Repeated Implementation with Finite Mechanisms and Complexity

Jihong Lee[*]
Seoul National University

Hamid Sabourian[†]
University of Cambridge

September 2012

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

---

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

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

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. However, this latter implementation approach has been often criticized for employing abstract institutions to deliver its results, especially, the usage of unbounded integer games which rules out some undesired outcomes via an infinite chain of dominated actions (see the surveys of Moore [22], Jackson [12], Maskin and Sjöström [20], and Serrano [29], for example).

In a recent paper, Lee and Sabourian [17] extend the scope of implementation to 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 [19] 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 [26], Saijo [28]). The notion of efficiency represents a basic goal of an economic system and therefore the sufficiency results in Lee and Sabourian [17] offer strong implications. However, the results also take advantage of integer arguments to eliminate unwanted equilibria.[1]

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

[1]More recently, Mezzetti and Renou [21] have derived conditions for Nash implementation in any finite or infinite repeated implementation problems with complete information. They identify a dynamic monotonicity property which is equivalent to Maskin monotonicity in a static implementation setup, and implies efficiency in the range (to be defined below) in an infinitely repeated setup with sufficiently patient agents, as in Lee and Sabourian [17]. Integer games are also used to establish their sufficiency results.

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 [11] has shown, a finite mechanism that Nash implements a social choice function could invite unwanted mixed equilibria that strictly Pareto dominate the desired outcomes.

This paper explores the question of implementing efficient social choice functions in a repeated environment by considering only finite mechanisms. Our approach appeals to bounded rationality of the agents. In particular, 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.[2]

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 Lee and Sabourian [17] is that the mechanisms can themselves be made contingent on past histories in a way that, roughly put, each agent's individually rational repeated game 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.

In this paper, we will require each mechanism to be simple and finite but introduce a non-stationary sequence of mechanisms in terms of its path dependence, via enforcing different mechanisms at different histories. Our precise constructions generate, under minor qualifications, the following equilibrium features:

- Every pure strategy Nash, or subgame perfect, equilibrium repeatedly implements the efficient social choice function.

- Every mixed strategy Nash, or subgame perfect, equilibrium is strictly Pareto-dominated by the pure equilibria.

- Randomization can be eliminated altogether by invoking subgame perfection and an additional equilibrium refinement, based on a "small" cost associated with implementing a more complex strategy.

[2] 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.

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 are less plausible from the efficiency perspective (as alluded to by Jackson [11]); 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 complex pattern, 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 contrasts with the more standard equilibrium notion 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, for instance, the survey of Chatterjee and Sabourian [7]).

Various refinements have been explored in the one-shot implementation literature to obtain more permissive results than what is implied by Maskin monotonicity. The role of credibility in the context of extensive form mechanisms was studied by Moore and Repullo [23] and Abreu and Sen [4]. Jackson, Palfrey and Srivastava [13] and Sjöström [30] adopt the notion of undominated Nash equilibrium to derive strong implementation results using only finite mechanisms and allowing for mixed strategies. These latter objectives were also addressed by Abreu and Matsushima [1][2] who consider virtual implementation in iteratively undominated strategies. In contrast to these works, our paper is concerned with the problem of (exact) repeated implementation with randomly evolving preferences.

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. We offer some discussion on how to extend our results in Section 5 before concluding in Section 6. Appendices 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 Lee and Sabourian [17] (henceforth, 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 $u_i : A \times \Theta \rightarrow \mathbb{R}$.

An SCF f in an implementation problem $\mathcal{P}$ is a mapping $f : \Theta \rightarrow A$, 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 and, 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 $v(f') = v(f)$.*

(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 $v(f') = v(f)$.

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 mechanism is defined as $g = (M^g, \psi^g)$, where $M^g = M_1^g \times \cdots \times M_I^g$ is a cross product of message spaces and $\psi^g : M^g \to A$ is an outcome function such that $\psi^g(m) \in A$ for any message profile $m = (m_1, \ldots, m_I) \in M^g$.[3]

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

Mechanism g is *finite* if $\|M_i^g\| < \infty$ for every agent i. Let G be the set of all feasible mechanisms.

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

Given mechanism $g = (M^g, \psi^g)$, define $\mathcal{E}^g \equiv \{(g, m)\}_{m \in M^g}$, and let $\mathcal{E} = \cup_{g \in G} \mathcal{E}^g$. Then, $H^t = \mathcal{E}^{t-1}$ (the $(t-1)$-fold Cartesian product of $\mathcal{E}$) represents the set of all possible 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 $R : H^\infty \to G$. 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 $\mathbf{h}$.

Then, we can write each agent i's mixed (behavioral) strategy as a mapping $\sigma_i : \mathbf{H}^\infty \times G \times \Theta \to \triangle (\cup_{g \in G} M_i^g)$ such that $\sigma_i(\mathbf{h}, g, \theta) \in M_i^g$ for any $\mathbf{h} \in \mathbf{H}^\infty$, $g \in G$ and $\theta \in \Theta$, where $\triangle$ before a set means the set of all probability measures over the set. Let Σ_i be the set of all such strategies, and let $\Sigma \equiv \Sigma_1 \times \cdots \times \Sigma_I$. A strategy profile is denoted by $\sigma \in \Sigma$. We say that σ_i is a Markov (history-independent) strategy if and only if $\sigma_i(\mathbf{h}, g, \theta) = \sigma_i(\mathbf{h}', g, \theta)$ for any $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^\infty$, $g \in G$ and $\theta \in \Theta$. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_I)$ is Markov if and only if σ_i is Markov for each i.

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 $A^{\mathbf{h},\theta}$.

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

LS propose the following two notions of Nash repeated implementation.

Definition 2 *(a) An SCF f is payoff-repeatedly implementable in Nash 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, $t \geq \tau$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$ on the equilibrium path.*

(b) An SCF f is repeatedly implementable in Nash 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 Nash 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 establish that, with some minor qualifications, an SCF is payoff-repeatedly implementable (repeatedly implementable) in Nash equilibrium if it is efficient in the range (strictly efficient in the range); also, efficiency in the range is necessary for Nash repeated implementation when the agents are sufficiently patient.

2.3 Obtaining Target Payoffs

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

Formally, let $d(i, N)$ denote a dictatorial mechanism in which agent i is the dictator over the set of outcomes N, or simply (i, N)-dictatorship; formally, $d(i, N) = (M, \psi)$ is such that $M_i = N \subseteq A$, $M_j = \{\emptyset\}$ for all $j \neq i$ and $\psi(m) = m_i$ for all $m \in M$. Note that, if N is a singleton set, $d(i, N)$ implements a constant SCF.

For any $N \subseteq A$, let $\Upsilon^i(N, \theta) \equiv \{\arg\max_{a \in N} u_i(a, \theta)\}$ represent the set of agent i's best outcomes among the set N in state θ, and define $v_i^j(N) = \sum_{\theta \in \Theta} p(\theta) \max_{a \in \Upsilon^j(N,\theta)} u_i(a, \theta)$ as i's maximum one-period payoff if j is the dictator and always acts rationally given the outcome set N. Clearly, $v_i^i(N)$ then is i's unique maximal one-period payoff. In addition, let

$$\Phi = \left\{(i, N) \in I \times 2^A : u_j(a, \theta) = u_j(a', \theta) \ \forall a, a' \in \Upsilon^i(N, \theta) \ \forall \theta \in \Theta \ \forall j \in I\right\},$$

that is, Φ is the set of all (i, N)-dictatorships that yield unique payoff profiles.[4]

Now, for any $d(i, N)$ with $(i, N) \in \Phi$, let $v^i(N)$ denote the unique payoff profile and $W = \{v^i(N)\}_{(i,N)\in\Phi}$ denote the set of all payoff profiles from such dictatorial mechanisms. By applying the algorithm of Sorin [31] to our setup, we have the following.

Lemma 1 *Fix any $W^* \subset W$ and suppose that $\delta \in \left(1 - \frac{1}{|W^*|}, 1\right)$. Then, for any payoff profile $w \in co(W^*)$, there exists a history-independent regime that generates a unique (discounted average) payoff profile equal to w.*

Proof. Note that $v^i(N)$ is the one-period payoff profile under (i, N)-dictatorship when i acts rationally. Thus, for any $w \in co(W^*)$, there exists a history-independent sequence of dictatorial mechanisms, or payoffs belonging to W^*, such that the corresponding discounted average payoff profile equals w if $\delta > 1 - \frac{1}{|W^*|}$. See Sorin [31] (or Lemma 3.7.1 of Mailath and Samuelson [18]) for details. ■

[4]Note that with strict preferences every (i, N)-dictatorship yields unique payoffs.

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 from the material of this section. Also, 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.[5] In our main analysis below, we shall make the same assumption. Alternatively, when the agents are sufficiently patient, intertemporal incentives can be designed to support truth-telling, which we discuss in Appendix A.2.

3.1 Regime Construction

Suppose that $I = 2$, and fix an SCF f that is efficient in the range. Consider dictatorial mechanisms over the range of the SCF, $d(1, f(\Theta))$ and $d(2, f(\Theta))$. Suppose that both of these dictatorships yield unique payoffs, i.e. $(i, f(\Theta)) \in \Phi$ for all i, and also that $v^1(f(\Theta)) \neq v^2(f(\Theta))$ since otherwise the problem is trivial. For each $i = 1, 2$, it is clearly the case that $v_i^i(f(\Theta)) \geq v_i(f)$ and, hence, efficiency in the range implies that $v_i^j(f(\Theta)) \leq v_i(f)$ for $j \neq i$. Thus, by Lemma 1, for each $i, j = 1, 2$, $i \neq j$, there exists a history-independent regime S^i that alternates the two dictatorial mechanisms in a way that yields a unique (discounted average) payoff profile $w^i = (w_i^i, w_j^i)$ such that $w_i^i = v_i(f)$ if $\delta \in \left(\frac{1}{2}, 1\right)$. Since f is efficient in the range it must be that, for $j \neq i$, $w_j^i \leq w_j^j$. Moreover, this inequality is strict if $v(f)$ cannot be obtained by a convex combination of the payoffs from $d(1, f(\Theta))$ and $d(2, f(\Theta))$.

In what follows, we assume that $\delta \in \left(\frac{1}{2}, 1\right)$ and the condition below.

Condition ϕ. (i) $(i, f(\Theta)) \in \Phi$ for all i.

(ii) $v(f) \neq \gamma v^1(f(\Theta)) + (1 - \gamma) v^2(f(\Theta))$ for all $\gamma \in [0, 1]$.

Thus, for each i regime S^i can be constructed with payoffs $w_i^i = v_i(f)$ and $w_j^i < w_j^j$, $j \neq i$.

[5] This condition is originally from Dutta and Sen [8] and is weaker than the "bad outcome" condition in Moore and Repullo (1990).

We make several observations about condition ϕ. First, both parts of condition ϕ are generic: strict preferences would imply that all dictatorships yield unique payoffs and hence ensure part (i), and part (ii) is true as long as $v(f)$ is an extreme point of the convex hull of the set of feasible payoffs that can be generated from the range of f (recall that f is assumed to be efficient in the range). Second, when part (ii) fails to hold so that $v(f)$ can be obtained exactly by a convex combination of $d(1, f(\Theta))$ and $d(2, f(\Theta))$, we can not only implement the desired payoff profile at the outset but also, by Fudenberg and Maskin [9], one can actually alternate the two dictatorships in such a way that the continuation payoffs at any date are arbitrarily close to $v(f)$ if the players are sufficiently patient. In terms of payoff-repeated implementation, therefore, condition ϕ imposes almost no additional loss of generality. Third, there are other ways to construct S^i with desired payoff properties. For instance, if there exists some $\tilde{a} \in A$ such that $v_i(\tilde{a}) < v_i(f)$ for all i, S^i can be built by alternating $d(i, f(\Theta))$ and constant enforcement of $\tilde{a}$.

Given condition ϕ, we can further alternate the two dictatorships to construct a set of regimes $X(t)$ for each $t = 1, 2, \ldots$ and another regime Y that respectively induce unique payoff profiles $x(t)$ and y satisfying the following condition:

$$w_1^2 < y_1 < x_1(t) < w_1^1 \quad \text{and} \quad w_2^1 < x_2(t) < y_2 < w_2^2. \tag{1}$$

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

As mentioned earlier, we invoke an additional condition to guarantee existence of a desired equilibrium in the regime that we construct to obtain the sufficiency results. 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 strict 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 $\theta, \theta' \in \Theta$.

Next, we define the following extensive form 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 $\mathcal{Z} \equiv \{0, 1, 2\}$.

Figure 1: Regime construction

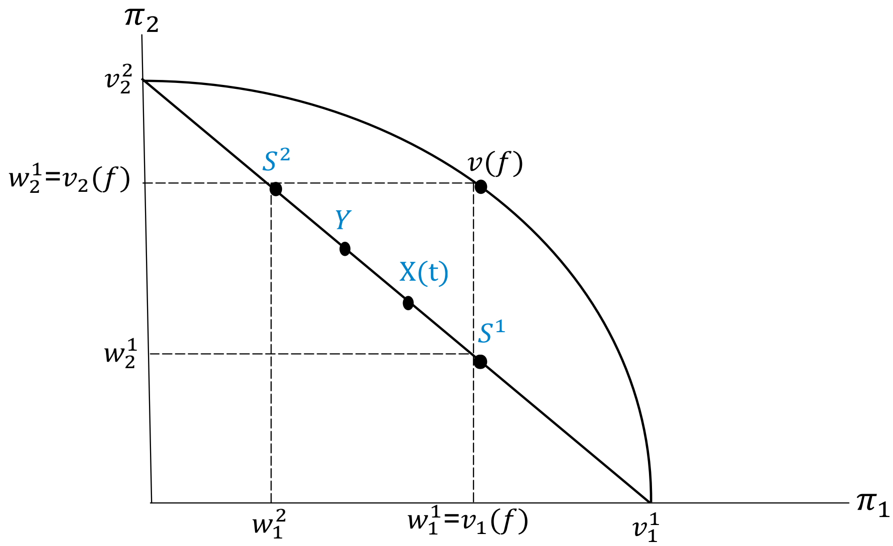

The outcome function of this mechanism depends solely on the agents' announcement of states in Stage 1 and is given below:

(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 $z_1 > 0$ and $z_2 = 0$ ($z_1 = 0$ and $z_2 > 0$), then the continuation regime is S^1 (S^2).

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, the implemented outcome depends solely on the announcement of states, while the integers dictate the continuation mechanism. In contrast to the corresponding constructions in LS, here we allow the agents to choose integers from only a finite set and also invoke a two-stage sequential structure. The latter change, as will be clarified shortly, enables us to define the notion of complexity of a strategy in a natural way.

Second, announcement of any non-zero integer effectively ends the strategic part of the game. Our regime is similar to that of LS in that 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 achieve our new objectives. 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 2 (Rule A.3(ii)) the continuation regimes could actually be different across periods. This feature will later be used to facilitate our refinement arguments.

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

We can now define strategies and payoffs for the above regime as follows. With slight abuse of notation, a mixed (behavioral) strategy of agent $i = 1, 2$ in regime R^e is the mapping $\sigma_i : \mathbf{H}^\infty \times D \rightarrow (\triangle\Theta) \cup (\triangle\mathcal{Z})$ such that, for any $\mathbf{h} \in \mathbf{H}^\infty$, $\sigma_i(\mathbf{h}, d) \in \triangle\Theta$ if $d \in D_\theta$ and $\sigma_i(\mathbf{h}, d) \in \triangle\mathcal{Z}$ if $d \in D_z$. 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 $\mathbf{h} \in \mathbf{H}^\infty$, i.e. when the history is such that the mechanism to be played is g^e.

3.2 Nash Equilibria

We begin the analysis of the above regime by establishing existence of a Nash 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 2 *Regime R^e admits a Nash equilibrium, σ^*, in Markov strategies such that, for any t, $\mathbf{h} \in \mathbf{H}^t$ and $\theta \in \Theta$ on the equilibrium path, (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 $\theta \in D_\theta$, and $\sigma_i^*(\mathbf{h}, (\theta, \underline{\theta})) = 0$ for all $\mathbf{h} \in \mathbf{H}^\infty$ and $\left(\theta, \underline{\theta}\right) \in D_z$.[6] 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.

To show that σ^* is a Nash equilibrium, 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 $\left(\theta, \underline{\theta}\right)$. 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 $v_i(f)$. ■

Our next Lemma is concerned with the players' equilibrium behavior whenever they face Stage 2 (the integer part) of mechanism g^e on the equilibrium path. 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.

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

Lemma 3 *Consider any Nash equilibrium of regime* R^e. *Fix any* t, $\mathbf{h} \in \mathbf{H}^t$ *and* $d = (\theta, \underset{\sim}{\theta}) \in D_z$ *on the equilibrium path. Then, one of the following must hold at* $(\mathbf{h}, d)$*:*

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

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

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.

Figure 2: Continuation payoffs

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

First, from Figure 2, the inequalities of (1) imply that any equilibrium with pure strategy at the relevant history must play 0. In such a case, 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 $w_i^i = v_i(f)$ 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.

Second, we show that if the players are mixing over integers then zero cannot be chosen. Since $x_i(t) > w_i^j$ and $y_i > w_i^j$ for $i, j = 1, 2$, the transition rules imply that each agent prefers to announce 1 than to announce 0 if the other player is announcing a positive integer for sure. It then follows that if agent i attaches a positive weight to 0 then the other agent j must also do the same, and i's continuation payoff is at least $v_i(f)$, 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 3, we can also show that if the players were to mix over integers at any on-the-equilibrium history 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 Nash equilibria of our regime can then be summarized as follows.

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

(a) Each player i announces 0 for sure at any $(\mathbf{h}, d) \in \mathbf{H}^\infty \times D_z$ 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 player i mixes between 1 and 2 at some $d \in D_z$ 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 2 imply that we obtain payoff-repeated implementation from period 2. 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 Refinement: Credibility and Complexity

Our characterization of Nash equilibria 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 arguments based on complexity considerations to select the former. Note first that, if we apply subgame perfection, the statements of Lemma 3 above can be readily extended to hold for any on- or *off*-the-equilibrium history after which the agents find themselves in the integer part of mechanism g^e; that is, in a subgame perfect equilibrium (SPE) of regime R^e, at *any* $(\mathbf{h}, d) \in \mathbf{H}^\infty \times D_z$ they must either choose 0 for sure or mix between 1 and 2. Also, the Nash equilibrium identified in Lemma 2 is itself an SPE.

In order to facilitate our complexity arguments, 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

$$x_1(t') \neq x_1(t'') \text{ and } x_2(t') \neq x_2(t'') \text{ for some } t', t''. \tag{2}$$

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

Clearly, this additional feature does not alter Lemmas 2 and 3, or their SPE extensions. However, it implies for any SPE that, if the agents mix over integers at some period t on or off the equilibrium path, each i's mixing probability, given by $\frac{x_i(t)-y_i}{x_i+x_i(t)-2y_i}$, is not constant across periods.[7]

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

Definition 3 *For any 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 mechanism g^e, σ_i' always behaves (randomizes) the same way while σ_i does not. Formally, there exists some $d' \in D \equiv D_\theta \cup D_z$ 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* $d \in D$, $d \neq d'$.

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

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

[7] Our results are unaffected by making $X(\cdot)$ dependent on the entire history and not just its date. See Section 6 for further discussion on this issue.

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 [6], Sabourian [27] and Gale and Sabourian [10] 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 SPEs as follows.

Definition 4 *A strategy profile σ is a weak perfect equilibrium with complexity cost (WPEC) of regime R^e if σ is an SPE and for each 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.[8] 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.

These complexity considerations imply that mixing over integers can no longer be part of equilibrium behavior in our regime.

Lemma 4 *Fix any WPEC of regime R^e. Also, fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d \in D_z$ (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 $d \in D_z$ depending on what happened in the past. But then,

[8] The two exceptions in the existing literature are Kalai and Neme [14] and Sabourian [27]. The notion of WPEC was first introduced by [27].

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 2 continues to hold with WPEC. Thus, combining the previous lemmas, we establish the following main result.

Theorem 1 *Suppose that $I = 2$ and $\delta \in \left(\frac{1}{2}, 1\right)$. 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 WPEC exists and (ii) every WPEC σ satisfies $\pi_i^{\mathbf{h}}(\sigma, R) = v_i(f)$ for any i, $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$.*

Proof. This follows immediately from Lemmas 2-4. ■

Notice that the extent of implementation achieved in Theorem 1 is stronger than that of Definition 2 above since, here, we obtain the desired payoffs at every on- and *off*-the-equilibrium history after period 1.

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 the following.

Definition 5 *An SCF f is strongly efficient if it is efficient and there does not exist a random SCF $\xi : \Theta \to \triangle(A)$ such that $v(\xi) = v(f)$; f is strongly efficient in the range if it is efficient in the range and there does not exist $\xi : \Theta \to \triangle(f(\Theta))$ such that $v(\xi) = v(f)$.*

Corollary 1 *Suppose that, in addition to the conditions in Theorem 1, f is strongly efficient in the range. Then, there exists a regime R such that (i) a WPEC exists and (ii) every WPEC σ satisfies $A^{\mathbf{h},\theta}(\sigma, R) = \{f(\theta)\}$ for any $t \geq 2$, $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$ and $\theta \in \Theta$.*

Proof. It suffices to show part (ii). 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 3 and 4, we know that, for any i,

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

which implies that $\sum_{\theta\in\Theta,a\in f(\Theta)} p(\theta)r(a,\theta)u_i\left(a,\theta\right) = v_i(f)$. Since f satisfies strong efficiency in the range, part (ii) of the claim follows. ■

3.4 Further Equilibrium 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 [3], Sabourian [27], Lee and Sabourian [16]) and to require the strategies to be minimally complex mutual best responses only *on* the equilibrium path.

Definition 6 *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 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 [6], Sabourian [27] and Gale and Sabourian [10]). 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 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, since WPEC is itself a PEC, all our previous WPEC results above remain valid under PEC. Additionally, we show that every PEC in regime R^e must be Markov.

Lemma 5 *Every PEC, σ, of regime R^e is Markov: for all 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. See Appendix A.1. ■

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

Theorem 2 *Suppose that $I = 2$ and $\delta \in \left(\frac{1}{2}, 1\right)$. 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 every PEC σ satisfies $\pi_i^{\mathbf{h}^t}(\sigma, R) = v_i(f)$ for any i, $t \geq 1$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$.*

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.

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 (by alternating two dictatorships) that generates a unique payoff profile $w^i = (w_i^i, w_j^i)$ such that $w_i^i = v_i(f)$ and $w_j^i \leq v_j(f)$. This was possible as long as the SCF f was efficient in the range, the two dictatorships $d(1, f(\Theta))$ and $d(2, f(\Theta))$ yielded unique payoffs, and $\delta > \frac{1}{2}$. With almost no loss of generality, we considered the case where the latter inequality was strict.

With three or more agents, we also need to be able to construct, for each agent i, regime S^i with the correct payoff property. We assume the following condition.

Condition χ. For each i, there exists $w^i = (w_1^i, \ldots, w_I^i) \in co(W)$ such that $v_i(f) = w_i^i \geq w_i^j$ for all $j \neq i$; moreover, for some $k, l \in I$, $w_k^j < w_k^k$ for all $j \neq k$ and $w_l^j < w_l^l$ for all $j \neq l$.

By Lemma 1 above, any payoff profile $w \in co(W)$ could be generated as a repeated game payoff of a regime that appropriately alternates some dictatorial mechanisms if $\delta \in \left(1 - \frac{1}{|W|}, 1\right)$. Assuming that δ indeed satisfies this condition (which we shall do throughout below), condition χ immediately implies that for each agent i there exists a regime S^i such that i obtains a payoff equal to the target level $v_i(f)$ but every other agent derives a payoff weakly less than his target, while for at least two agents the latter inequality is always strict. As before, let w^i denote the payoff profile associated with regime S^i.

We have already discussed that, with $I = 2$, condition χ imposes almost no restriction on our problem under efficiency. Let us present one case that guarantees condition χ with $I \geq 3$. The following lemma also demonstrates that we may not actually need a discount factor as large as $1 - \frac{1}{|W|}$ (as in Lemma 1) to obtain condition χ; here, the regime S^i could be constructed by alternating just two mechanisms and hence $\delta > \frac{1}{2}$ would be sufficient.

Lemma 6 *Consider an SCF f such that $(i, f(\Theta)) \in \Phi$ for all i and $v_i^j(f(\Theta)) \leq v_i(f)$ for all i, j, $i \neq j$. Suppose also that there exists some $\tilde{a} \in A$ such that $v_i(\tilde{a}) \leq v_i(f)$ for all i. Then, for any i, there exists $w^i \in co(W)$ such that*

(i) $w_i^i = v_i(f)$; and

(ii) $w_j^i \leq v_j(f)$ for any $j \neq i$, with this inequality being strict if either $v_j(f) > v_j(\tilde{a})$ or $v_j(f) > v_j^i(f(\Theta))$.

Proof. For each i, since $v_i^i(f(\Theta)) \geq v_i(f) \geq v_i(\tilde{a})$, there must exist $\alpha^i \in [0, 1]$ such that $v_i(f) = \alpha^i v_i^i(f(\Theta)) + (1 - \alpha^i)v_i(\tilde{a})$. Let $w^i = \alpha^i v^i(A) + (1 - \alpha^i)v(\tilde{a})$. Clearly, w^i satisfies (i) for all i.

To show (ii), consider any $j \neq i$. Then by construction $w_j^i \leq \max\{v_j^i(f(\Theta)), v_j(\tilde{a})\}$. Since by assumption $v_j(f) \geq v_j(\tilde{a})$ and $v_j(f) \geq v_j^i(f(\Theta))$, we have $w_j^i \leq v_j(f)$. Furthermore, the last inequality is strict if either $v_j(f) > v_j(\tilde{a})$ or $v_j(f) > v_j^i(f(\Theta))$. ■

Now, consider any SCF that is efficient in the range and satisfies condition χ. Henceforth, fix k and l as any two agents for whom the inequalities are strict in condition χ. Then, we can show the existence of the following (history-independent) regimes.

Lemma 7 *Suppose that f is efficient in the range and satisfies condition χ. 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 payoff profiles $w^C, x(t), y \in co(W(f))$ satisfying the following conditions:*[9]

$$w_k^l < y_k < x_k(t) < w_k^k \quad and \quad w_l^k < x_l(t) < y_l < w_l^l \tag{3}$$

$$x_k(t') \neq x_k(t'') \quad and \quad x_l(t') \neq x_l(t'') \quad for\ some\ t', t'' \tag{4}$$

$$w_k^C < w_k^k \quad if\ C \neq \{k\} \quad and \quad w_l^C < w_l^l \quad if\ C \neq \{l\} \tag{5}$$

$$w_i^C \geq w_i^{C\backslash\{i\}} \quad for\ all\ i \in C. \tag{6}$$

Proof. To construct these regimes, first let $y = \mu w^k + (1-\mu)w^l$ for some $\mu \in (0,1)$ and set $x(t) = \lambda(t)w^k + (1-\lambda(t))w^l$ such that $\{\lambda(t) : \lambda(t) \in (\mu, 1)\ \forall t\}$ and $\lambda(t') \neq \lambda(t'')$ for some t', t''. Also, for any $C \subseteq I$, let $w^C = \frac{1}{|C|}\sum_{i\in C} w^i$, where w^i is given by condition χ. Since $w_j^j > w_j^i$ for all $j \neq i$, these payoffs satisfy (3)-(6). Furthermore, since for each i, $w^i \in co(W)$ can be obtained as a convex combination of some dictatorships, it follows that $w^C, x(t), y \in co(W)$ can be all obtained by regimes that appropriately alternate the same dictatorships. ■

Using the constructions in Lemma 7, we extend the regime construction for the case of $I = 2$ to the case of $I \geq 3$. 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 $\{0, 1\}$.

Next, using the constructions in Lemma 7 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

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

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.

We define histories, partial histories (within period), strategies and continuation payoffs similarly to their two-agent counterparts. Also, the definitions of complexity and WPEC can be defined analogously here to the two-agent case above.

4.2 Results

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

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

(a) Each player $i \in I$ announces 0 for sure at any $(\mathbf{h}, d) \in \mathbf{H}^\infty \times D_z$ 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) Players k and l mix between 1 and 2 at some $d \in D_z$ 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 Nash 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 (3) 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 (5) 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 (6), 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 (4), 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 3 *Suppose that $I \geq 3$ and $\delta \in \left(1 - \frac{1}{|W|}, 1\right)$. If an SCF f is efficient in the range and satisfies condition χ, there exists a regime R such that (i) a WPEC exists and (ii) every WPEC σ satisfies $\pi_i^{\mathbf{h}}(\sigma, R) = v_i(f)$ for any i, $t \geq 2$ and $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$.*

Proof. See Appendix A.1. ■

As in Corollary 1 for the case of $I = 2$, we can strengthen Theorem 3 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

5.1 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 captures this by saying that a strategy that at every date t responds identically to some partial history d, independently of the previous history of play $\mathbf{h}$ 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 treatment of history in our complexity measure appeals to the presence of memory cost for recalling history of actions before the current period;

the partial history within the period can be interpreted as some stimuli that involve no cost of recalling.

This asymmetric treatment of history of outcomes before and within a period in our definition of complexity is motivated because, in repeated interactions, substantial time lags often exist between periods and it may require costly memory to condition the current action on the play of previous periods whereas, within a period, 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 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. One possibility is to differentiate the partial histories that occur in the two stages of the sequential mechanism. More specifically, we can assume that when the players are asked to announce a state at the beginning of the first stage of each period, nature's move θ is simply some stimulus that is known at no cost, while in the second stage of each period, when making integer announcements, recalling partial history that previously occurred is costly due to a significant time lag between stages. This would mean that we need to modify our previous definition of complexity such that conditioning behavior on the partial history of play in the first stage of the mechanism is more complex than not doing so. A complexity measure that reflects this idea is follows.

Definition 7 *For any 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 $d' \in D_\theta \equiv \Theta$ such that

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

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

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

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

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

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

Definition 7 may be a plausible complexity criterion not only because the knowledge of nature's move θ at the beginning of each period may be costless information, but also because it influences the players' immediate payoffs, whereas the partial history at the beginning of the second stage of the mechanism only affects future payoffs. 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.

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 4 that if a WPEC were to involve mixing over integers at any history, deviating to always announcing integer 1 at every information set in the second stage would always generates 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 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 $l \in \{\theta, z\}$ 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. Even if there is no significant time lag between decision turns, however, we have already argued that this approach may be less plausible than Definition 3 or 7 because nature's moves may be just simple costless stimuli and also are immediately payoff-relevant.

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 and 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 in environments where a strict one-period punishment is possible for deviations from truth-telling. With $I = 2$, a stronger self-selection condition with strict inequalities would achieve this. 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 [24]). Then, we could alter $\widehat{R}^e$ above 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. Note that the transition rules are entirely independent of actions of Stage 1 and, hence, all our characterization results remain unaffected; on the other hand, the Markov strategy profile in which the true state and integer 0 are always announced is a strict WPEC.

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 the bad outcome or its equivalent. In Appendix A.2, we construct such a regime for the case of $I = 2$ (without self-selection) which delivers the same results with Definition 8.

5.2 One-Shot Mechanisms

Our analysis is built on regime constructions employing sequential mechanisms. As should be clear from our WPEC arguments above, the use of sequential mechanisms allows us to treat complexity of behavior regarding states and integers differently and invoke a minimal partial order over the set of strategies.

Definition 9 *For any 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 mechanism g^e and some set of histories, σ_i' always behaves (randomizes) the same way while σ_i does not. Formally, there exists some $d' \in D \equiv D_\theta \cup D_z$ and $\mathbf{H}' \subseteq \mathbf{H}^\infty$ with the following properties:*

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

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

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

5.3 Complete Complexity Orders

The above notions of complexity are based on partial ordering over strategies. We could also adopt complete complexity orders such as counting the number of "states of the minimal automaton" implementing the strategy (e.g. Abreu and Rubinstein [3]) or the number of "continuation strategies" induced by the strategy (e.g. Kalai and Stanford [15], Lee and Sabourian [16]). Another measure that can be used is the "collapsing state condition" (Binmore, Piccione and Samuelson [5]).

To see how our WPEC results can be obtained with these complete orders, consider the notion of state complexity in which the number of states of an automaton counts the number of continuation strategies that the underlying strategy induces. By appropriately defining the output function of the automaton, we can then make the corresponding state complexity measure to be consistent with any of our partial orders above. For instance, an automata definition that includes all the partial histories (both at the first and second stages of the sequential mechanism) as arguments of the output function would correspond to Definition 3, while an output function that does not involve any type of partial history would lead to state complexity in line with Definition 8.

6 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 Lee and Sabourian [17], 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 $\{X(t)\}_{t=0}^{\infty}$, 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 (2) 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 $\{X(t)\}_{t=0}^{\infty}$ 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.

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 [25]).[10] 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.

A Appendix

A.1 Omitted Proofs

Proof of Lemma 3

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.

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

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 postive 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 3, it suffices to show that any mixing over integers in equilibrium must occur in period 1. Suppose not; so, there exists a Nash equilibrium σ such that, for some $t > 1$, there exist $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_z$ that occur on the equilibrium path at which the players are mixing over integers.

First, note that by Lemma 3 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 $d' \in D_z$, 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 $(d', \underset{\sim}{z}) \in D_z \times \mathcal{Z}^2$ 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

$$\pi_i^{\mathbf{h}^t}(\sigma, R^e) = \sum_{(d',\underset{\sim}{z}) \in D_z \times \mathcal{Z}^2} r(d', \underset{\sim}{z}) \left[(1-\delta) u_i(a^{\mathbf{h}^t,d'}, d') + \delta \pi_i^{\mathbf{h}^t,d',\underset{\sim}{z}} \right] = v_i(f). \tag{7}$$

Lemma 3 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 (7) 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 4

Suppose not. Then, by the SPE extension of Lemma 3 discussed at the beginning of Section 3.3, there exists a WPEC, σ, such that, at some t, $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_z$, 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 3, there are two cases to consider at $(\mathbf{h}, d)$.

First, suppose that j plays 0 for sure. Then, by part (a) of Lemma 3, 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 3, j mixes between 1 and 2 and i is also indifferent between choosing 1 and 2.

Proofs of Proposition 2 and Theorem 3

These results are proved by the following lemmas.

Lemma 8 *Regime $\widehat{R}^e$ admits a Nash equilibrium, 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 2. ■

Lemma 9 *Consider any Nash equilibrium of regime $\widehat{R}^e$. Fix any t, $\mathbf{h} \in \mathbf{H}^t$ and $d \in D_z$ on the equilibrium path. 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 $(\mathbf{h}, d) \in \mathbf{H}^t \times D_z$. 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.

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. We derive contradiction by considering the following subcases.

Subcase 1A: $z_k > 0$ 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 (3), 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.

Subcase 1B: $z_k > 0$ 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 (3), 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 (3), 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 (5), 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 $\underset{\sim}{z}_{-k} \neq (0, \ldots, 0)$. 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 (3), we have $x_k > y_k > w_k^l$.

Second, suppose that $z_l = 0$. In this case, since $\underset{\sim}{z}_{-k} \neq (0, \ldots, 0)$, 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 (5), 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 (6), $w_i^C \leq w_i^{C \cup \{i\}}$. 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 postive 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, (3) implies that the other must also do the same. But, by applying (3) 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 (3), 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. ■

If we consider SPEs instead of Nash equilibria, Lemma 9 can be extended to hold for any on- or off-the-equilibrium history/partial history. Also, the equilibrium profile identified in Lemma 8 is itself an SPE and, also, a WPEC. This, together with our next Lemma, completes the proof of Theorem 3.

Lemma 10 *Fix any WPEC of regime $\widehat{R}^e$. Also, fix any t, $\mathbf{h} \in \mathbf{H}^\infty$ and $d \in D_z$. Then, every agent announces zero for sure at this history.*

Proof. Suppose not. Then, by the SPE extension of Lemma 9 as discussed above, there exists a WPEC, σ, such that, at some t, $\mathbf{h}^t \in \mathbf{H}^t$ and $d \in D_z$, 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 9, 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 9, 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 9, agents k and l mix between 1 and 2; thus, i is indifferent. ■

Proof of Lemma 5

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 4, we know that $d' \in D_\theta$; 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.

A.2 Alternative Regime Construction

Here, we offer an alternative construction for the two agent case that does not involve the self-selection condition but instead requires sufficiently large δ. The construction and its equilibrium properties below can be similarly extended to the case of $I \geq 3$.

The construction involves the following two mechanisms. Let $g'(1)$ denote an extensive-form mechanism such that:

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 is such that a constant outcome, $f\left(\tilde{\theta}\right)$ for some arbitrary but fixed $\tilde{\theta} \in \Theta$, is always implemented.

Let g' be an extensive-form mechanism such that:

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 $\mathcal{Z}$.

The outcome function is given below:

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

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

Note that this mechanism differs from g^e in Section 3 in that it does not invoke the self-selection condition when the agents announce different states.

Next, we define regime R' inductively with mechanism $g'(1)$ enforced in period 1 and the transition rules below.

Period 1:

Let (z_1, z_2) be the integers announced. The transition rules in period 1 are as follows.

Rule C.1: If $z_1 = z_2 = 0$, the mechanism next period is g'.

Rule C.2: If $z_1 > 0$ and $z_2 = 0$ ($z_1 = 0$ and $z_2 > 0$), the continuation regime is S^1 (S^2).

Rule C.3: Suppose that $z_1, z_2 > 0$. Then, we have the following:

Rule C.3(i): If $z_1 = z_2 = 1$, the continuation regime is $X \equiv X(\tilde{t})$ for some arbitrary $\tilde{t}$, with the payoffs henceforth denoted by x.

Rule C.3(ii): If $z_1 = z_2 = 2$, the continuation regime is $X(1)$.

Rule C.3(iii): If $z_1 \neq z_2$, the continuation regime is Y.

Period $t \geq 2$:

Consider any date $t \geq 2$. Let (θ_1, θ_2) and (z_1, z_2) be the states and integers announced in period t. The transitions rules are as follows.

Rule D.1: If $\theta_1 \neq \theta_2$, the continuation regime is X.

Rule D.2: If $\theta_1 = \theta - 2$ and $z_1 = z_2 = 0$, the mechanism next period is g'.

Rule D.3: If $\theta_1 = \theta_2$, $z_1 > 0$ and $z_2 = 0$ ($z_1 = 0$ and $z_2 > 0$), the continuation regime is S^1 (S^2).

Rule D.4: Suppose that $\theta_1 = \theta_2$ and $z_1, z_2 > 0$. Then, we have the following:

Rule D.4(i): If $z_1 = z_2 = 1$, the continuation regime is X.

Rule D.4(ii): If $z_1 = z_2 = 2$, the continuation regime is $X(t)$.

Rule D.4(iii): If $z_1 \neq z_2$, the continuation regime is Y.

This regime modifies R^e in Section 3 in the following way. In the first period, the planner enforces a constant outcome but the integer play generates essentially identical transition rules as in R^e. In any period after the first, the agents play a sequential revelation mechanism with integers g', but the transition rules when playing g' is identical to the corresponding features of R^e only if the two agents announce the same state in Stage 1; otherwise, the continuation regime is X, which generate continuation payoffs strictly dominated by $v(f)$. With some abuse of notation, let us define the set of histories, partial histories, strategies and payoffs as before.

To examine the set of equilibria of regime R', note first that the statements of Lemma 3, or their SPE extensions, can be extended here as follows: conditional on any history beyond the first period at which mechanism g' is played *and the same state announced in Stage 1*, the two players must either report 0 for sure and obtain $v(f)$, or uniquely mix between 1 and 2 and obtain strictly less than $v(f)$; for $t = 1$, this also holds by similar arguments. Second, as in Lemma 4, in any WPEC, the agents must play 0 for sure on- or off-the-equilibrium. (Note that, following any partial history involving disagreement in Stage 1, the integer play does not affect the continuation game because of Rule D.1.)

In the next lemma, we establish that the players must always report the same state *after the first period*. The basic intuition is that, otherwise, the continuation payoff of each agent i falls short of $v_i(f)$ and hence a deviation would occur in the previous period's integer stage. Note that this argument could not work if disagreement occurred in the first period; in order to avoid such coordination failure, we implement a constant outcome in period 1.

Lemma 11 *Fix any WPEC of regime R'. Also, fix any $t \geq 2$ and $\mathbf{h}^t \in \mathbf{H}^t$. Then, the agents always report the same state for sure.*

Proof. Let $r(\theta, \underset{\sim}{\theta})$ denote the probability with which partial history $(\theta, \underset{\sim}{\theta}) \in D_z$ occurs at $\mathbf{h}^t$ under the given WPEC, and let $a^{\mathbf{h}^t, \theta, \underset{\sim}{\theta}}$ represent the corresponding outcome. Also, let $\underset{\sim}{\Theta}' = \{ \ (\theta^1, \theta^2) \in \Theta^2 \mid \theta^1 = \theta^2 \}$ denote the set of state profiles in which the players agree and $\underset{\sim}{\Theta}'' = \Theta^2 \backslash \underset{\sim}{\Theta}'$ denote the set of state profiles in which they disagree.

By definition, at $\mathbf{h}^t$ mechanism g' is played. Therefore, in the previous period $t-1$ of history $\mathbf{h}^t$, one of the following must be true: (i) the same mechanism was in force, and moreover, by the transition rules of R', the players announced the same state in Stage 1 and integer 0 in Stage 2, or (ii) the players were in period 1 and announced 0 for sure. Then, by previous arguments, it must be that $\pi_i^{\mathbf{h}^t} = v_i(f)$ for all i.

Next, at $\mathbf{h}^t$, if the agents report the same state in Stage 1, by applying the arguments of Lemma 4, the agents report zero for sure in Stage 2. It then follows that the continuation payoff profile after any $d = (\theta, \underset{\sim}{\theta})$ is $v(f)$ if $\underset{\sim}{\theta} \in \underset{\sim}{\Theta}'$ and, by Rule D.1, x if $\underset{\sim}{\theta} \in \underset{\sim}{\Theta}''$.

Therefore, we can write each i's continuation payoff at $\mathbf{h}^t$ as

$$\begin{aligned}\pi_i^{\mathbf{h}^t} &= \sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}'} r(\theta,\underset{\sim}{\theta})\left[(1-\delta)u_i(a^{\mathbf{h}^t,\theta,\underset{\sim}{\theta}},\theta)+\delta v_i(f)\right] + \sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}''} r(\theta,\underset{\sim}{\theta})\left[(1-\delta)u_i(a^{\mathbf{h}^t,\theta,\underset{\sim}{\theta}},\theta)+\delta x_i\right] \\ &= (1-\delta)\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\Theta^2} r(\theta,\underset{\sim}{\theta})u_i(a^{\mathbf{h}^t,\theta,\underset{\sim}{\theta}},\theta)+\delta\left[v_i(f)\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}'} r(\theta,\underset{\sim}{\theta}) + x_i\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}''} r(\theta,\underset{\sim}{\theta})\right].\end{aligned}$$

Since $\pi_i^{\mathbf{h}^t} = v_i(f)$ and $x_i < v_i(f)$ for all i, if $\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}''} r(\theta,\underset{\sim}{\theta}) \neq 0$ then it must be that $\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\Theta^2} r(\theta,\underset{\sim}{\theta})\ u_i(a^{\mathbf{h}^t,\theta,\underset{\sim}{\theta}},\theta) > v_i(f)$ for all i. But this is not feasible with f being efficient in the range. It therefore follows that $\sum_{\theta\in\Theta,\underset{\sim}{\theta}\in\underset{\sim}{\Theta}''} r(\theta,\underset{\sim}{\theta}) = 0$. ■

Next, define

$$\bar{\delta} = \max_{i\in I}\left\{\frac{\max_{a,a'\in f(\Theta),\theta\in\Theta}\{u_i(a,\theta)-u_i(a',\theta)\}}{\max_{a,a'\in f(\Theta),\theta\in\Theta}\{u_i(a,\theta)-u_i(a',\theta)\}+(v_i(f)-x_i)}\right\}\in(0,1).$$

We obtain existence below.

Lemma 12 *If $\delta > \bar{\delta}$, regime R' admits a WPEC.*

Proof. Consider the following repeated game strategy profile. In period 1, each player announces 0 for sure. From period 2, each player always reports the true state followed by integer 0.

To see that this profile constitutes an SPE, by Rules C.2 and D.3, neither player wants to deviate at any integer-reporting stage; by Rule D.1, and since $\delta > \bar{\delta}$, deviation from the prescribed state-reporting strategy is not profitable. It is also clear that this SPE is a WPEC itself. ■

A.3 Construction with One-Shot Mechanisms

Suppose that $I = 2$. Define g^* as a one-shot mechanism in which, for each $i = 1, 2$, $M_i = \Theta \times \mathcal{Z}$ and the outcome function is such that:

(i) If m_1 and m_2 are such that $\theta_1 = \theta_2 = \theta$ and $z_i = 0$ for some i, $f(\theta)$ is implemented.

(ii) If m_1 and m_2 are such that $\theta_1 \neq \theta_2$ and $z_i = 0$ for some i, an outcome from the set $L_1(\theta_2) \cap L_2(\theta_1)$, as defined by self-selection in the range, is implemented.

(iii) If m_1 and m_2 are such that $z_1 > 0$ and $z_2 > 0$, a constant outcome, $f\left(\tilde{\theta}\right)$ for some arbitrary but fixed $\tilde{\theta} \in \Theta$, is always implemented.

Next, regime R^* is defined inductively as follows. First, mechanism g^* is played in $t = 1$. Second, if, at date $t \geq 1$, g^* is the mechanism played with $m_i = (\theta_i, z_i)$ being the message announced by $i = 1, 2$, the continuation mechanism or regime at the next period is given by the transition rules below:

Rule E.1: If $z_1 = z_2 = 0$, then the mechanism next period is g^*.

Rule E.2: If $z_1 > 0$ and $z_2 = 0$ ($z_1 = 0$ and $z_2 > 0$), then the continuation regime is S^1 (S^2).

Rule E.3: If $z_1, z_2 > 0$, then we have the following:

Rule E.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 E.3(ii): If $z_1 = z_2 = 2$, the continuation regime is $X(t)$.

Rule E.3(iii): If $z_1 \neq z_2$, the continuation regime is Y.

For this regime, with slight abuse of notation, let $\mathbf{H}^t$ continue to denote the set of histories at the beginning of period t. Then, we write player i's strategy as $\sigma_i : \mathbf{H}^\infty \times \Theta \rightarrow \triangle(\Theta \times \mathcal{Z})$; for $\mathbf{h} \in \mathbf{H}^\infty$, $\pi_i^{\mathbf{h}}(R^*, \sigma)$ denotes i's continuation payoff in this regime under strategy profile σ.

Define $\overline{u} = \max_{i \in I, a \in A, \theta \in \Theta} u_i(a, \theta)$ and $\underline{u} = \min_{i \in I, a \in A, \theta \in \Theta} u_i(a, \theta)$. Also, for each i and $j \neq i$, define $E_i = \{w_i^i, w_i^j, x_i(t), y_i\}_{t \in \mathbb{Z}_{++}}$ and $\underline{e}_i = \min_{e, e' \in E_i} |e - e'|$; let $\underline{e} = \min_{i \in I} \underline{e}_i$. Then, define

$$\bar{\delta} = \frac{\overline{u} - \underline{u}}{\overline{u} - \underline{u} + \underline{e}}.$$

Let us consider Nash equilibria of the above regime.

Lemma 13 *Suppose that $\delta \in (\bar{\delta}, 1)$. Consider any Nash equilibrium of regime R^*, and fix any t, $\theta \in \Theta$ and $\mathbf{h} \in \mathbf{H}^t$ on the equilibrium path. For each $i = 1, 2$, let $r_i(\theta, z)$ denote the probability with which i's equilibrium strategy plays message (θ, z) at $(\mathbf{h}, \theta)$. Then, one of the following must hold:*

(a) For each i, $\sum_{\theta \in \Theta} r_i(\theta, 0) = 1$ and his continuation payoff at the next period is $v_i(f)$.

(b) For each i, $\sum_{\theta \in \Theta} r_i(\theta, 1) + \sum_{\theta \in \Theta} r_i(\theta, 2) = 1$ such that $\sum_{\theta \in \Theta} r_i(\theta, 1) > 0$ and $\sum_{\theta \in \Theta} r_i(\theta, 2) > 0$; the mixing probabilities depend on $x(t)$; i's continuation payoff at the next period is less than $v_i(f)$.

Proof.

Fix any $\delta > \bar{\delta}$ and any Nash equilibrium of R^*. Fix any t, $\theta \in \Theta$ and $\mathbf{h} \in \mathbf{H}^t$ on the equilibrium path. Let z_i denote the integer that i ends up choosing at $(\mathbf{h}, \theta)$ and Π_i denote i's continuation payoff at the next period if both agents announce zero at the given history.

Also, for any i, whenever we mention a deviating strategy that is identical to the equilibrium strategy everywhere except that at $(\mathbf{h}, \theta)$ it announces integer z', we mean the following: letting $r_i'(\theta, z)$ be the probability of playing (θ, z) under the deviating strategy at $(\mathbf{h}, \theta)$, we have $\sum_{\theta \in \Theta} r_i'(\theta, z') = 1$ and $r_i'(\theta, z') = \sum_{z \in \mathcal{Z}} r_i(\theta, z)$ for any θ.

We consider two cases in turn.

<u>*Case 1*</u>: No player randomizes over integers, i.e. $\sum_{\theta \in \Theta} r_i(\theta, 0) = 1$, $\sum_{\theta \in \Theta} r_i(\theta, 1) = 1$ or $\sum_{\theta \in \Theta} r_i(\theta, 2) = 1$ for all i.

In this case we show that each i must play 0 for sure, i.e. $\sum_{\theta \in \Theta} r_i(\theta, 0) = 1$. 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.

<u>Subcase 1A</u>: $z_i > 0$ and $z_j = 0$.

The continuation regime at the next period is S^i (Rule E.2). Thus, j's equilibrium continuation payoff at $(\mathbf{h}, \theta)$ is at most $(1 - \delta)\overline{u} + \delta w_j^i$. Consider j deviating to another strategy identical to the equilibrium strategy except that it announces the positive integer other than z_i at this history. By (iii) of the outcome function of g^*, and by Rule E.3(iii), the corresponding continuation payoff is $(1 - \delta) f(\tilde{\theta}) + \delta y_j$. Since $y_j > w_j^i$ by construction, and since $\delta > \bar{\delta}$, the deviation is profitable. This is a contradiction.

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

The continuation regime is either X, $X(t)$ or Y (Rule E.3). Also, the current period's outcome is $f(\tilde{\theta})$. 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^* 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, consider i deviating to another strategy identical to the equilibrium strategy except that it announces a positive integer at this history. By (i) and (ii) of g^*, and by Rule E.2, such a deviation does not alter the current period's payoff but leads to a continuation payoff at the next period equal to $v_i(f)$, a contradiction. The rest follows as in the proof of Lemma 3 since f is efficient in the range.

Case 2: Some player randomizes over integers.

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 at the next period is w_i^j. If he announces 1, by Rules A.3(i) and A.3(iii), the continuation payoff at the next period is $x_i > w_i^j$ or $y_i > w_i^j$. Since $\delta > \bar{\delta}$, the current period's payoff does not matter.

Claim 2: Suppose that agent i announces 0 with positive probability, i.e. $\sum_{\theta \in \Theta} r_i(\theta, 0) > 0$. 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. Let us show the first part of this claim by way of contradiction. So, suppose that $\sum_{\theta \in \Theta} r_j(\theta, 0) = 0$. Consider i deviating to a strategy identical to the equilibrium strategy except that it announces integer 1 for sure at this history. By Claim 1, the deviation is profitable, a contradiction against the assumption that $\sum_{\theta \in \Theta} r_i(\theta, 0) > 0$. The latter parts of the claim follow immediately.

We now show that, in this Case 2, both players choose a postive integer for sure. To show this suppose otherwise; then some player chooses 0 with 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^* involves only outcomes in the range of f. Therefore, since f is efficient in the range, we have a contradiction.

In this case, therefore, for each i, $\sum_{\theta \in \Theta} r_i(\theta, 1) + \sum_{\theta \in \Theta} r_i(\theta, 2) = 1$ and, clearly, the mixing probabilities must depend on $x(t)$. 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 $v_i(f)$. $\blacksquare$

We adopt the following complexity definition.

Definition 10 *For any i and any pair of strategies $\sigma_i, \sigma_i' \in \sigma_i$ of regime R^*, we say that σ_i is more complex than σ_i' if there exist some $\mathbf{H}' \subseteq \mathbf{H}^\infty$ and $\theta' \in \Theta$ with the following properties:*

(i) $\sigma_i'(\mathbf{h}, \theta) = \sigma_i(\mathbf{h}, \theta)$ *for all* $\mathbf{h} \in \mathbf{H}^\infty \backslash \mathbf{H}'$ *and all* θ.

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

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

It is straightforward to identify that R^* has a WPEC in which each agent always announces the true state and integer 0 (for any δ). The next lemma characterizes WPECs of R^*.

Lemma 14 *Suppose that $\delta \in (\bar{\delta}, 1)$. Consider any WPEC of regime R^*, and fix any t, $\theta \in \Theta$ and $\mathbf{h} \in \mathbf{H}^t$ (on or off the equilibrium path). Then, each agent announces zero for sure at this history.*

Proof. Suppose not. Then, since the results of Lemma 13 must hold with SPE for any on- or off-the-equilibrium history, there exists a WPEC, σ, such that, at some t, θ and $\mathbf{h}^t \in \mathbf{H}^t$, the two agents play integer 1 or 2 for sure with the mixing probabilities determined by $x(t)$. 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 $\sigma_i(\mathbf{h}^t, \theta) \neq \sigma_i(\mathbf{h}^{t'}, \theta)$ for some $\mathbf{h}^{t'} \in \mathbf{H}^{t'}$, or $\sigma_i(\mathbf{h}^t, \theta) \neq \sigma_i(\mathbf{h}^{t''}, \theta)$ for some $\mathbf{h}^{t''} \in \mathbf{H}^{t''}$. Without loss of generality, assume the former.

Now, consider any $i = 1, 2$ deviating to another strategy σ_i' that is identical to the equilibrium strategy σ_i except that $\sigma_i'(\mathbf{h}, \theta) = \sigma_i'(\mathbf{h}', \theta)$ for all $\mathbf{h}, \mathbf{h}' \in \mathbf{H}^t \cup \mathbf{H}^{t'}$ such that integer 1 is announced with probability 1. Since σ_i' is less complex than σ_i according to Definition 10, we obtain a contradiction by showing that $\pi_i^{\mathbf{h},\theta}(\sigma_i', \sigma_{-i}, R^*) = \pi_i^{\mathbf{h},\theta}(\sigma, R^*)$ for any $\mathbf{h} \in \mathbf{H}^t \cup \mathbf{H}^{t'}$.

Fix any $\mathbf{h} \in \mathbf{H}^t \cup \mathbf{H}^{t'}$ and consider history $(\mathbf{h}, \theta)$. First, suppose that j plays 0 for sure. Then, by part (a) of the SPE extension of Lemma 13, i also plays 0 for sure and obtains a continuation payoff equal to $v_i(f)$ at the next period in equilibrium. By (i) and (ii) of the outcome function of g^*, the deviation does not alter the current outcomes and, by Rule E.2, it also induces the same continuation payoff $v_i(f)$ at the next period. Second, suppose that j is mixing over integers at this history. Then, by part (b) of the SPE extension of Lemma 13, j mixes between 1 and 2 in equilibrium. By (iii) of g^*, the deviation does not alter the current outcomes, while i is indifferent between choosing 1 and 2 in terms of the next period's continuation payoff. ■

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 H. Matsushima (1994): "Exact Implementation," *Journal of Economic Theory*, 64, 1-19.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[20] 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.

[21] Mezzetti, C. and L. Renou, "Repeated Nash Implementation," *Mimeo.*

[22] 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.

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

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

[25] 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.

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

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

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

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

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

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