad and \quad w_l^C < w_l^l \quad if\ C \neq \{l\} \tag{6}$$

$$w_i^C > w_i^{C \backslash \{i\}} \quad for\ all\ i \in C. \tag{7}$$

Proof. Given the two fixed agents k and l, and the corresponding regimes S^k and S^l, we construct regimes $\{X(t)\}_{t=1,2,\ldots}$ and Y to satisfy (4) and (5) similarly to their two-agent counterparts.

To construct S^C, fix any $C \subseteq I$, and let $\overline{w}^C = \frac{1}{|C|}\sum_{i \in C} w^i$. Since $w_j^j > w_j^i$ for all $i, j \in I$, $j \neq i$, $\{\overline{w}^C\}_C$ satisfy the strict inequalities described in (6) and (7). Furthermore, since w^i is generated by a convex combination of v^i and $v(\tilde{a})$, $\overline{w}^C$ can be written as a convex combination of $v^1, \ldots, v^I$ and $v(\tilde{a})$. Then, since $\delta > 1 - \frac{1}{|I|+1}$, and by part (b) of Lemma 1, there must exist a history-independent regime that induces a unique Nash payoff profile that is arbitrarily close to $\overline{w}^C$. Since $\{\overline{w}^C\}_C$ satisfy the strict inequalities as in (6) and (7), there in turn exists a regime S^C which induces a unique Nash payoff profile w^C such that (6) and (7) are satisfied. ■

Using the constructions in Lemma 6, we extend the regime construction for the case of $I = 2$ to the case of $I \geq 3$. Fix throughout $\delta > 1 - \frac{1}{|I|+1}$.

First, define the sequential mechanism $\hat{g}^e$ as follows:

- Stage 1 - Each agent i announces a state from Θ.
- Stage 2 - Each of agents k and l announces an integer from the set $\{0, 1, 2\}$; each $i \in I \backslash \{k, l\}$ announces an integer from the set $\{0, 1\}$.

The outcome function of this mechanism again depends solely on the action of Stage 1 and is given below:

(i) If at least $I - 1$ agents announce θ, then $f(\theta)$ is implemented.

(ii) Otherwise, $f(\tilde{\theta})$ for some arbitrary but fixed $\tilde{\theta}$ is implemented.

It is important to note that this mechanism extends mechanism g^e above by allowing *only two* agents to choose from $\{0, 1, 2\}$ while all the remaining agents choose from just $\{0, 1\}$.

[17] Note that when C consists of single player i, S^C means S^i and $w^C = w^i$.

Next, using the constructions in Lemma 6 above, we define new regime $\widehat{R}^e$ inductively as follows: (i) mechanism $\hat{g}^e$ is implemented at $t = 1$ and (ii) if, at some date t, $\hat{g}^e$ is the mechanism played with a profile of states $\underset{\sim}{\theta} = (\theta_1, \ldots, \theta_I)$ announced in Stage 1 and a profile of integers $\underset{\sim}{z} = (z_1, \ldots, z_I)$ announced in Stage 2, the continuation mechanism/regime at the next period is as follows:

- Rule B.1: If $z_i = 0$ for all i, then the mechanism next period is $\hat{g}^e$.
- Rule B.2: If $z_k > 0$ and $z_l = 0$ ($z_k = 0$ and $z_l > 0$), then the continuation regime is S^k (S^l).
- Rule B.3: If $z_k, z_l > 0$, then we have the following:
  - Rule B.3(i): If $z_k = z_l = 1$, then the continuation regime is $X \equiv X(\tilde{t})$ for some arbitrary but fixed $\tilde{t}$, with the payoffs henceforth denoted by x.
  - Rule B.3(ii): If $z_k = z_l = 2$, then the continuation regime is $X(t)$.
  - Rule B.3(iii): If $z_k \neq z_l$, then the continuation regime is Y.
- Rule B.4: If, for some $C \subseteq I \backslash \{k, l\}$, $z_i = 1$ for all $i \in C$ and $z_i = 0$ for all $i \notin C$, then the continuation regime is S^C.

This regime extends the two-agent counterpart R^e by essentially maintaining all the features for two players (k and l) and endowing the other agents with the choice of just 0 or 1. However, the regime prioritizes these two selected agents when determining the transition of mechanism: note from Rules B.2 and B.3 that if either k or l plays a non-zero integer the integer choices of other players are irrelevant to transitions. We emphasize that the size of the integer set in our construction is actually independent of the number of agents. Histories, partial histories (within period), strategies and continuation payoffs are defined similarly to their two-agent counterparts.

4.2 Results

We present below our main results for the case of $I \geq 3$. First, we obtain properties of SPEs of regime $\widehat{R}^e$ above that parallel Proposition 1 for the two-agent case.

Proposition 2 *Consider any SPE σ of regime $\widehat{R}^e$. Then, one of the following must hold:*

(a) *Each agent $i \in I$ announces 0 for sure at any $(\mathbf{h}, d) \in \mathbf{H}^\infty \times D_2$ on the equilibrium path, and $\pi_i^{\mathbf{h}}(\sigma, \widehat{R}^e) = v_i(f)$ for any $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t$ on the equilibrium path.*

(b) *Agents k and l mix between 1 and 2 at some $d \in D_2$ in period 1 on the equilibrium path and, for each $i \in I$, the continuation payoff at the next period is less than $v_i(f)$; hence, for each $i \in I$, $\pi_i(\sigma, \widehat{R}^e) < v_i(f)$ if δ is sufficiently large.*

Proof. See Appendix A.1. ■

Characterizing the set of equilibria of regime $\widehat{R}^e$ is more involved but essentially yield the same set of results as in the case of two agents. The key feature of our regime construction with $I \geq 3$ that extends the previous ideas with $I = 2$ is that $\widehat{R}^e$ treats two (arbitrary but fixed) agents asymmetrically. Let us offer a brief sketch of our arguments for Proposition 2.

First, suppose that the players choose pure strategies over integers. We want to show that in this case the agents must all play 0. On the one hand, note that when either of the two selected agents k and l announces a positive integer the integer choice of any other agent does not matter at all (Rules B.2 and B.3). Thus, the inequalities in (4) imply that, by the analogous arguments in the two-agent case above, k and l must report 0 in equilibrium. On the other hand, if k and l both announce 0 and another agent reports integer 1, by (6) and Rule B.4, either k or l could profitably deviate by announcing a positive integer himself.

Second, suppose that some player randomizes over integers. We want to show that in this case k and l must mix between integers 1 and 2. Suppose otherwise, so that either k or l plays 0 with positive probability. Similarly to the two-agent case above, our construction here is such that both k and l strictly prefer to announce a positive integer over 0 if there is another player (possibly other than themselves) announcing a positive integer. Therefore, if either k or l chooses 0 with positive probability then every other agent must do the same and the corresponding continuation payoff for k or l is at least $v_k(f)$ or $v_l(f)$, with the inequality being strict if another agent chooses a positive integer with positive probability. Furthermore, using (7), we can show that it is also true for any agent i other than k or l that his continuation payoff from choosing 0 is at least $v_i(f)$ if every other agent announces 0 with positive probability. Combining these observations leads to a contradiction against the assumption that the SCF is efficient in the range.

Introducing complexity considerations to regime $\widehat{R}^e$ yields that the players must always play 0 for sure in any WPEC. Given (5), again, the arguments are similar to those for the two-agent case. Also, the regime admits a Markov equilibrium in which the agents always tell the truth and announce 0. Thus, together with Proposition 2, we next obtain the following.

Theorem 2 *Suppose that $I \geq 3$ and $\delta \in \left(1 - \frac{1}{|I|+1}, 1\right)$. If an SCF f is efficient in the range and satisfies condition ϕ^*, f is payoff-repeatedly implementable in WPEC from period 2.*

Proof. See Appendix A.1. ■

As in Corollary 1 for the case of $I = 2$, we can strengthen Theorem 2 to outcome implementation by additionally invoking strong efficiency in the range. Also, we can introduce the stronger refinement notion of PEC in the same way as in the case of $I = 2$ and obtain repeated implementation from period 1.

5 Alternative Complexity Measures and Simultaneous Mechanisms

5.1 More Complete Complexity Orders and Simultaneous Mechanisms

Our notion of complexity in Definition 3 (as well as the alternative notions introduced in Definitions 7 and 8 below) is a partial order over strategies as it compares only strategies that always play the same action with those that do not. We could make such ordering more complete and obtain the same set of results with any notion of complexity that satisfies the conditions specified in Definition 3 or else. For example, take the two-agent case and consider two strategies that are identical everywhere except that in the integer part of the mechanism g^e the first strategy always announces integer 0 or 1, while the second strategy either announces integer 0, 1 or mixes between integers 0 and 1 depending on the partial history. According to the definitions discussed in the previous subsection, one cannot rank the two strategies whereas it may be argued that the first strategy is less complex than the second because the range of choices is smaller for the first. Let us present a measure of complexity capturing this idea which is a more complete order than that in Definition 3. This measure is used in Gale and Sabourian [12].

Definition 6 *For any $i \in I$ and any pair of strategies $\sigma_i, \sigma_i' \in \Sigma_i$, we say that σ_i is more complex than σ_i' if the strategies are identical everywhere except, after some partial history in some mechanism and some set of histories, σ_i' always behaves (randomizes) the same way while σ_i does not. Formally, there exist some $g' \in G$, $d' \in D$ and $\mathbf{H}' \subseteq \mathbf{H}^\infty$ with the following properties:*

(i) $\sigma_i'(\mathbf{h}, g, d) = \sigma_i(\mathbf{h}, g, d)$ *if either* $\mathbf{h} \in \mathbf{H}^\infty \backslash \mathbf{H}'$, *or* $\mathbf{h} \in \mathbf{H}'$ *and* $(g, d) \neq (g', d') \in G \times D$.

(ii) $\sigma_i'(\mathbf{h}, g', d') = \sigma_i'(\mathbf{h}', g', d')$ *for all* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}'$.

(iii) $\sigma_i(\mathbf{h}, g', d') \neq \sigma_i(\mathbf{h}', g', d')$ *for some* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}'$.

(iv) $\sigma_i(\mathbf{h}, g', d') \neq \sigma_i(\mathbf{h}', g', d')$ *for all* $\mathbf{h} \in \mathbf{H}'$ *and all* $\mathbf{h}' \notin \mathbf{H}'$.

This definition allows us to compare strategies that differ only on some set of histories $\mathbf{H}'$.[18] Clearly, if a strategy σ_i is more complex than σ_i' according to Definition 6 then it is also more complex according to Definition 3. Hence, all our characterization results hold with this new definition.

The choice of complexity measure is also related to our constructive arguments. In particular, our analysis is built on regime constructions employing sequential mechanisms since, as should be clear from our WPEC arguments above, this allows us to treat complexity of behavior regarding states and integers differently, and invoke a minimal partial order over the set of strategies as in Definition 3. If one invokes a stronger notion of complexity such as Definition 6, our results can also be obtained with constructions involving simultaneous mechanisms as long as the agents are sufficiently patient; see Section 3 of the Supplementary Material.

Finally, it is also possible to adopt complete complexity orders such as counting the number of "states of the minimal automaton" implementing the strategy (e.g. Abreu and Rubinstein [2]),[19] or the "collapsing state condition" (e.g. Binmore, Piccione and

[18] Note that condition (iv) is assumed here in order to ensure that the complexity partial order is asymmetric, as in Gale and Sabourian [12].

[19] This measure is equivalent to counting the number of "continuation strategies" induced by the strategy (e.g. Kalai and Stanford [19], Lee and Sabourian [20]).

Samuelson [3]), and obtain all our results above. To gain some intuition, consider an automaton implementing a WPEC strategy in regime R^e which randomizes over integers at some history; by condition (3), it follows that the automaton must use at least two different states in order to induce the behavior in the integer part of mechanism g^e. But then, we derive a contradiction as there exists a simpler automaton that does not affect the corresponding player's continuation payoff anywhere: the alternative automaton employs just a single state to play integers in which it announces integer 1.[20]

5.2 Cost of Recalling History

The basic idea behind our complexity measure is that if a strategy conditions its actions less on what happened in the "past" than another strategy then the strategy is simpler than the other. Definition 3 above captures this by saying that, whenever playing some mechanism, a strategy that at every date t responds identically to some partial history d, independently of the previous history of play $\mathbf{h}$ before t, is less complex than one that responds differently to the same partial history d but is identical everywhere else.

According to this definition, the "past" that matters for complexity of a strategy in any given period is not what happened within the period but the play that precedes it. Thus, a simple strategy may still announce different messages at different partial histories. An intuitive justification for such an asymmetric treatment of history of outcomes before and within a period in our definition of complexity is the following: there is a memory cost for recalling history of actions before the current period whereas the partial history within the period is just some stimulus that involves no cost of recalling. For example, in repeated interactions, substantial time lags often exist between periods so that costly memory is needed to condition the current action on the play of previous periods while the delay between receiving information about the partial history and taking an action is insignificant. Since our regimes involve two-stage sequential mechanisms in each period, this justification of Definition 3 also means that the time lag between players' turns across the two stages of each period are inconsequential, or at least less important than the distance between two periods, in terms of complexity of a strategy.

Such sharp asymmetric treatment, however, may not always be reasonable and one

[20]Note that the transition from this single state is not an issue since the regime activates a permanent sequence of dictatorships and constant outcomes, and the game effectively ends, once a positive integer is announced.

may want to ask how robust our results are to a less stark treatment of the history before the period and the partial history within it. We offer two extensions for the strategies in our regimes R^e and $\widehat{R}^e$. One possibility is to differentiate the partial histories that occur in the two stages of the sequential mechanism on the grounds that there is a significant time lag between stages. More specifically, we can assume that in each period, while nature's move θ is simply some stimulus that is known at no cost when the players are asked to announce a state at the beginning of the first stage, recalling partial history that occurred in the first stage when making integer announcements in the second stage is costly, due to the significant time lag between stages. In the presence of such costs, we need to modify our previous definition of complexity so that conditioning behavior on the partial history of play in the first stage of the mechanism is also more complex than not doing so. A complexity measure that reflects this idea is introduced below.

Definition 7 *For any $i \in I$ and any pair of strategies $\sigma_i, \sigma_i' \in \Sigma_i$, we say that σ_i is more complex than σ_i' if one of the following holds:*

(i) There exists $d' \in D_1$ such that the three conditions in Definition 3 are satisfied.

(ii) $\sigma_i'(\mathbf{h}, d) = \sigma_i(\mathbf{h}, d)$ *for all* $\mathbf{h} \in \mathbf{H}^\infty$ *and all* $d \in D_1$,

$\sigma_i'(\mathbf{h}, d) = \sigma_i'(\mathbf{h}', d')$ *for all* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ *and all* $d, d' \in D_2$, *and*

$\sigma_i(\mathbf{h}, d) \neq \sigma_i(\mathbf{h}', d')$ *for some* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ *and some* $d, d' \in D_2$.

Notice that the continuation payoff at the beginning of each first stage before the players announce states may depend on the partial history at that stage (the payoff in the current stage depends on the true state θ), whereas the continuation payoff at the beginning of the second stage does not depend on the partial history at the the second stage. Assuming that players have the knowledge of current and future payoffs would therefore mean that one should treat nature's move differently from other partial histories for its payoff relevance. This provides another justification for the different treatments in Definition 7 of the partial histories across the two stages .

It is straightforward to verify that our WPEC results above are not affected in any way by imposing Definition 7 instead. In particular, recall from the proof of Lemma 5 that if a WPEC were to involve mixing over integers at any history, deviating to another strategy that is otherwise the same except that it always announces integer 1 at every information set in the second stage would always generate the same payoff. This deviating strategy is simpler than the equilibrium strategy according the new definition as well.

Yet another approach to the treatment of the partial histories would be to treat information at any decision node identically and 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.

Definition 8 *For any $i \in I$ and any pair of strategies $\sigma_i, \sigma_i' \in \Sigma_i$, we say that σ_i is more complex than σ_i' if there exists $l \in \{1, 2\}$ with the following properties:*

(i) $\sigma_i'(\mathbf{h}, d) = \sigma_i(\mathbf{h}, d)$ *for all* $\mathbf{h} \in \mathbf{H}^\infty$ *and all* $d \notin D_l$*.*

(ii) $\sigma_i'(\mathbf{h}, d) = \sigma_i'(\mathbf{h}', d')$ *for all* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ *and all* $d, d' \in D_l$*.*

(iii) $\sigma_i(\mathbf{h}, d) \neq \sigma_i(\mathbf{h}', d')$ *for some* $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$ *and some* $d, d' \in D_l$*.*

Under this definition, complexity-averse players may want to economize even on the responsiveness of their behavior to nature's move θ at the beginning of each period. As we have already argued, this approach may not be as plausible as Definition 3 or 7 because nature's moves are payoff-relevant and therefore assuming that it is costly to recall them amounts to assuming that the knowledge of continuation payoffs are costly to acquire.

With Definition 8, our characterization of WPECs of the regimes R^e for $I = 2$ and $\widehat{R}^e$ for $I \geq 3$ actually remain valid via identical arguments. However, these regimes may not admit an equilibrium since the players may find it beneficial to economize on the complexity of their reports in the first stage by make unconditional state announcements. To see this, consider the type of strategies that we have used to obtain existence in which the true state is always announced. Here, a unilateral deviation from truth-telling leads to either one-period outcome according to self-selection when $I = 2$ or no change in the outcome when $I \geq 3$. Thus, in the latter case, deviating to always announcing the same state may reduce complexity cost without affecting payoffs; in the former case, each player faces the same incentive if the self-selection condition holds with equality.

In order to obtain our WPEC results on the basis of this alternative complexity measure, we therefore need to have equilibria where such a deviation generates a *strict* reduction in the continuation payoff. This would be possible if deviations from truth-telling induced a strict one-period punishment. With $I = 2$, this would be the case if the self-selection condition held with strict inequalities.

With more than two players, suppose that there exists a "bad outcome" $\tilde{a} \in A$ such that $u_i(\tilde{a}, \theta) < u_i(f((\theta), \theta)$ for all i and θ (e.g. zero consumption in a market; see Moore and Repullo [32]). Then, we could alter $\widehat{R}^e$ above and obtain our WPECs implementation results with Definition 8 by simply modifying the outcome function of its stage mechanism as follows: whenever all agents announce the same state θ in Stage 1, $f(\theta)$ is implemented while, otherwise, the bad outcome $\tilde{a}$ is implemented. Such a modification does not change the decision problems regarding integers and hence all our characterization results remain unaffected; the Markov strategy profile in which the true state and integer 0 are always announced is a WPEC with Definition 8 because any deviation from truth-telling induces the bad outcome.

If the agents are sufficiently patient, another way to obtain the same results with Definition 8 is to modify the regime in a way that strict punishment for deviation from truth-telling arises in the continuation game, rather than from one-shot incentives via strict self-selection (in the range) or bad outcome. In the Supplementary Material (Section 2), we formally provide such a construction for the two-agent case.[21]

6 Relation to the Existing Literature

LS consider the same repeated implementation problem as this paper and show how to repeatedly implement an SCF f satisfying efficiency in the range in Nash equilibrium, using the following regime.[22] In a given period, the agents simultaneously report a state and a non-negative integer. The profile of state announcements determines the outcome implemented in the current period, and the profile of integer announcements induces the continuation regime as follows: (I) If every player announces zero, the same mechanism is repeated next period; (II) if all but one agent, say i, announces zero integer, the continuation regime S^i is such that i obtains a payoff exactly equal to $v_i(f)$, his expected payoff from the target SCF f; and (III) in every other case the agent announcing the highest integer becomes dictator forever thereafter. The last transition rule corresponds

[21] This is the same construction mentioned at the beginning of Section 3 for deriving the results there without self-selection.

[22] Similar constructions are used by Lee and Sabourian [23] who consider repeated implementation with incomplete information.

to the so-called integer game arguments, and it ensures that there is no, pure or mixed, equilibrium in which any player reports a positive integer on the outcome path. Given this, the second part of transition offers each i an opportunity to pre-emptively deviate should his continuation payoff fall below $v_i(f)$.

The constructions proposed in this paper have a similar structure in that the agents are also asked to report two messages (albeit sequentially), where the former dictates the current implementation and the latter the continuation play. In contrast to LS, the support of the second message is finite, consisting of at most three numbers 0, 1 and 2. With such a structure the integer arguments cannot be applied to exclude undesirable equilibria. Hence, we replace transition rule (III) in LS by transitions to continuation regimes described by $\{X(t)\}_{t=1}^{\infty}$ and Y, while retaining the opportunities of pre-emptive deviation that ensure at least the target payoff $v_i(f)$ for each player i. The continuation regimes $\{X(t)\}_{t=1}^{\infty}$ and Y are designed precisely to ensure that (a) every pure strategy equilibrium entails announcements of 0, and (b) mixing occurs only between 1 and 2; furthermore, any equilibrium mixing would involve non-stationary behavior that can be replaced by a simpler strategy without any payoff consequence on- or off-path.

As in the constructions of LS, when there is an odd-one-out i reporting a non-zero integer, our regimes activate S^i in the next period. However, to achieve (a) and (b) above, each continuation regime S^i in our setup not only delivers payoff $v_i(f)$ to agent i but also a payoff strictly worse than $v_j(f)$ to every other agent j. With two agents, there is almost no additional loss of generality relative to LS in achieving this, since condition ϕ in Lemma 2 would be true generically if f instead satisfied condition ω in LS. With three or more agents, we invoke condition ϕ^* (see Lemma 6).

The payoff structure of S^i in this paper which keeps one agent's payoff fixed at the target level while lowering the payoffs of others strictly below their targets has a flavor of the *separability* assumption invoked by Jackson, Palfrey and Srivastava [17] and Sjöström [39] in one-shot implementation. Separability involves the following two elements: (i) there exists some *bad outcome* $a^* \in A$ such that $u_i(a^*, \theta) < u_i(f(\theta), \theta)$ for any i and θ, and (ii) for any $a \in A$, i and θ, there exists some $a' \in A$ such that $u_i(a', \theta) = u_i(a, \theta)$ and $u_j(a', \theta') = u_j(a^*, \theta')$ for all $j \neq i$ and θ'. Separability is satisfied in economic environments like an allocation problem without externalities, where zero consumption vector would satisfy (i) and a consumption vector with only a single positive element would correspond to (ii).

The second part of separability shares some similarity to the payoff structure of S^i in our paper. However, note that we require S^i to lower another agent j's payoff below $v_j(f)$, and not necessarily to a "worst" level, as stipulated by outcome a^* in the separable case. Also, in our repeated setup, the ability to construct regime S^i with such a payoff structure is derived from condition ϕ for the case of $I = 2$, or from condition ϕ^* when $I \geq 3$. The first part of conditions ϕ and ϕ^* are in turn implied by part (i) of separability (the former, for example, only assume an outcome whose *expected* payoff is less than the target payoff for each agent). The second part of our conditions concerns the payoff that a player derives from the SCF as compared to the case when another player is dictator.

Despite the above similarity, our objectives and arguments differ from those of Jackson, Palfrey and Srivastava [17] and Sjöström [39]. In particular, our paper is concerned with the problem of repeated implementation with randomly evolving preferences, and we adopt a novel approach that invokes individuals' bounded rationality to deal with unwanted mixed strategy equilibria and enhance the scope of implementability beyond LS. The objective of Jackson, Palfrey and Srivastava [17] and Sjöström [39] is to show that in separable one-shot environments any social choice function can be implemented in undominated Nash equilibrium with finite mechanisms.[23] In their setup, this is done for example by requiring each agent to report either a profile of states or outcomes (but not both). Since the mechanism is such that each agent guarantees implementation of a^* from announcing outcomes it follows from part (i) of separability that announcing outcomes is dominated by announcing states; the other part of separability is used to make truth-telling the only undominated choice of state announcement.[24]

Could we exploit a similar separability argument to derive that announcing zero is the only undominated strategy in our constructions? The difficulty in our setup is that our arguments are applied to continuation payoffs; in particular, we do not know *a priori* the continuation payoffs when the agents all announce zero. This is pinned down only after the players' equilibrium integer choices have been derived.

[23]Refinement arguments allow these authors to obtain their one-shot implementation results without Maskin monotonicity. Other approaches in this direction include using dynamic mechanisms and subgame perfection (e.g. Moore and Repullo [31]) and virtual implementation (e.g. Abreu and Matsushima [1]).

[24]In a recent paper, Mezzetti and Renou [28] identify sufficient conditions for mixed Nash implementation with finite mechanisms in separable environments. In their construction, "odd-one-out" leads to deviating opportunities similar to ours but they assume that such opportunities can be constructed using outcomes that are themselves desired social choices. See also Goltsman [14].

Mezzetti and Renou [29] characterize the repeated Nash implementation setup of LS in terms an alternative property that they call *dynamic monotonicity.* Their characterization applies to both infinitely and finitely repeated setups. The sufficiency results of Mezzetti and Renou [29] also use mechanisms that involve unbounded integers and are therefore open to the same criticisms against LS that the present paper addresses. In this paper, we do not consider dynamically monotone SCFs and consider only infinitely repeated implementation of efficient SCFs via bounded mechanisms. In the infinite setup with sufficiently patient agents, there is not much of a loss in this approach as dynamic monotonicity is generically equivalent to the notion of efficiency invoked by both LS and the current paper.[25] It may however be possible to extend the ideas and techniques developed in this paper to the case of finitely repeated, or even one-shot, implementation of dynamically monotone social choices.

While a full investigation is beyond the scope of this paper, there are several issues to note in this regard. On the one hand, some of our key arguments, in particular, constructing $\{S^i\}_{i\in I}$, $\{X(t)\}_{t=1}^{\infty}$ and Y, can be exactly replicated in any finite setup if the planner can randomly choose mechanisms (see footnote 11). The main idea of complexity definition, that stationary behavior is simpler than otherwise, can be applied to any finitely repeated setup, as well as any (within-period) extensive form mechanisms. On the other hand, repeatedly implementing a dynamically monotone SCFs via finite mechanisms may not be so straightforward as the regimes proposed by Mezzetti and Renou [29] are more complex than those used in LS.[26] For instance, they ask players to report their continuation payoffs in addition to states, outcomes and unbounded integers. Also, our arguments cannot be directly applied to the one-shot environment since the ideas of pre-emptive deviation to bound equilibrium payoffs and using complexity to eliminate non-stationary behavior are based on multiple rounds of play.

Finally, our paper contributes to a growing body of papers that study the implications of individual bounded rationality for implementation in static environments. This literature has so far introduced agents who make mistakes (Eliaz [10]), are partially honest (Matsushima [27], Dutta and Sen [9]), choose best responses adaptively (Cabrales and Serrano [4]) and have context-dependent preferences (Saran [37], de Clippel [7]), among others. Complexity considerations in mechanism design have been analyzed in a

[25]Non-generically, the concepts are different depending on the precise details of conditions.

[26]Mezzetti and Renou [29] actually obtain their characterization by restricting attention to pure strategies as well as allowing for random choice of mechanisms, in contrast to LS and this paper.

principal-agent context by Glazer and Rubinstein [13]. Their model of bounded rationality is different from ours, but interestingly, their results turn out to offer a broadly similar message: faced with boundedly rational agent, complexity of the mechanism may actually help the principal's cause.

7 Conclusion

In summary, this paper explores how to repeatedly implement an efficient social choice function when the agents have a preference for less complex strategies at the margin. We identify some minor conditions under which such implementation is achieved with using only finite mechanisms and allowing for mixed strategies. Compared with LS, when faced with complexity-averse agents, the freedom to set different mechanisms at different histories gives the planner an additional leverage to deter undesirable (mixing) behavior even if the mechanisms themselves are simple. Another feature of our constructions is that all mixed equilibria are strictly Pareto-dominated by pure equilibria which attain the desired outcome paths.

The key feature in our regime constructions driving the WPEC results is the non-stationarity of continuation regimes $\{X(t)\}_{t=1}^{\infty}$, activated if two players announce the same positive integer 1 or 2 at each period t. Although each mechanism in our regimes is simple and does not employ integer games, one may suggest that the non-stationary sequence of regimes poses another kind of implausible design. Our response is two-fold. First, the criticism leveled at integer games is not about the implausibility of unboundedness *per se* but rather about the fact that integer games kill off unwanted equilibria by strategies that are themselves dominated. Our constructions achieve full implementation without appealing to such arguments. Second, as specified in condition (3) above, our WPEC results require $X(t)$ to be distinct at just two dates. Thus, the degree of non-stationarity or complexity in our regime constructions needs not be overly demanding.

Regarding the latter point, however, it is worth pointing out that greater non-stationarity in $\{X(t)\}_{t=1}^{\infty}$ also means more complex mixed strategy SPEs, and therefore, strengthens the agents' incentives to economize on complexity cost associated with such behavior. In general, the planner could even write $X(\cdot)$ as a function of the entire (publicly observable) history instead of just its date.

Another related issue that can be raised against our complexity analysis is why we consider a preference for less complexity 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. In a similar vein, for the planner the cost of implementing a more complex regime could be warranted if it led to better implementation results.

Our notions of complexity apply to any regimes defined with extensive form mechanisms and therefore would be used to address the necessary conditions for repeated implementation in WPEC or other related solution concepts. LS establish the following result: if an SCF is repeatedly implementable via some regime but the SCF is strictly Pareto-dominated by another SCF whose range belongs to that of the former, then the regime admits another equilibrium that achieves the superior payoffs by triggering the original equilibrium with lower payoffs upon any deviation. An interesting future research question would be to see whether a similar argument can be extended to the notion of WPEC.

A more broad lesson from our analysis is that complexity may help the planner's cause: by deliberately constructing a complex institution, the planner may guide the agents to adopt desired strategies if they are simple while other equilibria involve complex behavior. In our particular exploration, the agents are assumed to have preference for less complex strategies at the margin, where complexity is concerned with the degree of history-dependence of behavior. The complexity of regime that exploited these traits was manifested in the non-stationarity of the sequence of mechanisms enforced. Indeed, one can find many real world cases of complex institutions that have survived the test of time (for an illuminating example, see the voting protocol for electing the Doge of Venice between 1268 and 1797; Mowbray and Gollmann [33]).[27] A potentially fruitful direction for future research would be to uncover other relationships between complexity and mechanism/institution design beyond the premises of this paper.

[27] The authors thank Romans Pancs for the example.

A Appendix

A.1 Omitted Proofs

Proof of Lemma 4

For each $i = 1, 2$, let Π_i denote i's continuation payoff at the next period if both agents announce zero at the given history $(\mathbf{h}, d)$. 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.

Case 1: No player randomizes.

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 A.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 A.3(iii)). This is a contradiction.

Subcase 1B: $z_i > 0$ and $z_j > 0$.

The continuation regime is either X, $X(t)$ or Y (Rule A.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 contradiction.

Thus, both players choose 0 for sure at this history, and g^e must be the mechanism at the next period. We next show that $\Pi_i = v_i(f)$ for all i. For this, suppose first that $\Pi_i < v_i(f)$ for some i. But then, by Rule A.2, i could deviate at this history $(\mathbf{h}, d)$ by announcing a positive integer and obtain a continuation payoff equal to $v_i(f)$, a contradiction. It therefore follows that $\Pi_i \geq v_i(f)$ for all i. Then, suppose that $\Pi_i > v_i(f)$ for some i. But, since regime R^e only employs outcomes from the set $f(\Theta)$, and since f is efficient in the range, it must be that $\Pi_j < v_j(f)$ for $j \neq i$. This contradicts that $\Pi_i \geq v_i(f)$ for all i.

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 A.2, his continuation payoff is w_i^j. If he announces 1, by Rules A.3(i) and A.3(iii), the continuation payoff is $x_i > w_i^j$ or $y_i > w_i^j$.

Claim 2: Suppose that agent i announces 0 with positive probability. Then the other agent j must also announce 0 with 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 positive probability, it must then be that j also chooses 0 with positive probability. Hence, we obtain that $\Pi_i \geq v_i(f)$ with the inequality being strict if j plays a positive integer with positive probability.

We now show that, in this Case 2, both players choose a positive integer for sure. To show this suppose otherwise; then some player chooses 0 with positive probability. By Claim 2, the other player must also play 0 with 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 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. Note also that regime R^e involves only outcomes in the range of f. Therefore, since f is efficient in the range, we have 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 4, it suffices to show that any mixing over integers in equilibrium must occur in period 1. Suppose not; so, there exists an SPE σ such that, for some $t > 1$, there exist $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_2$ that occur on the equilibrium path at which the players are mixing over integers.

First, note that by Lemma 4 the players must have all announced 0 for sure in the previous period and, moreover, $\pi_i^{\mathbf{h}^t}(\sigma, R^e) = v_i(f)$ for all $i = 1, 2$. Second, for any $d' \in D_2$, we can apply similar reasoning to show that $\pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}}(\sigma, R^e) = v_i(f)$ for all i if $\underset{\sim}{z} = (0, 0)$ and $(\mathbf{h}^t, d', \underset{\sim}{z})$ occurs on the equilibrium path.

Next, let $r(d', \underset{\sim}{z})$ denote the probability of $(d', \underset{\sim}{z}) \in D_2 \times \mathcal{Z}^2$ occurring at $\mathbf{h}^t$ under σ, 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

$$\pi_i^{\mathbf{h}^t}(\sigma, R^e) = \sum_{(d', \underset{\sim}{z}) \in D_2 \times \mathcal{Z}^2} r(d', \underset{\sim}{z}) \left[(1 - \delta) u_i(a^{\mathbf{h}^t, d'}, d') + \delta \pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}}(\sigma, R^e) \right] = v_i(f). \tag{8}$$

Lemma 4 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 $\pi_i^{\mathbf{h}^t, d', \underset{\sim}{z}} < v_i(f)$ for all i. Thus, since we assume that mixing over positive integers occurs after d, it follows from (8) 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 in the range.

Proof of Lemma 5

Suppose not. Then, by Lemma 4, there exists a WPEC, σ, such that, at some t, $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_2$, the two agents play integer 1 or 2 for sure and each i plays 1 with probability $\frac{x_i(t) - y_i}{x_i + x_i(t) - 2y_i}$. Furthermore, by construction, there exist t' and t'' such that $x(t') \neq x(t'')$ and, therefore, it follows that, for all i, we have either $\sigma_i(\mathbf{h}^t, d) \neq \sigma_i(\mathbf{h}^{t'}, d)$ for some $\mathbf{h}^{t'} \in \mathbf{H}^{t'}$, or $\sigma_i(\mathbf{h}^t, d) \neq \sigma_i(\mathbf{h}^{t''}, d)$ for some $\mathbf{h}^{t''} \in \mathbf{H}^{t''}$.

Now, consider any $i = 1, 2$ deviating to another strategy σ_i' that is identical to the equilibrium strategy σ_i except that, for all $\mathbf{h} \in \mathbf{H}^\infty$, $\sigma_i'(\mathbf{h}, d)$ prescribes announcing 1 with probability 1. Since σ_i' is less complex than σ_i, we obtain a contradiction by showing that $\pi_i^{\mathbf{h}}(\sigma_i', \sigma_{-i}, R^e) = \pi_i^{\mathbf{h}}(\sigma, R^e)$ for all $\mathbf{h} \in \mathbf{H}^\infty$. To do so, fix any history $\mathbf{h}$ and suppose that the given partial history d occurs after $\mathbf{h}$. Given Lemma 4, there are two cases to consider at $(\mathbf{h}, d)$.

First, suppose that j plays 0 for sure. Then, by part (a) of Lemma 4, i also plays 0 for sure and obtains a continuation payoff equal to $v_i(f)$ in equilibrium. By Rule B.2 of the regime, the deviation also induces the same continuation payoff $v_i(f)$. Second, suppose that j is mixing. Then, by part (b) of Lemma 4, j mixes between 1 and 2 and i is also indifferent between choosing 1 and 2.

Proofs of Proposition 2 and Theorem 2

These results are proved by the next three lemmas.

Lemma 7 *Regime $\widehat{R}^e$ admits an SPE, σ^*, in Markov strategies such that, for any t, $\mathbf{h} \in \mathbf{H}^t$ and $\theta \in \Theta$, (i) $g^{\mathbf{h}}(\sigma^*, R^e) = g^e$ and (ii) $A^{\mathbf{h},\theta}(\sigma^*, R^e) = \{f(\theta)\}$.*

Proof. The proof is similar to that of Lemma 3. ∎

Lemma 8 *Consider any SPE of regime $\widehat{R}^e$. Fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d \in D_2$. Then, one of the following must hold at $(\mathbf{h}, d)$:*

(a) Each $i \in I$ announces 0 for sure and his continuation payoff at the next period is equal to $v_i(f)$.

(b) Each $i \in \{k, l\}$ announces 1 or 2 for sure, with the probability of choosing 1 equal to $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i} \in (0,1)$. Furthermore, for all $j \in I$, the continuation payoff at the next period is less than $v_j(f)$.

Proof. For each i, let Π_i denote i's continuation payoff at the next period if all agents announce zero at the fixed history $(\mathbf{h}, d) \in \mathbf{H}^t \times D_2$. 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 shall prove the claim by considering each case separately.

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

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. We derive contradiction by considering the following subcases.

<u>Subcase 1A</u>: $z_k > 0$ and $z_l = 0$, or $z_k = 0$ and $z_l > 0$.

Consider the former case; the latter case can be handled analogously. The continuation regime at the next period is S^k (Rule B.2). But then, since $y_l > w_l^k$ by (4), l can profitably deviate by choosing a strategy identical to the equilibrium strategy except that it announces the positive integer other than z_l at this history, which activates the continuation regime Y instead of S^k (Rule B.3(iii)). This is a contradiction.

<u>Subcase 1B</u>: $z_k > 0$ and $z_l > 0$.

The continuation regime is either $X, X(t)$ or Y (Rule B.3). Suppose that it is X or $X(t)$. By (4), we have $y_l > x_l(t')$ for all t'. But then, l can profitably deviate by choosing a strategy identical to the equilibrium strategy except that it announces the positive integer other than z_l at this history, which activates Y (Rule B.3(iii)). This is a contradiction. Similarly, since $x_k > y_k$ by (4), when the continuation regime is Y, player k can profitably deviate and we obtain a similar contradiction.

Subcase 1C: For some $C \subseteq I\backslash\{k,l\}$, $z_i = 1$ for all $i \in C$ and $z_i = 0$ for all $i \notin C$.

The continuation regime is S^C (Rule B.4). By (6), we have $w_j^j > w_j^C$ for $j \in \{k,l\}$. But then, j can profitably deviate by choosing a strategy identical to the equilibrium strategy except that it announces a positive integer at this history, which activates S^j (Rule B.2).

Thus, all players choose 0 for sure at this history, and $\hat{g}^e$ must be the mechanism at the next period. We next show that $\Pi_i = v_i(f)$ for all i. Suppose not. First, suppose that $\Pi_i < v_i(f)$ for some i. But then, i could deviate at this history $(\mathbf{h}, d)$ by announcing a positive integer and obtain a continuation payoff equal to $v_i(f)$, a contradiction. It therefore follows that $\Pi_i \geq v_i(f)$ for all i. In the continuation game, we either have implementation of outcomes in the range of f or end up activating continuation regimes from $\{S^i\}_{i=1}^I \cup \{X(t)\}_{t=1}^\infty \cup Y$ whose payoffs are all Pareto-dominated by $v(f)$. Therefore, if f is efficient in the range $\Pi_i = v_i(f)$ for all i.

Case 2: Some player randomizes.

We proceed by establishing the following claims.

Claim 1: For each agent k or l, the continuation payoff (at the next period) from announcing 1 is greater than that from announcing 0, if there exists another player announcing a positive integer.

Proof of Claim 1. Consider k and any $z_{-k} \neq (0,\ldots,0)$. The other case for l can be proved identically. There are two possibilities:

First, suppose that $z_l > 0$. In this case, if k announces zero, by Rule B.2, his continuation payoff is w_k^l. If he announces 1, by Rules B.3(i) and B.3(iii), the continuation payoff is x_k or y_k. But, by (4), we have $x_k > y_k > w_k^l$.

Second, suppose that $z_l = 0$. In this case, since $z_{-k} \neq (0,\ldots,0)$, there must exist a non-empty set $C \subseteq I\backslash\{k,l\}$ such that $z_i = 1$ for all $i \in C$ and $z_i = 0$ for all $i \in I\backslash\{C \cup k\}$.

Then if k announces 0, by Rule B.4, his continuation payoff is w_k^C, whereas if he announces 1, by Rule B.2, the continuation payoff is w_k^k. But, by (6), we have $w_k^k > w_k^C$.

Claim 2: If agent k or l announces zero with positive probability, then every other agent must also announce zero with positive probability.

Proof of Claim 2. Suppose not. Then, suppose that k plays 0 with positive probability but some $i \neq k$ chooses 0 with zero probability. (The other case for l can be proved identically.) But then, by Claim 1, the latter implies that k obtains a lower continuation payoff from choosing 0 than from choosing 1. This contradicts the supposition that k chooses 0 with positive probability.

Claim 3. Suppose that some agent $i \in \{k, l\}$ announces 0 with positive probability. Then, $\Pi_i \geq v_i(f)$ with this inequality being strict if some other agent announces a positive integer with positive probability.

Proof of Claim 3. For any agent $i \in \{k, l\}$, by Claim 1, playing 1 must always yield a higher continuation payoff i than playing 0, except when all other agents play 0. Since i plays 0 with positive probability, the following must hold:

(i) If all others announce 0, i's continuation payoff when he announces 0 must be no less than that he obtains when he announces 1, i.e. $\Pi_i \geq v_i(f)$.

(ii) If some other player attaches a positive weight to a positive integer, i's continuation payoff must be greater when he chooses 0 than when he chooses 1 in the case in which all others choose 0, i.e. $\Pi_i > v_i(f)$.

Claim 4: For each agent $i \in I \backslash \{k, l\}$, the continuation payoff from announcing zero is no greater than that from announcing 1, if there exists another player announcing a positive integer.

Proof of Claim 4. For each $i \in I \backslash \{k, l\}$, the continuation payoff is independent of his choice if $z_k > 0$ or $z_l > 0$. So, suppose that $z_k = z_l = 0$. Then if i chooses 1 he obtains w_i^C, for some $C \in I \backslash \{k, l\}$ such that $i \notin C$, while he obtains $w_i^{C \cup \{i\}}$ from choosing 1. By (7), $w_i^C \leq w_i^{C \cup \{i\}}$. Thus, the claim follows.

Claim 5. For each agent $i \in I \backslash \{k, l\}$, $\Pi_i \geq v_i(f)$ if all players announce 0 with positive probability.

Proof of Claim 5. Note that, if $z_j = 0$ for all $j \neq i$, i obtains Π_i from choosing 0 and obtains $v_i(f)$ from choosing 1. Since, by assumption, i announces 0 with positive probability, announcing 0 must be weakly preferred to either positive integer. The claim then follows immediately from the previous claim.

Claim 6. Both k and l choose a positive integer for sure.

Proof of Claim 6. Suppose otherwise; then some $i \in \{k,l\}$ chooses 0 with positive probability. Then, by Claim 2, every other agent must play 0 with positive probability. By Claims 3 and 5, this implies that $\Pi_j \geq v_j(f)$ for every j. Moreover, since in this case there is randomization, some player must be choosing a positive integer with positive probability. Then, by appealing to Claim 3 once again, we must also have that at least one of the inequalities $\Pi_k \geq v_k(f)$ or $\Pi_l \geq v_l(f)$ is strict. Since f is efficient in the range, this is a contradiction.

Claim 7. Both k and l choose each of the integers 1 and 2 with positive probability.

Proof of Claim 7. Suppose not; then by the previous claim one of either k or l must choose one of the positive integers for sure. But then, (4) implies that the other must also do the same. But, by applying (4) once again, this induces a contradiction (the argument is exactly the same as in Subcase 1B of Case 1 with no randomization).

Given the last two claims, simple computation verifies that both agents k and l must be playing 1 with unique probability as in the statement. The continuation payoffs, for each $i \in I$, when k or l chooses a positive integer are $x_i, x_i(t)$ or y. Moreover, by (4), each of these payoffs is less than $v_i(f)$. Therefore, it follows that, in this case, the continuation payoff at the next period must be less than $v_i(f)$ for all i. ■

Lemma 9 *Fix any WPEC of regime $\widehat{R}^e$. Also, fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d \in D_2$. Then, every agent announces zero for sure at this history.*

Proof. Suppose not. Then, by Lemma 8 above, there exists a WPEC, σ, such that, at some t, $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_2$, each $i \in \{k,l\}$ plays integer 1 or 2 for sure and integer 1 is chosen with probability $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i}$. Furthermore, by construction, there exist t' and t'' such that $x_k(t') \neq x_k(t'')$ and $x_l(t') \neq x_l(t'')$. Thus, it follows that, for each $i \in \{k,l\}$, we have either $\sigma_i(\mathbf{h}^t, d) \neq \sigma_i(\mathbf{h}^{t'}, d)$ for some $\mathbf{h}^{t'} \in \mathbf{H}^{t'}$, or $\sigma_i(\mathbf{h}^t, d) \neq \sigma_i(\mathbf{h}^{t''}, d)$ for some $\mathbf{h}^{t''} \in \mathbf{H}^{t''}$.

Now, consider any $i \in \{k,l\}$ deviating to another strategy σ_i' that is identical to the equilibrium strategy σ_i except that, for all $\mathbf{h} \in \mathbf{H}^\infty$, $\sigma_i'(\mathbf{h}, d)$ prescribes announcing 1 for sure. Since σ_i' is less complex than σ_i, we obtain a contradiction by showing that $\pi_i^{\mathbf{h}}(\sigma_i', \sigma_{-i}, R^e) = \pi_i^{\mathbf{h}}(\sigma, R^e)$ for all $\mathbf{h} \in \mathbf{H}^\infty$. To this end, fix any history $\mathbf{h}$ and suppose that the given partial history d occurs at $\mathbf{h}$. Given Lemma 8, there are two cases to consider at $(\mathbf{h}, d)$.

First, suppose that every agent plays 0 for sure. Then, by part (a) of Lemma 8, i also plays 0 for sure and obtains a continuation payoff equal to $v_i(f)$ in equilibrium. By Rule B.2 of the regime, the deviation also induces the same continuation payoff $v_i(f)$. Otherwise, by part (b) of Lemma 8, agents k and l mix between 1 and 2; thus, i is indifferent. ■

A.2 PEC and Period 1

Lemma 10 *Every PEC, σ, of regime R^e is Markov: for all $i \in I$, $\sigma_i(\mathbf{h}', d) = \sigma_i(\mathbf{h}'', d)$ for all $\mathbf{h}', \mathbf{h}'' \in \mathbf{H}^\infty$ and all $d \in D$.*

Proof. Suppose not. Then, there exists some PEC, σ, such that $\sigma_i(\mathbf{h}', d) \neq \sigma_i(\mathbf{h}'', d')$ for some $i, \mathbf{h}', \mathbf{h}''$ and d'. By Lemma 5, we know that $d' \in D_1$; let $d' = \theta$.

Consider i deviating to another strategy σ_i' that is identical to σ_i except that, irrespective of past history, (i) whenever $d = \theta$, it does what the equilibrium strategy does in period 1 after the given partial history, and (ii) whenever $d = (\theta, \underset{\sim}{\theta})$ for any $\underset{\sim}{\theta} \in \Theta^2$, i.e. any Stage 2 partial history following realization of the given state $\tilde{\theta}$, it announces 1. Formally, for all $\mathbf{h} \in \mathbf{H}^\infty$ and $\underset{\sim}{\theta} \in \Theta^2$, $\sigma_i'(\mathbf{h}, \tilde{\theta}) = \sigma_i(\emptyset, \theta)$ and $\sigma_i'(\mathbf{h}, \theta, \underset{\sim}{\theta}) = 1$ (where the latter slightly abuses notation to denote a pure strategy).

Clearly, σ_i' is less complex than σ_i. Furthermore, the deviation alters neither i's one-period payoff in period 1 at $\tilde{\theta}$ nor, by Rule A.2 of the regime, and since the opponent player's equilibrium strategy announces 0, his continuation payoff as of period 2 on the equilibrium path. This contradicts the assumption of PEC. ■

Together with Theorem 1, this lemma immediately implies the following.

Theorem 3 *Suppose that $I = 2$ and $\delta \in \left(\frac{3}{4}, 1\right)$. If an SCF f is efficient in the range, and satisfies self-selection in the range and condition ϕ, there exists a regime R such that (i) a PEC exists and (ii) every PEC σ satisfies $\pi_i^{\mathbf{h}^t}(\sigma, R) = v_i(f)$ for any $i \in I$, $t \geq 1$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$.*

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

[4] Cabrales, A. and R. Serrano (2011): "Implementation in Adaptive Better-Response Dynamics: Towards a General Theory of Bounded Rationality in Mechanisms," *Games and Economic Behavior*, 73, 360-374.

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

[6] Chatterjee, K. and H. Sabourian (2009): "Game Theory and Strategic Complexity," in *Encyclopedia of Complexity and System Science*, ed. by R. A. Meyers, Springer.

[7] de Clippel, G. (2014): "Behavioral Implementation,"forthcoming in *American Economic Review*.

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

[9] Dutta, B. and A. Sen (2012): "Nash Implementation with Partially Honest Individuals," *Games and Economic Behavior*, 74, 154-169.

[10] Eliaz, K. (2002): "Fault-Tolerant Implementation," *Review of Economic Studies*, 69, 589-610.

[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] Glazoer, J. and A. Rubinstein (2014): “Complex Questionnaires,” forthcoming in *Econometrica*.

[14] Goltsman, M. (2011): “Nash Implementation Using Simple Mechanisms without Undesirable Mixed-Strategy Equilibria,” *Mimeo*.

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

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

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

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

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

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

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

[22] Lee, J. and H. Sabourian (2011b): “Supplement to ‘Efficient Repeated Implementation’,” *Econometrica Supplemental Material*, 79, http://www.econometricsociety.org/ecta/Supmat/8859_proofs.pdf.

[23] Lee, J. and H. Sabourian (2013): “Repeated Implementation with Incomplete Information,” *Mimeo*.

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

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

[26] 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.

[27] Matsushima, H. (2008): "Role of Honesty in Full Implementation," *Journal of Economic Theory*, 139, 353-359.

[28] Mezzetti, C. and L. Renou (2012a): "Mixed Nash Implementation with Finite Mechanisms," *Mimeo.*

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

[30] Moore, J. (1992): "Implementation, Contracts, and Renegotiation in Environments with Complete Information," in *Advances in Economic Theory, Sixth World Congress, Vol. 1*, ed. by J-J. Laffont, Cambridge University Press.

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

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

[33] Mowbray, M. and D. Gollmann (2007): "Electing the Doge of Venice: Analysis of a 13th Century Protocol," *Proceedings of the 20th IEEE Computer Security Foundations Symposium*, 295-310.

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

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

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

[37] Saran, R. (2011): "Menu-Dependent Preferences and Revelation Principle," *Journal of Economic Theory*, 146, 1712-1720.

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

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

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