ORIGINAL PAPER

Repeated implementation and complexity considerations

Hannu Vartiainen

Received: 11 November 2005 / Accepted: 22 October 2007 / Published online: 22 November 2007 © Springer-Verlag 2007

Abstract A game form is commitment-free if single actions of players do not have physical consequences, i.e., affect the continuation game. Such game can be thought to represent inifinite interaction with complete patience. A choice rule can be Nash implemented via a commitment-free mechanism if and only if it coincides with the feasible set of a normal form game. However, when players are complexity averse (in the lexicographic sense), then any Nash implementable choice rule becomes available.

Keywords Repeated interaction · Commitment-freeness · Nash implementation · Complexity considerations

JEL Classification C70 · D71 · D78

1 Introduction

In many real life problems, an arbitrator can only implement a choice rule via conversation, a game in which messages are exchanged unboundedly long time, without players clearly seeing the deadline or the final message stage. What matters is the agreement that players reach (or do not reach) in the long run. One could imagine this being the case, e.g., in political interaction or in marriage counseling.

H. Vartiainen (🖂)

H. Vartiainen Turku School of Economics, Turku, Finland

I am very grateful to Matt Jackson and a referee for their comments. I also thank Klaus Kultti and Hannu Salonen for useful conversations.

Yrjö Jahnsson Foundation, Ludviginkatu 3-5, 00130 Helsinki, Finland e-mail: hannu.vartiainen@yjs.fi

Aumann and Hart (2003) model conversation as a *long cheap talk* game: a player may have payoff relevant information but an outcome is chosen only after an infinite sequence of messages has been exchanged. The difference between finitely and infinitely repeated cheap talk game is analogous to the difference between finitely and infinitely repeated prisoners' dilemma: In the former case, the existence of a final round unravels everything backwards and ruins hope for cooperation. In the infinite case, however, unraveling need not take place. Thus *not being able to commit* to playing the stage game one more time, which is the case under long cheap talk, allows wider scope of cooperation.

In the scenario of Aumann and Hart (2003), the two players send messages infinitely long before they choose their physical actions in a bimatrix game. We apply their idea to the implementation context and assume that it is the planner who chooses, according to a prefixed rule, the outcome after an infinite exchange of messages. The rule according to which the outcome is implemented. This rule, i.e. the *mechanism*, is restricted by the condition that we dub *commitment-freeness*: no single message of any player affects how the implemented outcome depends on the messages. Commitment-freeness reflects the idea that the planner cannot commit to not let the players play again the message game. Due to this property alone, all terminal histories of a commitment-free game are infinitely long, and the implemented outcome is not sensitive to finite deviations. However, the implemented outcome may be sensitive to an *infinite* deviation. Hence, commitment-freeness is a measurability restriction on the outcome function.

Formally, implementation via a commitment-free mechanism takes place in three phases: (i) the initial phase t = 0 when the mechanism is designed and the players' preferences are determined, (ii) the communication phase t = 1, 2, ... in which players send infinite sequences of public messages, and (iii) the implementation phase $t = \omega + 1$ when an outcome is chosen based on the realized stream of messages. The planner's game design problem is to associate outcomes to all infinite streams of messages in a way that desirable Nash equilibria, and only them, are induced. Choice rules that are implementable via a commitment-free mechanism are said to be commitment-freely implementable.

The standard way of studying commitment in game theory is to do it through repeated games. Indeed, one may think a commitment-free game a repeated game where an outcome is implemented in each period with small probability (rather than future being discounted). Assuming that the limit is well defined, the players' payoffs can be characterized by the *limit of the means* criterion (cf. Aumann and Shapley 1976; see also Osborne and Rubinstein 1995; Aumann 1997). This criterion puts no weight on payoffs in a single period.^{1,2}

The fundamental problem with commitment-freeness in terms of implementation is that there is no way of preventing players' coordination. Indeed, a version of the folk theorem applies: A set of outcomes is commitment-freely implementable if this

¹ Limit of the means gives same preference ordering over the streams of payoffs as does the standard discounting criterion when the discount factor approaches unity. Moreover, the discounted infinite stream payoff converges as the factor approaches unity exactly when limit of the means exists (see Binmore 1998).

 $^{^2}$ Limit of the means is not defined for all payoff streams. However, in the domain of strategies that are implementable via finite state machines it is well defined. See the discussion in the final section of the paper.

set agrees with a feasible set (the set of minmax dominating outcomes) of a normal form mechanism. The converse also holds; the feasible set of any normal form mechanism can be commitment-freely implemented. For example, no Pareto optimal rule is implementable via a commitment-free game in the unrestricted preference domain (with the outcome set containing at least two elements).

However, the folk theorem requires players to use complex strategies that may not be easily induced by players with bounded computational capacity. It is well known that even weak *complexity aversion* can reduce dramatically the number of equilibria in repeated games. For example, the folk theorem collapses.³ The main results of this paper concern implementable rules when the players are computationally bounded, i.e., have tendency to avoid complex strategies. Complexity of a strategy is modeled in terms of the size of the smallest (Moore) machine that induces the strategy. A complexity averse player has a slight preference for strategy that is inducable via a smaller machine.⁴

The planner can exploit players' complexity aversion when designing the mechanism. We show that any *Nash implementable* choice rule can be Nash implemented via a commitment-free game form. This is done by constructing a canonical commitmentfree mechanism that Nash implements any Nash implementable outcome when players are complexity averse. Hence, under complexity aversion commitment inability is not a restriction on implementation. One way to interpret this is that complexity aversion is an alternative assumption that game theory can be based on, rather than commitment ability.

Finally, we show that the canonical commitment-free mechanism can be interpreted in the language of repeated games. We construct a game form that implements an outcome in each period t = 1, 2, ... as a function of the message history until t, and argue that the limit of the means payoffs generated from this game form are equivalent to the canonical mechanism.

Literature Nash equilibria of commitment-free games are related to the concept of *program equilibria* by Tennenholz (2004). Program equilibrium is based on the idea that computers are not only programmed to implement strategies but also to read how other computers are programmed. Hence an equilibrium strategy can be made directly contingent on the other players' program. Tennenholz shows that all minmax dominating strategies can be implemented. One difference is that we are not, a priori, restricted to computable strategies. The second is that our players do not scan other players' programs, only actions. This guarantees that a play path can always be determined.⁵ It would be interesting to explore whether the results of this paper hold, or appropriate versions of them, under implementation in program equilibrium.

Commitment is not a new question the implementation literature. A common theme in the literature is that a mechanism is an intermediate institution, a communication device, whose output is merely a suggestion that can be changed. The approach closest to this one is Jackson and Palfrey (1998, 2001) who study *voluntary* (stationary)

³ Rubinstein (1986) is a pioneering work. See also Abreu and Rubinstein (1988), Banks and Sundaram (1990), and Chatterjee and Sabourian (2000).

⁴ In the lexicographic sense.

⁵ In Tennenholz (2004) the existence is guaranteed by a natural restriction on the programming language.

implementation.⁶ In their framework, any player may veto an outcome of the mechanism and opt instead to play the same mechanism over again, without a bound.⁷ As a consequence, only mutually beneficial deviations are effective. Assuming stationary strategies, this imposes a restriction on mechanisms and leads to a condition that is more stringent than Nash implementation. However, the Pareto correspondence can be implemented via such voluntary process.

Kalai and Ledyard (1998) study dominant strategy implementation in a repeated framework. Their focus is on learning. They show that learning along the play path may allow the planner to achieve socially optimal outcome in the long run, provided that players discount future payoffs.

There are notably differences to this paper. First, we allow general strategies. Second, what drives the Jackson–Palfrey-model is that players *can* commit to an agreement. Since a move to the next round must be mutually accepted, there is an a priori tendency towards players' coordination. Finally, we derive the structure of the game from the no commitment-condition.

Section 2 defines the implementation concepts. Section 3 gives the results. The final section concludes with discussion.

2 The concepts

The set up Let the set of players be $\{1, 2\}$ with generic distinct elements *i* and *j*. There is a finite set *A* of feasible pure alternatives containing at least two elements. State θ that defines players' preferences is drawn from the set Θ . For profile $\theta \in \Theta$, denote by $u_i(\cdot, \theta)$ player *i*'s vNM utility function over *A*.

A social choice correspondence (SCC) f is a mapping $f : \Theta \to A$ (not necessarily single valued). Assume that $f(\Theta) = A$.

Nash implementation We are interested in *full* implementation where the set of SPE outcomes of the implementing mechanism coincides with the desired SCC in all states. A *normal form mechanism* $\Gamma = \langle S, g \rangle$ consists of a *strategy space* $S = S_1 \times S_2$ and an *outcome function* $g : S \to A$. Pair (Γ, θ) constitutes a *normal form game*. Denote by $NE(\Gamma, \theta)$ the set of Nash equilibria of (Γ, θ) . Then $s \in NE(\Gamma, \theta)$ if

 $u_i(g(s), \theta) \ge u_i(g(s'_i, s_i), \theta)$, for all $s'_i \in S_i$, for all i = 1, 2.

Mechanism Γ *Nash implements* f if $g(NE(\Gamma, \theta)) = f(\theta)$, for all $\theta \in \Theta$. If there is a mechanism Γ which Nash implements f, then f is *Nash implementable* (cf. Maskin 1999).

Commitment-free mechanism Suppose implementation is not restricted to normal form games. An *extensive mechanism with simultaneous moves* is defined by a

⁶ Baliga et al. (1997) study implementation when the designer is unable to commit to the mechanism. Maskin and Moore (1999) restrict implementation with a renegotiation procedure.

⁷ For an application of a similar idea to the principal-agent framework, see Ma et al. (1988).

triple $\overline{\Gamma} = \langle H, \mathbf{S}, \overline{g} \rangle$, where the set of finite histories H consists of strings of pairs of actions—histories—such that (i) $\emptyset \in H$, (ii) $h \in H$ and $h' \subset h$ imply $h' \in H$, (iii) $h^1 \subset \cdots \subset h^k$ for all $k = 1, 2, \ldots$ implies $h^k \in H$.⁸ At h, players' choice sets $S_h = S_{h1} \times S_{h2}$ satisfy the parity $S_h = \{s : (h, s) \in H\}$. The set of all strategies is $\mathbf{S} = \times_{h \in H} S_h$ and player *i*'s strategy set is $\mathbf{S}_i = \times_{h \in H} S_{hi}$. Each strategy $\mathbf{s} = (s(h))_{h \in H}$ of \mathbf{S} induces a unique *terminal* history

$$h(\mathbf{s}) = (s_1, \ldots) = (s(\emptyset), s(s(\emptyset)), s(s(\emptyset), s(s(\emptyset))), \ldots)$$

Denote by \overline{H} the set of terminal histories, endowed with the product topology. The outcome is conditional on the terminal history and is defined by the measurable function $\overline{g} : \overline{H} \to A$. Every nonterminal history *h* specifies a *subgame* $\overline{\Gamma}(h) = [H^h, \mathbf{S}^h, \overline{g}^h]$, the branch of $\overline{\Gamma}$ with origin *h*, where H^h, \mathbf{S}^h , and \overline{g}^h are herited from *H*, **S**, and \overline{g} , respectively, in the obvious way.⁹

Commitment to an action means that the action has physical consequences. Hence, mechanism Γ reflecting no-commitment is equivalent to saying that no single action has physical consequences. We say that the two mechanisms $\Gamma = [H, \mathbf{S}, \bar{g}]$ and $\Gamma' = [H', \mathbf{S}', \bar{g}']$ are equivalent, written $\Gamma = \Gamma'$, if there is a 1–1 relation φ between \bar{H} and \bar{H}' such that $\bar{g}(\bar{h}) = \bar{g}'(\varphi(\bar{h}))$ for all terminal histories $\bar{h} \in \bar{H}$.

Condition 1 Extensive mechanism $\overline{\Gamma}$ is commitment-free if $\overline{\Gamma}(h, s) = \overline{\Gamma}(h, s')$ for all $s, s' \in S_h$, for all $h \in H$.

That is, no single action should ever affect the continuation game. However, the condition does *not* require the continuation game form to be constant across histories. It is important that the condition does not impose a restriction on any choice that cannot be regretted afterwards, i.e., on infinite histories.

As an example of a nontrivial game that is commitment-free, let 1 and 2 move in a sequential order. At odd stages 1 makes an offer a and at even stages 2 accepts or rejects. If there is stage k from which onwards 1 always offers a and 2 always accepts, then a is implemented. In all other cases, some d is implemented. This "Coasian" bargaining game is studied by Vartiainen (2003).

Mechanism Γ is *trivial* if $\bar{g}(h(\mathbf{S}))$ is singleton.

Mechanism Γ has a *stage structure* if there is a set S_k such that $S_{(s_0,...,s_k)} = S_k$ for all $(s_0, \ldots, s_k) \in H$, for all $k = 0, \ldots$. Thus if Γ has a stage structure, then the structure of the continuation game is independent of the past history. Moreover, if Γ has a stage structure, then the set of terminal histories can be represented by a Cartesian product $\times_{k=0}^{K} S_k$, where *K* is finite or infinite. In such case, it is meaningful to say that a terminal history differs at a particular, say *k*th, stage.

Lemma 1 $\overline{\Gamma} = \langle H, \mathbf{S}, \overline{g} \rangle$ is commitment-free and nontrivial only if it has a stage structure, every terminal history is infinite and $\overline{g}(\overline{h}) = \overline{g}(\overline{h}')$ whenever terminal histories \overline{h} and \overline{h}' differ at most finitely many stages. Conversely, if $\overline{\Gamma}$ has these properties, then it is commitment-free.

⁸ For convenience, use the notation $h \subset h'$ if $h = (s, \ldots, s')$ and $h' = (s, \ldots, s', \ldots, s'')$.

⁹ For a more detailed construction of an extensive form, see, e.g., Osborne and Rubinstein (1995).

The proof is in Appendix. Since an outcome of a commitment-free mechanism Γ is independent of finite deviations, no nonterminal history gives any advice of the relevant features of the other player's strategies.

It is convenient to represent the extensive mechanism $[H, \mathbf{S}, \bar{g}]$ by the equivalent normal form $[\mathbf{S}, \bar{g} \circ \bar{h}]$. If a commitment-free mechanism $[\mathbf{S}, \bar{g} \circ \bar{h}]$ Nash implements f, then f is commitment-freely Nash implementable.

3 Structure of equilibria

This section studies the structure of Nash equilibria under a fixed θ since, due to the structure of strategy sets, the key issues with commitment-free implementation pop up already in this case (all proofs of this section are in Appendix). For simplicity, we remove θ from the notation.

Given a normal form mechanism $\Gamma = [S, g]$ and a payoff vector $v = (v_1, v_2)$, denote the outcomes that generate both i = 1, 2 payoff at least v_i by

$$D_{\Gamma}(v) = \{a \in g(S) : u_i(a) \ge v_i, \text{ for all } i\}.$$

Denote player *i*'s minmax-payoff by

$$v_i(\Gamma) = \min_{s_i} \max_{s_i} u_i(g(s_j, s_i)), \text{ for all } i.$$
(1)

Then set $D_{\Gamma}(v(\Gamma))$ is called the *feasible set* of normal form mechanism $\Gamma = [S, g]$.

First we show that the set of Nash equilibrium outcomes of a commitment-free game coincides with its feasible set.

Lemma 2 If *E* is the set of Nash equilibrium outcomes induced by a commitment-free mechanism $\bar{\Gamma} = [\mathbf{S}, \bar{g} \circ \bar{h}]$ under preferences θ , then $E = D_{\bar{\Gamma}}(v(\bar{\Gamma}))$.

The next lemma establishes a sufficient condition for commitment-free Nash implementation: The feasible set of a normal form game coincides with the set of Nash equilibrium outcomes of some commitment-free game.

Lemma 3 If, under preferences θ , $E = D_{\Gamma}(v(\Gamma))$ for a normal form mechanism Γ , then E is the set of Nash equilibrium outcomes induced by some commitment-free mechanism.

Constructing a commitment-free mechanism from a normal form mechanism is not completely straightforward. Since the set of outcomes is finite (unlike in standard repeated games where convex combinations of payoffs can be generated), and the set of strategies is a continuum, there are "too many" strategies to map to the set of outcomes. The main difficulty lies in doing this in a way that preserves commitment-freeness.

Thus, collecting Lemmas 2 and 3, we have a tight characterization. The set of Nash equilibrium outcomes under a commitment-free mechanism coincides with the feasible set of some normal form mechanism.

Theorem 1 *E* is the set of Nash equilibrium outcomes under a commitment-free mechanism under preferences θ if and only if $E = D_{\Gamma}(v(\Gamma))$ for a normal form mechanism Γ .

Thus to verify the implementability of set E via a commitment-free mechanism it is necessary and sufficient to focus on normal form mechanisms.

4 Implementation

Now we associate our observations concerning the structure of Nash equilibria of commitment-free games to implementation. Analogously to the previous notation, let, for any $\theta \in \Theta$,

$$D_{\Gamma}(v,\theta) = \{a \in g(S) : u_i(a,\theta) \ge v_i, \text{ for all } i\}.$$

Denote player *i*'s minmax-payoff in Γ under θ by

$$v_i(\Gamma, \theta) = \min_{s_i} \max_{s_i} u_i(g(s_j, s_i), \theta), \text{ for all } i$$

Then $D_{\Gamma}(v(\Gamma, \theta), \theta)$ is the feasible set of Γ under θ .

With this notation, it is now straightforward to extend Theorem 1 to the state space $\Theta.^{10}$

Theorem 2 Choice rule f is commitment-freely Nash implementable if and only if $f(\theta) = D_{\Gamma}(v(\Gamma, \theta), \theta)$, for all $\theta \in \Theta$, for some normal form mechanism Γ .

To present a simple sufficient condition, call choice rule f^{IR} individually rational correspondence with respect to $d \in A$ if

$$f^{IR}(\theta) = \{a \in A : u_i(a, \theta) \ge u_i(d, \theta), \text{ for all } i\}, \text{ for all } \theta.$$

Construct a simple veto-game form: all players suggest an outcome. If all agree on a, then a is implemented. Otherwise, d is implemented. This game form induces a feasible set $f^{IR}(\theta)$.

Proposition 1 Choice rule f^{IR} is commitment-freely Nash implementable, for any $d \in A$.

Not all Nash implementable rules are implementable by commitment-free mechanism. It is well known that the intersection of the strong *Pareto rule*,

$$f^{P}(\theta) = \{a \in A : \text{there is no } b \text{ s.t.} u_{i}(a, \theta) < u_{i}(b, \theta), \text{ for all } i\}, \text{ for all } \theta,$$

and f^{IR} , i.e.,

. .

$$f^{PIR}(\theta) = f^{P}(\theta) \cap f^{IR}(\theta), \text{ for all } \theta \in \Theta,$$

¹⁰ Since Theorem 1 is a necessary and sufficient condition under any θ .

is Nash implementable (cf. Maskin 1985). Consider the following scenario: there are at two players 1 and 2 and three outcomes *a*, *b*, and *c*.¹¹ Let rule f^{PIR} be defined on all strict preferences, and f^{IR} with respect to *c*. Suppose that there is a normal form mechanism $\Gamma = \langle S, g \rangle$ such that f^{PIR} always agrees with the feasible set of Γ . Let 1 and 2 agree on the top alternative under θ :

$$u_1(a,\theta) > u_1(b,\theta) > u_1(c,\theta),$$

$$u_2(a,\theta) > u_2(b,\theta) > u_2(c,\theta).$$

Since $f^{PIR}(\theta) = \{a\}$ and since f^{PIR} agrees with the feasible set of Γ , there is *i* such that

$$\min_{s_j} \max_{s_i} u_i(g(s_j, s_i), \theta) > u_i(b, \theta).$$
(2)

Define θ' such that

$$u_i(a, \theta') > u_i(b, \theta') > u_1(c, \theta'),$$

 $u_j(b, \theta') > u_j(a, \theta') > u_2(c, \theta').$

Then $f^{PIR}(\theta') = \{a, b\}$ (since c is Pareto dominated by both a and b which in turn are not Pareto dominated). By the definition of feasible set,

$$u_{i}(b,\theta) \geq \min_{s_{j}} \max_{s_{i}} u_{i}(g(s_{j},s_{i}),\theta')$$

=
$$\min_{s_{j}} \max_{s_{i}} u_{i}(g(s_{j},s_{i}),\theta)$$
(3)

But (2) and (3) are in conflict. Thus f^{PIR} cannot agree with the feasible set of Γ . Note that defining f^{IR} with respect to *a* or *b* would give the result as well (since then too $f^{PIR}(\theta') = \{a, b\}$).

Proposition 2 Rule f^{PIR} is not commitment-freely Nash implementable.

5 Complexity considerations

The previous section demonstrates that communication through costless actions creates more opportunities for coordination, and hence more undesirable equilibria. Thus the lack of commitment ability restricts choice rules that can be implemented.

Cooperation requires players to use contingent strategies. Such strategies are complex as they have to be conditioned on the past observations. But would players use such strategies if they have even small preference for simplicity? This section explores the consequences of complexity aversion to implementable choice rules.

To capture the complexity considerations, we view players implementing their strategies via *finite state* (Moore) *machines*.¹² Possible machines are defined with respect

¹¹ The argument can be easily adjusted to many players and many outcomes case.

¹² See Rubinstein (1986) for a seminal treatment of complexity issues in repeated games. For further discussion, see Abreu and Rubinstein (1988); Kalai and Stanford (1988).

to the underlying mechanism, say $\overline{\Gamma} = \langle H, \mathbf{S}, \overline{g} \rangle$. For simplicity, let $S_{ih} = S_i$ for all $h \in H$. Player *i*'s machine is a tuple $m_i = (Q_i, q_i^0, \sigma_i, \tau_i)$, where

- Q_i is the set of states,
- $q_i^0 \in Q_i$ is the initial state,
- $\sigma_i : Q_i \to S_i$ is the output function,
- $\tau_i : Q_i \times S_j \to Q_i$ is the transition function.

We analyze commitment-free games with choice set S_i , i = 1, 2, that is fixed across time. Machine m_i induces a strategy $\{s_i^t\}_{t=0}^{\infty}$, denoted by $s[m_i]$, if the following hold:

$$s_i^t = \sigma_i(q_i^t)$$
 and $q_i^{t+1} = \tau_i(q_i^t, s_i^t)$, for all $t = 0, 1, \dots$

That is, in the first stage, player *i*'s machine m_i is in state q_i^0 and generates a choice $s_i^0 = \sigma_i(q_i^0)$. Given the other player's choice s_j^0 and state q_i^0 , the machine moves to state $q_i^1 = \tau_i(q_i^0, s_j^0)$. In the second stage, the machine generates a choice $s_i^1 = \sigma_i(q_i^1)$, etc. The infinite stream of actions is defined recursively given the other player's stream of actions.

Denote the set of all machines available to player *i* by M_i .¹³ Since a machine specifies the player's action after all histories, it does implement a strategy. Without restrictions on Q, any strategy can be induced by some machine, hence the machine description is without loss of generality. The key feature of the machine description is that it allows explicit modeling of complexity considerations. As a proxy of the complexity of the strategy we take the size of the state space Q_i of the machine, or size of the machine for short, that implements it. We denote the size of machine m_i by $|m_i|$. Complexity aversion is then accounted by the negative sensitivity of players' payoffs to the number $|m_i|$. We assume that complexity aversion has the weakest possible form; it is *lexicographic*.

Formally, denote by s[m] the pair of strategies that are induced by the pair of machines $m = (m_1, m_2)$. Also denote by $q = (q_1, q_2)$ a typical pair of states in the combined state space $Q = Q_1 \times Q_2$ of the machines m, and $\sigma(q) = (\sigma_1(q_1), \sigma_2(q_2))$ the action taken by them under state $q = (q_1, q_2)$. Finally, denote by $\bar{s}[m] = (\bar{s}_1[m], \bar{s}_2[m]) = (s, s', s'', ...)$ the path of actions that materializes with s[m], and $\bar{q}[m] = (\bar{q}_1[m], \bar{q}_2[m]) = (q, q', q'', ...)$ the path of states that are activated along the play. That is, $\tau(q, \sigma(q)) = q'$ for any two successive elements q, q' in $\bar{q}[m]$ and $s = \sigma(q)$, and $s' = \sigma(q')$ for the corresponding successive elements s, s' in $\bar{s}[m]$. The implemented outcome depends only on $\bar{s}[m]$. To simplify notation, let us abbreviate and write $\bar{g}(m)$ instead of $\bar{g}(\bar{s}[m])$.

The time line of a commitment-free *machine game* is: (i) State $\theta \in \Theta$ is realized and a commitment-free mechanism $\overline{\Gamma}$ is designed. (ii) Players simultaneously choose their machines. (iii) the chosen machines induce strategies within the mechanism and an outcome is implemented accordingly. Machines $m = (m_1, m_2)$ form a *machine equilibrium* of $\overline{\Gamma}$ under θ , written simply $m \in NE(\overline{\Gamma}, \theta)$, if for all i = 1, 2 and for all

¹³ Indexed by *i* since *i*'s machine implements actions in the set S_i and takes inputs from the set S_j .

machines $m'_i \in M_i$,

$$u_i(g(m), \theta) > u_i(g(m_j, m'_i), \theta), \text{ or } u_i(g(m), \theta) = u_i(g(m_j, m'_i), \theta) \text{ and } |m'_i| \ge |m_i|.$$

That is, a player always prefers an outcome that generates him a higher utility but if there are two outcomes of equal value, then he prefers strategy that is less complex, i.e., inducable by a smaller machine.

The simplest strategy is the one that is implemented by a one-state machine. Many Nash equilibria are too complex to be implemented by such machines. Before establishing our main results, we make some preliminary observations concerning machine games.

Since the machines are finite, the play of the game must eventually repeat itself in a *cycle*. Thus given the machines $m = (m_1, m_2)$, the path of states $\bar{q}[m]$ can be decomposed into two phases: an *introductory phase* $IP = (q^0, \ldots, q^{T_0})$ of length $T_0 \ge 0$ and a *cycling phase* $CP = (q^{T_0+1}, \ldots, q^{T_0+T})$ of length $T \ge 1$.

Lemma 4 Any *m* induces a path of states $\bar{q}[m] = (IP, CP, CP, \ldots)$, where IP is a (possibly empty) finite sequence of elements in the combined state space Q of the machines m, and CP is a finite sequence of elements in Q, and $IP \cap CP = \emptyset$.

Proof Let $Q_1 \times Q_2 = Q$. Define $\tau = (\tau_1, \tau_2)$ and $\sigma = (\sigma_1, \sigma_2)$. Let $\phi = \tau \circ \sigma$: $Q \to Q$ such that $(\tau_1(q_1, \sigma_2(q_2)), \tau_2(q_2, \sigma_1(q_1))) = \phi(q_1, q_2)$. Starting from $q^0 = (q_1^0, q_2^0)$, generate $\phi(q^t) = q^{t+1} \in Q$, for $t = 0, \ldots$ Since all Q_1, Q_2 are finite sets, Q is finite. By the pigeonhole principle, there is the smallest $T_0 \in \{0, \ldots\}$ and $T \in \{1, \ldots\}$ such that $(q_1^{T_0}, q_2^{T_0}) = (q_1^{T_0+T}, q_2^{T_0+T})$. Since ϕ depends only on the current $q^t, (q_1^{T_0}, q_2^{T_0}) = (q_1^{T_0+rT}, q_2^{T_0+rT})$, for any $r \in \{0, \ldots\}$. Since there is no $T' < T_0$ with this property, $q^t \neq q^k$ for all $t < T_0 \leq k$.

Any state appearing in the introductory phase is visited only once, but any state appearing in the cycling phase is visited infinitely often.

Corollary 1 Any *m* induces a path of *i*'s states $\bar{q}_i[m] = (IP_i, CP_i, CP_i, \ldots)$, where IP_i is a (possibly empty) finite sequence of states in Q_i , and CP_i is a finite sequence of states in Q_i , and $IP_i \cap CP_i = \emptyset$.

First we note that a machine equilibrium cannot support states that are never used.

Lemma 5 Let $m \in NE(\overline{\Gamma}, \theta)$. Then Q_i coincides with the distinct elements in the path of states $\overline{q}_i[m]$.

Proof If a machine m_i of a player *i* has a state q_i that is not used when m_1 and m_2 operate, then *m* cannot be a Nash equilibrium since q_i can be eliminated without affecting the physical process, i.e., $\bar{s}[m]$. Since *i* prefers a smaller machine given that the outcome remains the same, this cannot be the case.

An immediate consequence of Corollary 1 and Lemma 5 is that any state in the *i*'s cycling phase can be reached from any state in Q_i . Since *i*'s current state q_i is always

either in the initial phase or in the cycling phase, player j needs to just follow the actions that he would have taken would he follow the equilibrium path. Then the play eventually goes through the all states in i's cycling phase.

Corollary 2 Let $m \in NE(\bar{\Gamma}, \theta)$. If $q_i \in CP_i$, then, for any $q'_i \in Q_i$ there is $\{s_j^1, \ldots, s_j^K\}$ of *j*'s actions such that $\tau_i(s_j^k, q_i^k) = q_i^{k+1}$ for all $k = 1, \ldots, K$ and $q_i^1 = q'_i$ imply $q_i^{K+1} = q_i$.

The next lemma shows that if there is a state in player i's machine such that once the state is received i will never change his action no matter what choices the other player makes, then that state is the unique element in the cycling phase of i.

Lemma 6 Let $m \in NE(\overline{\Gamma}, \theta)$. If, for any q_i , all finite sequences $\{s_j^1, \ldots, s_j^K\}$ of j's actions imply $\sigma_i(q_i) = \sigma_i(q_i^{K+1})$ for q_i^1, \ldots, q_i^{K+1} such that $\tau_i(s_j^k, q_i^k) = q_i^{k+1}$ for all $k = 1, \ldots, K$ and $q_i^1 = q_i$, then $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$.

Proof Suppose that the conditions hold for m_i . Then, given that q_i is reached, no sequence of actions of j changes the action made by i at q_i . Thus unless $\tau'(s_j, q_i) = q_i$ for all s_j , one could replace τ with τ' that would do that and remove all states reached after q_i in $\bar{q}_i[m]$ without affecting its performance, i.e., the path $\bar{s}[m]$. But this this violates the assumption that m forms a Nash equilibrium.

5.1 Nash implementability implies machine implementability

Moore and Repullo (1990) and Dutta and Sen (1991) establish a necessary and sufficient condition for two-person Nash implementation. The condition, which is called $\mu 2$, is defined with respect to a collection of sets $\{C_i(\theta, a)\}$ such that $C_i(\theta, a) \subseteq L_i(\theta, a)$ for all $a \in f(\theta)$, for all $\theta \in \Theta$, for all i = 1, 2, and a function $e : \Theta \times A \times \mathbb{N} \times A \to A$ such that $e(\theta^1, a^1, \theta^2, a^2) \in C_1(\theta^2, a^2) \cap C_2(\theta^1, a^1)$.

Choice rule f satisfies $\mu 2$ if there are sets $\{C_i(\theta, a)\}$ and function e such that the following hold:¹⁴

- (i) If an outcome $a\theta$ -maximizes both players' payoff in A, then $a \in f(\theta)$.
- (ii) If an outcome $a\theta$ -maximizes *i*'s payoff in A, and *j*'s in $C_i(\theta', a')$, then $a \in f(\theta)$.
- (iii) If an outcome $a\theta$ -maximizes *i*'s payoff in $C_i(\theta', a')$ and *j*'s in $C_j(\theta', a')$, then $a \in f(\theta)$.
- (iv) If an outcome $a = e(\theta^1, a^1, \theta^2, a^2)\theta$ -maximizes 1's payoff in $C_1(\theta^2, a^2)$, and 2's in $C_2(\theta^1, a^1)$, then $a \in f(\theta)$.

Thus if f is implementable, then there are sets $\{C_i(\theta, a)\}$ such that (i)–(iv) are met by f. To prove the other direction, Moore and Repullo (1990) and Dutta and Sen (1991) employ the following *canonical mechanism* Γ^* : Both players i choose $(\theta^i, a^i, n^i, b^i) \in \Theta \times A \times \mathbb{N} \times A$ such that $a^i \in f(\theta^i)$. The outcome function g^* is defined by the following rules:¹⁵

¹⁴ Recall that we assume $f(\Theta) = A$.

¹⁵ Ties are broken in favor of 1.

(a) If $(\theta^1, a^1) = (\theta^2, a^2) = (\theta, a)$ and $n^i \ge n^j = 0$, then $g^*(s) = a$. (b) If $(\theta^1, a^1) \ne (\theta^2, a^2)$ and $n^i > n^j = 0$, then

$$g^*(s) = \begin{cases} b^i, & \text{if } b^i \in C_i(\theta^j, a^j), \\ e(\theta^1, a^1, \theta^2, a^2), & \text{otherwise.} \end{cases}$$

(c) If $n^i > n^j > 0$, then $g^*(s) = b^i$.

We show that any Nash implementable choice rule can be implemented via a commitment-free machine game by using a dynamic version of the canonical mechanism. Given the collection $\{C_i(\theta, a)\}$ of sets and the choice rule f, construct the following commitment-free mechanism Γ^M . In each stage $t = 0, 1, \ldots$, each player i announces a tuple $(\theta_t^i, a_t^i, n_t^i, b_t^i) \in \Theta \times A \times \mathbb{N} \times A$ such that $a_t^i \in f(\theta_t^i)$. Outcome function g^M is defined on the set of terminal histories $(\Theta \times A \times \mathbb{N} \times A)^{\infty}$, and it aggregates the content of strategies by using the following measures.

Given strategy $\{s_t^i\}_{t=0}^{\infty} = \{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^{\infty}$, denote a set of messages that *i* chooses infinitely many times by

$$X^{i} = \left\{ (\theta^{i}, a^{i}) : \sum_{t=0}^{\infty} \mathbf{1} \{ (\theta^{i}_{t}, a^{i}_{t}) = (\theta^{i}, a^{i}) \} = \infty \right\}.$$

Since *i*'s machine is finite, there has to be an s^i that is chosen infinitely many times. Thus X^i is nonempty. Denote the highest integer that *i* announces infinitely many times by

$$\bar{n}^i = \max\left\{n : \sum_{t=0}^{\infty} \mathbf{1}\{n_t^i = n\} = \infty\right\} \cup \{0\}.$$

Since sequence $\{n_t^i\}_t$ is countable, there is a well defined $\bar{n}^i \in \mathbb{N}$.¹⁶ Finally, for any $b \in A$, let

$$\rho(b) = \lim_{T \to \infty} \frac{\sum_{t=0}^{T} \mathbf{1}\{b_t^i = b \text{ and } n_t^i = \bar{n}^i\}}{\sum_{t=0}^{T} \mathbf{1}\{n_t^i = \bar{n}^i\}}.$$

That is, $\rho(b)$ is the limit of the frequency of times that *i* chooses *b* together with \bar{n}^i . Since \bar{n}^i is well defined, and since the sequence $\{n_t^i, b_t^i\}_t$ is generated by a finite machine, the limit exists for each *b*. Let

$$b^i = \arg \max_{b \in A} \rho(b).$$

Since A is finite, a maximum exists.

¹⁶ Otherwise $\{n_t^i\}$ would contain infinitely many disjoint subsequences. Since $\{n_t^i\}$ is countable, this is not feasible.

Thus strategy $\{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^{\infty}$ uniquely defines a triple $(X^i, \bar{n}^i, b^i) \in 2^{\Theta \times A} \times \mathbb{N} \times A$. If X^i contains a single element, we slightly abuse the notation and denote this element by (θ^i, a^i) . Given $\{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^{\infty}$, the outcome function g^M is defined as follows.¹⁷

(A) If $|X^1| = |X^2| = 1$, $(\theta^1, a^1) = (\theta^2, a^2)$, and $\bar{n}^i \ge \bar{n}^j = 0$, then $g^M(s) = a$. (B) If $|X^1| = |X^2| = 1$, $(\theta^1, a^1) \ne (\theta^2, a^2)$, and $\bar{n}^i > \bar{n}^j = 0$, then

$$g^{M}(s) = \begin{cases} b^{i}, & \text{if } b^{i} \in C_{i}(\theta^{j}, a^{j}), \\ e(\theta^{1}, a^{1}, \theta^{2}, a^{2}), & \text{otherwise.} \end{cases}$$

(C) If $|X^1| = |X^2| = 1$, and $\bar{n}^i > \bar{n}^j > 0$, then $g^M(s) = b^i$. (D) If $|X^j| > |X^i| = 1$, then $g^M(s) = a^i$. (E) If $|X^i| > |X^j| > 1$, then $g^M(s) = b^i$.

 Γ^M is a commitment-free mechanism: the continuation game is always independent of the players' current actions. No finite change in choices affect the implemented outcome. Any deviation with material consequences must induce a play path that differs in infinitely many times from the initial path.

Note the analogy between (a–c) and (A–C) in g^* and g^M , respectively. The additional features of g^M are D and E that relate to the cases when X_1 and X_2 are not single valued. This can be interpreted as an integer game-construction in a new dimension.

Our aim is to show that the commitment-free mechanism Γ^M and the canonical mechanism Γ^* induce the same set of Nash equilibria in each state, when the former is played via machines.

Lemma 7 Let Γ^* Nash implement f. Then Γ^M implements f in machine equilibrium.

The "if"-part follows from the observation that a constant strategy that chooses the same message unconditionally of the history of the play (and the machine inducing it uses only one state) is a best response against a constant strategy. When such strategies are used, only Rules (A–C) of game Γ^M apply. Thus a pair of constant strategies form a machine equilibrium in Γ^M if the analogous strategies form a Nash equilibrium in Γ^* .

The more involved "only if"-part of the proof is based on the idea that unless there is a state in the player's cycling phase when he is totally unresponsive to changes in the other player's strategy, then he can be "fooled" into a cycle that, by Rule D or E, lets his opponent choose the implemened outcome (Claim 2 in the proof). This implies both players have a tendency to choose unconditional, constant strategies. Comparing game Γ^M (without D and E that employ nonconstant strategies) and Γ^* one observes that constant strategies in the former are analogous to the stratgies of the latter. Thus, given that players use constant strategies, equilibrium behaviors in the former and in the latter are also analogous. Since Γ^* is a canonical mechanism in the context of two-player Nash implementation, the following sufficient condition is thus obtained.

¹⁷ Break integer ties in favor of 1.

Theorem 3 Choice rule f is implementable in machine equilibrium of a commitmentfree game if f is Nash implementable.

It is well known in the literature that small changes in the underlying implementation framework can have a large impact on the class of implementable rules. How much more can be implemented in machine equilibrium than in Nash equilibrium is left as an open problem.

Caveat Theorem 7 is sensitive to the assumption that complexity costs are lexicographic. With any strictly positive costs, it may no longer hold that it is profitable for a player i to induce the other player j into a cycle á la Claim 2. Without such property, j may not need to protect himself by committing to a constant strategy in the cycling phase. It is also clear that with high complexity costs the result would again apply since players would be inhibited from using conditional strategies. Exploring the consequences of complexity costs in the middle area should be interesting.

5.2 Repeated implementation under limit of the means

While analytically unproblematic, that commitment-free mechanisms never end is inconvenient from the heuristic perspective. A natural way to interpret this feature is to view a commitment-free game as a reduced form description of a repeated interaction where the planner seeks to implement an outcome with completely patient players in every period $0, 1, \ldots$. Let us formalize complete patience in the language of the limit of the means-criterion which places no weight on the payoffs that are materialized in finitely many periods.

Formally, given a stream of outcomes $\bar{a} = \{a_t\}_{t=0}^{\infty} \in A$, the limit of the means-criterion gives player *i* a payoff

$$u_i(\bar{a},\theta) = \lim_{T \to \infty} \frac{\sum_{t=0}^T u_i(a,\theta)}{T},$$

whenever the limit is defined. We now show how the canonical commitment-free mechanism can be altered so that it matches the current repeated implementation framework.

Construct a game Γ^r that implements an outcome in A in every period $T = 0, \ldots$, conditional on the submitted messages in periods $0, \ldots, T$. Let strategy sets be defined as in the game Γ^M : at each period t, player i announces a tuple $(\theta_t^i, a_t^i, n_t^i, b_t^i) \in \Theta \times A \times \mathbb{N} \times A$ such that $a_t^i \in f(\theta_t^i)$.

Given the announcements $\bar{s}^T = \{s_t^i\}_{t=0}^T = \{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^T$ until period T, denote a set of messages that i has chosen in proportion higher than $1/(\sqrt{T} - C)$ by

$$X^{i}(T) = \left\{ (\theta^{i}, a^{i}) : \sum_{t=0}^{T} \mathbf{1}\{ (\theta^{i}_{t}, a^{i}_{t}) = (\theta^{i}, a^{i}) \} \ge \sqrt{T} - C \right\}, \text{ for } C = |A \times \Theta|^{2}.$$

Set $\{(\theta_t^i, a_t^i)\}_{t=0}^T$ contains at most $|A \times \Theta|$ distinct elements. Thus, when *T* is higher than $|A \times \Theta|^2$, $\{(\theta_t^i, a_t^i)\}_{t=0}^T$ repeats some element at least \sqrt{T} times. With this choice of constant *C*, set $X^i(T)$ is nonempty for all *T*. Note that *C* does not affect the convergence properties of $X^i(T)$.

Denote the highest integer that *i* has announced in proportion of at least $1/\sqrt{T}$ of all messages in periods $0, \ldots, T$,

$$\bar{n}^i(T) = \max\left\{n : \sum_{t=0}^T \mathbf{1}\{n_t^i = n\} \ge \sqrt{T}\right\} \cup \{0\}.$$

Finally, let, as in the previous section,

$$\rho(b,T) = \frac{\sum_{t=0}^{T} \mathbf{1}\{b_t^i = b \text{ and } n_t^i = \bar{n}^i\}}{\sum_{t=0}^{T} \mathbf{1}\{n_t^i = \bar{n}^i\}},$$

and

$$b^i(T) = \arg \max_{b \in A} \rho(b, T).$$

Thus, as *T* becomes large, a strategy $s^T = \{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^T$ well defines a triple $(X^i(T), \bar{n}^i(T), b^i(T)) \in 2^{\Theta \times A} \times \mathbb{N} \times A$. Moreover, when $|X^i(T)| = 1$, it well defines (θ^i, a^i) . An outcome is implemented in every period. Given, $\{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^T$, for any $T = 0, 1, \ldots$, the outcome function *o* is defined on $\bigcup_{T=0}^{\infty} \{(\theta_t^i, a_t^i, n_t^i, b_t^i)\}_{t=0}^T$ as follows.¹⁸

- 1. If $|X^{1}(T)| = |X^{2}(T)| = 1$, $(\theta^{1}, a^{1}) = (\theta^{2}, a^{2})$, and $\bar{n}^{i}(T) \ge \bar{n}^{j}(T) = 0$, then $o(s^{T}) = a$.
- 2. If $|X^1(T)| = |X^2(T)| = 1$, $(\theta^1, a^1) \neq (\theta^2, a^2)$, and $\bar{n}^i(T) > \bar{n}^j(T) = 0$, then

$$o(s^{T}) = \begin{cases} b^{i}(T), & \text{if } b^{i}(T) \in C_{i}(\theta^{j}, a^{j}), \\ e(\theta^{1}, a^{1}, \theta^{2}, a^{2}), & \text{otherwise.} \end{cases}$$

- 3. If $|X^{1}(T)| = |X^{2}(T)| = 1$, and $\bar{n}^{i}(T) > \bar{n}^{j}(T) > 0$, then $o(s^{T}) = b^{i}$.
- 4. If $|X^{j}(T)| \neq |X^{i}(T)| = 1$, then $o(s^{T}) = a^{i}$.
- 5. If $|X^{i}(T)| > |X^{j}(T)| > 1$, then $o(s^{T}) = b^{i}$.

Denote by ΔA the standard |A|-simplex. Given a strategy $\{s^t\}_{t=0}^{\infty}$, denote the average implemented outcome by

$$\bar{o}(\bar{s}^T) = \frac{\sum_{t=0}^T r(s^t)}{T} \in \Delta A.$$

¹⁸ Break integer ties in favor of 1.

Denote by $\{s^t[m]\}_{t=0}^{\infty}$ the strategy that is induced by machines $m = (m_1, m_2)$. Since *m* is finite,

$$\bar{o}(\bar{s}^{\infty}[m]) = \lim_{T \to \infty} \bar{r}(\bar{s}^{T}[m])$$

is well defined.

The next lemma argues that the limit of the average outcome induced by m coincides exactly with the outcome that m induces in g^M . Thus $\bar{o}(\bar{s}^{\infty}[m]) \in A$ for all m.

Lemma 8 $\bar{o}(\bar{s}^{\infty}[m]) = g^M(\bar{s}[m]), \text{ for all } m \in M_1 \times M_2.$

Proof Take any m. Denote

$$\bar{o}(\bar{s}_{t_0}^T[m]) = \frac{\sum_{t=t_0}^T r(s^t[m])}{T - t_0}$$

Since finite changes do not affect the average of an infinite sample,

$$\bar{o}(\bar{s}^{\infty}[m]) = \bar{o}(\bar{s}^{\infty}_0[m]) = \bar{o}(\bar{s}^{\infty}_{t_0}[m]), \text{ for all } t_0 = 1, 2, \dots$$

Since, by construction

$$\lim_{T \to \infty} X^{i}(T) = X^{i},$$
$$\lim_{T \to \infty} \bar{n}^{i}(T) = \bar{n}^{i},$$
$$\lim_{T \to \infty} b^{i}(T) = b^{i},$$

we have

$$\lim_{t \to \infty} o(s^{t+k}[m]) = g^M(\bar{s}[m]), \quad \text{for all } k = 1, 2, \dots$$

Thus, by transfinite induction,

$$\bar{o}(\bar{s}^{\infty}[m]) = \lim_{t_0 \to \infty} \bar{o}(\bar{s}^{\infty}_{t_0}[m])$$
$$= g^M(\bar{s}[m]).$$

 \Box

Machines $m = (m_1, m_2)$ form a machine equilibrium in Γ^r if for all *i* and for all machines m'_i ,

$$u_i(\bar{o}(\bar{s}^{\infty}[m]), \theta) > u_i(\bar{o}(\bar{s}^{\infty}[m_j, m'_i]), \theta) \text{ or} u_i(\bar{o}(\bar{s}^{\infty}[m]), \theta) = u_i(\bar{o}(\bar{s}^{\infty}[m_j, m'_i]), \theta) \text{ and } |m'_i| \ge |m_i|.$$

By this, Lemma 8, and Theorem 7 we immediately obtain a full characterization of the choice rules that are Nash implementable via machines in a repeated game with a limit of the means criterion:

Corollary 3 *Choice rule f is implementable in machine equilibrium of a (repeated) commitment-free game under limit of the means criterion if f is Nash implementable.*

Note that Γ^r is commitment-free as it has a stage structure, and only infinite deviations may affect the outcome. Thus it is only the *interpretation* of the moves that differentiates Γ^r from Γ^M , not the strategic substance of the game.

6 Discussion

That a game form fully describes relevant interaction among individuals—the property of games that is firmly rooted in the game theory language—is problematic for any attempts to model genuine *non*commitment. Noncommitment would mean, heuristically, that even the structure of the game form should not be unchangeable. In particular, one should be able to take back one's move that is found desirable. But this would require a change in the game form which is given from outside and cannot be affected by the players.

This paper proposes a way to analyze noncommitment in games. Our approach is inspired by Aumann and Hart (2003). We impose a commitment-freeness condition on game forms. The condition requires that the continuation game form should not be sensitive to actions, and should not require irrevocable actions. Any commitment-free game continues forever.¹⁹

Commitment-free interaction can be interpreted as cheap talk without a detailed view how the actions transform into outcomes. The question is rather of what communication strategies are consistent with what outcomes given that there is no final word.

Communication creates new opportunities to coordinate and, hence, new equilibria. (Nash) implementation with commitment-free games is, therefore, difficult. However, some equilibria require one to use complex strategies. The planner may therefore be able to take advantage of players' complexity aversion. We show that if players are complexity averse in the lexicographic sense, then implementation is not more difficult than Nash implementation. Hence commitment can be replaced with complexity aversion in the context of Nash implementation. The framework can be interpreted as a repeated implementation problem where the planner implements an outcome in every period 0, 1, . . . according to predesigned rules, and players are infinitely patient.

Appendix

Proof of Lemma 1 Let $\overline{\Gamma} = [H, \mathbf{S}, \overline{g}]$ be commitment-free and nontrivial. Since $\Gamma(h, s) = \Gamma(h, s')$ it follows that $S_{h,s} = S_{h,s'}$, for all $s, s' \in S_h$, for all $h \in H$.

¹⁹ Cf. the Bill Murray movie "Groundhog Day".

By induction, **S** has a stage structure: The set of terminal histories coincides with $\times_{k=1}^{K} S_k$ where *K* may be finite or infinite. Suppose *K* is finite. Since terminal history $(h, s) \in \times_{k=1}^{K} S_k$ is associated to an outcome a(h, s), subgame $\overline{\Gamma}(h, s)$ is trivial. But since $\overline{\Gamma}(h, s) = \overline{\Gamma}(h, s')$ by commitment-freeness, a(h, s) = a(h, s'), for all $s, s' \in S_K$. Thus also subgame $\overline{\Gamma}(h)$ is trivial. By induction, $\overline{\Gamma}(\emptyset)$, or $\overline{\Gamma}$, is trivial, a contradiction. Thus *K* must be infinite.

Identify terminal histories $\bar{h} = (s_1, ...)$ and $\bar{h}' = (s'_1, ...)$. Suppose that \bar{h} and \bar{h}' differ in stages $k_1, ..., k_L$ where L is finite. Since h and \bar{h}' agree up to stage $k_1 - 1$, it follows by commitment-freeness that $\bar{\Gamma}(s_1, ..., s_{k_1}) = \bar{\Gamma}(s'_1, ..., s'_{k_1})$. Since \bar{h} and \bar{h}' agree between stages $k_1 + 1$ and k_2 , it follows that $\bar{\Gamma}(s_1, ..., s_{k_2-1}) = \bar{\Gamma}(s'_1, ..., s'_{k_2-1})$. By induction, $\bar{\Gamma}(s_1, ..., s_{k_L}) = \bar{\Gamma}(s'_1, ..., s'_{k_L})$. Since \bar{h} and \bar{h}' agree from k_L onwards, $\bar{g}(\bar{h}) = \bar{g}(\bar{h}')$.

Conversely, if $\bar{\Gamma} = [H, \mathbf{S}, \bar{g}]$ has a stage structure, then for any h, there is an identity relation between $H|_{(h,s)}$ and $H|_{(h,s')}$, where $\bar{\Gamma}(h, s) = [H|_{(h,s)}, \mathbf{S}|_{(h,s)}, \bar{g}|_{(h,s)}]$ and $\bar{\Gamma}(h, s') = [H|_{(h,s')}, \mathbf{S}|_{(h,s')}, \bar{g}|_{(h,s')}]$. If $\bar{\Gamma}$ does not have finite terminal histories, and has the property that $\bar{g}(\bar{h}) = \bar{g}(\bar{h}')$ whenever terminal histories \bar{h} and \bar{h}' differ at most finitely many stages, then also $\bar{g}|_{(h,s)}(\bar{h}|_{(h,s)}) = \bar{g}|_{(h,s)}(\bar{h}|_{(h,s')})$, where $(h, s, \bar{h}|_{(h,s)})$ is a terminal history, since $\bar{g}(h, \bar{h}|_h) = \bar{g}|_h(\bar{h}|_h)$ for all h and since $(h, s, \bar{h}|_{(h,s)})$ and $(h, s', \bar{h}|_{(h,s')})$ differ in only one stage.

Proof of Lemma 2 Let set *E* form a Nash equilibria of a commitment-free game $\overline{\Gamma} = [\mathbf{S}, \overline{g} \circ \overline{h}]$. Let

$$v_i(\bar{\Gamma}) = \min_{\mathbf{s}_j} \max_{\mathbf{s}_i} u_i(\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))).$$

By Lemma 1, no finite deviation affects the form of the continuation game. Thus also

$$v_i(\bar{\Gamma}) = \min_{\mathbf{s}_j} \max_{\mathbf{s}_i} u_i(\bar{g}^h(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))), \text{ for any } h \in H.$$
(4)

Denote *i*'s best response against s_i at *h* by

$$BR_i^h(\mathbf{s}_j) \in \arg\max_{\mathbf{s}_i} u_i(g^h(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))).$$

Since A is finite, a best response exists.

We show that $D_{\overline{\Gamma}}(v(\overline{\Gamma})) = E$.

 $D_{\overline{\Gamma}}(v(\overline{\Gamma})) \subseteq E$: Take any $a \in D_{\overline{\Gamma}}(v(\overline{\Gamma}))$ Then there is a terminal history $\overline{s} = (\overline{s}_1, \overline{s}_2)$ such that $g(\overline{s}) = a$. We prove that history \overline{s} is playable in Nash equilibrium.

Construct, by (4), a *trigger* strategy of $j \neq i$ that punishes *i* with \mathbf{s}_j^h whenever *i* deviates from the path \bar{s} at *h*, where

$$\mathbf{s}_j^h \in \arg\min_{\mathbf{s}_j} u_i(\bar{g}^h(\bar{h}(\mathbf{s}_j, BR_i^h(\mathbf{s}_j)))).$$

Since A is finite, such \mathbf{s}_{j}^{h} exists. By Lemma 1, a finite deviation affects the payoff only via the effect on players' strategies. Thus player *i*'s payoff from a deviation is at most

$$v_i(\bar{\Gamma}) = u_i(\bar{g}^h(\bar{h}(\mathbf{s}_i^h, BR_i^h(\mathbf{s}_i^h)))).$$

Thus a deviation cannot be profitable.

 $E \subseteq D_{\overline{\Gamma}}(v(\overline{\Gamma}))$: No $a \notin D_{\overline{\Gamma}}(v(\overline{\Gamma}))$ can be supported in Nash equilibrium since otherwise some player is receiving less than what he can guarantee himself by responding optimally to the other player's strategy.

Proof of Lemma 3 Let $E = D_{\Gamma}(v(\Gamma))$ for normal form game $\Gamma = [S, g]$. Construct a commitment-free game as follows. Let the choice set of *i* be an infinite replica of S_i s. Define, for any $\bar{s}_i = \{\bar{s}_i(t)\}_{t=1}^{\infty} \in S_i^{\infty}$ and for any *i*, a set $C_i(\bar{s}_i)$ that contains all s_i 's that are repeated in \bar{s}_i for infinitely many times, i.e.,

$$C_i(\bar{s}_i) = \left\{ s_i : \left| \bar{s}_i^{-1}(s_i) \right| = \infty \right\}.$$

For all *i*, there is a function²⁰ $c_i : S_i^{\infty} \to S_i$ such that, for some $s_i^* \in S_i$,

$$c_i(\bar{s}_i) \in \begin{cases} C_i(\bar{s}_i), & \text{if } C_i(\bar{s}_i) \neq \emptyset, \\ s_i^*, & \text{if } C_i(\bar{s}_i) = \emptyset. \end{cases}$$

That is, c_i specifies an element from the subset of S_i that are repeated in \bar{s}_i for infinitely many times. If there is no such element in S_i , then c_i specifies some (any) fixed element of S_i . Note that $c_i(s_i, s_i, ...) = s_i$, for any $s_i \in S_i$.

To construct \bar{g} , denote the set of player *i*'s history dependent strategies by $\mathbf{S}_i = S_i \times \bigcup_{k=0}^{\infty} S^k$, where $S = S_1 \times S_2$, and by $\bar{h}(\mathbf{s}_1, \mathbf{s}_2)$ the realized path (\bar{s}_1, \bar{s}_2) in S^{∞} , as a result of using strategy $(\mathbf{s}_1, \mathbf{s}_2) \in \mathbf{S}_1 \times \mathbf{S}_2$. Construct an outcome function \bar{g} on S^{∞} such that

$$\bar{g}(\bar{s}_1, \bar{s}_2) = g(c_1(\bar{s}_1), c_2(\bar{s}_2)), \text{ for all } (\bar{s}_1, \bar{s}_2) \in S^{\infty}.$$
 (5)

Now $\overline{\Gamma} = [\mathbf{S}, \overline{g} \circ \overline{h}]$ is a commitment-free game as no finite change of moves affects the continuation game.

We now show that $E = D_{\overline{\Gamma}}(v(\overline{\Gamma}))$. By Lemma 2, it suffices to show that

$$v_i(\Gamma) = v_i(\bar{\Gamma}). \tag{6}$$

Denote by $\mathbf{1}(s_i) \in \mathbf{S}_i$ the simple strategy of *i* that chooses s_i after *all* histories. Denote also $\mathbf{1}(s) = (\mathbf{1}(s_1), \mathbf{1}(s_2)) \in \mathbf{S}$ and $\mathbf{1}(s_j) \in \mathbf{S}_j$. Since $\bar{h}(\mathbf{s}) \in S_1^{\infty} \times S_2^{\infty}$, and $c_i^{-1}(\cdot)$ partitions S_i^{∞} such that $(s_i, s_i, \ldots) \in c_i^{-1}(s_i)$ for all $s_i \in S_i$ for all *i*, it follows that for any $\mathbf{s} \in \mathbf{S}$ there is $s \in S$ such that $\bar{h}(\mathbf{s}) \in c^{-1}(s)$. Since, $(s, s, \ldots) \in c^{-1}(s)$ it follows

²⁰ Assuming the axiom of choice.

by (5) that $\bar{g}(\bar{h}(\mathbf{s})) = \bar{g}(\bar{h}(\mathbf{1}(s)))$. Similarly, for any $\mathbf{s}_j \in \mathbf{S}_j$ there is $s_j \in S_j$ such that $\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{1}(s_i))) = \bar{g}(\bar{h}(\mathbf{1}(s_j), \mathbf{1}(s_i)))$. Hence

$$\min_{\mathbf{s}_j} \max_{\mathbf{s}_i} u_i(\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))) \ge \min_{\mathbf{s}_j} \max_{s_i} u_i(\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{1}(s_i))))$$
$$= \min_{s_j} \max_{s_i} u_i(\bar{g}(\bar{h}(\mathbf{1}(s_j), \mathbf{1}(s_i)))).$$
(7)

Similarly,

$$\min_{\mathbf{s}_j} \max_{\mathbf{s}_i} u_i(\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))) \le \min_{s_j} \max_{\mathbf{s}_i} u_i(\bar{g}(\bar{h}(\mathbf{1}(s_j), \mathbf{s}_i)))$$
$$= \min_{s_i} \max_{s_i} u_i(\bar{g}(\bar{h}(\mathbf{1}(s_j), \mathbf{1}(s_i)))).$$

Thus, by (5),

$$\min_{s_j} \max_{s_i} u_i(\bar{g}(\bar{h}(\mathbf{1}(s_j), \mathbf{1}(s_i)))) = \min_{s_j} \max_{s_i} u_i(g(c_j(s_j, s_j, \ldots), c_i(s_i, s_i, \ldots)))$$
$$= \min_{s_j} \max_{s_i} u_i(g(s_j, s_i)),$$

and consequently,

$$\min_{\mathbf{s}_j} \max_{\mathbf{s}_i} u_i(\bar{g}(\bar{h}(\mathbf{s}_j, \mathbf{s}_i))) = \min_{s_j} \max_{s_i} u_i(g(s_j, s_i)),$$

as desired.

Proof of Lemma 7 Fix θ . By assumption, $f(\theta) = g^*(NE(\Gamma^*, \theta))$, for all θ . We show that $g^*(NE(\Gamma^*, \theta)) = g^M(NE(\Gamma^M, \theta))$.

" \subseteq ": Take $a \in g^*(NE(\Gamma^*, \theta))$. Then $a \in f(\theta)$. Construct m_1 and m_2 with $|m_1| = |m_2| = 1$ that choose, in every stage $t, s^i = (\theta, a, 0, b^i)$. Since $|m_i| = 1$, there are no cycles and hence $X_i = \{(\theta, a)\}$ for all *i*. By Rule A, *a* is implemented.

Deviation by, say, player 1 such that $X'_1 = \{(\theta^1, a^1, n^1, b^1)\}$ would, by Rule B, implement either b^1 or $e(\theta^1, a^1, \theta, a)$, both in $L_1(\theta, a)$. Thus the deviation does not pay. Deviation such that $|X'_1| > 1$ would, by Rule D, implement a. Thus the deviation does not pay. Hence $a \in g^M(NE(\Gamma^M, \theta))$.

"⊇": Suppose, to the contrary, that $m \in NE(\Gamma^M, \theta)$ but $g^M(m) \notin f(\theta) := g^*(NE(\Gamma^*, \theta))$.

Claim 1 For each *i*, if there is $q_i \in Q_i$ such that $\tau_i(s_j, q_i) = q_i$ for all $s_j \in S_j$, then q_i is the unique element in *i*'s cycling phase.

Proof If such q_i is in the initial phase, then, by Lemmas 5 and 6, the cycling phase would never be reached. Hence q_i is in the cycling phase. Once it is reached, no other element in the cycling phase is reached. Hence the length of the cycling phase is 1.

Claim 2 If there is no $q_i \in Q_i$ such that $\tau_i(s_j, q_i) = q_i$ for all $s_j \in S_j$, then $g^M(m)$ maximizes j's payoff, for $j \neq i$.

Proof Suppose that for all $q'_i \in Q_i$ there is $s_j \in S_j$ such that $\tau_i(s_j, q'_i) \neq q'_i$, for $j \neq i$. This holds, in particular, for q_i in the cycling phase. By Lemma 6 and Claim 1, j can induce i's action $\sigma_i(q_i^{K_1+1}) \neq \sigma_i(q_i)$ by choosing a sequence $\{s_j^1, \ldots, s_j^{K_1}\}$ of actions such that $\tau_i(s_j^k, q_i^k) = q_i^{k+1}$ for all $k = 1, \ldots, K_1$, and $q_i^1 = q_i$. Take *any* finite sequence of actions $\{s_j^{K_1+K_2+1}, \ldots, s_j^{K_1+K_2}\}$. By Corollary 2, there is a continuation sequence of actions $\{s_j^{K_1+K_2+1}, \ldots, s_j^{K_1+K_2+K_3}\}$ such that $\tau_i(s_j^k, q_i^k) = q_i^{k+1}$ for all $k = 1, \ldots, K_1 + K_2 + K_3$, and $q_i^{K_1+K_2+K_3+1} = q_i^1$. That is, $\{q_i^1, \ldots, q_i^{K_1+K_2+K_3}\}$ forms a cycle such that $\{\sigma_i(q_i^1), \ldots, \sigma_i(q_i^{K_1+K_2+K_3})\}$ is not a singleton, implying $|X^i| > 1$. By construction, $|X^i| \leq |Q_i|$. Choosing $\{s_j^{K_1+1}, \ldots, s_j^{K_1+K_2}\}$ of *distinct* actions and $K_2 > |Q_i|$, we have $|X^j| > |X^i| > 1$. Applying Rule E, and choosing b^j that maximizes j's payoff, it follows that unless $g^M(m)$ maximizes j's payoff, there is a profitable deviation. A contradiction to the assumption $m \in NE(\Gamma^M, \theta)$.

Claim 3 $g^M(m)$ maximizes *i*'s payoff, for some *i*.

Proof If $g^M(m)$ does not maximize either player's payoffs, then, by Claim 2, there is for both i = 1, 2 a state q_i such that $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$. Thus, by Claim 1, neither player's actions in the cycling phase affect the the other player's choices. Since infinite action streams are unconditional, i.e., player *i*'s actions are not conditioned on *j*'s actions, only Rules (A–C) of Γ^M apply, and they are identical with Rules (a–c) of Γ^* . Thus, $m \in NE(\Gamma^M, \theta)$ implies $g^M(m) \in g^*(NE(\Gamma^*, \theta))$. A contradiction to the initial hypothesis.

Claim 4 There is *i* and q_i such that $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$.

Proof Otherwise, by Claim 2, $g^M(m)$ maximizes *both* players' payoffs under θ . Thus, by the unanimity condition (i), $g^M(m) \in f(\theta)$. But this contradicts $f(\theta) = g^M(NE(\Gamma^*, \theta))$.

Claim 5 Rules D or E of Γ^M do not apply.

Proof It suffices to show $|X^1| = |X^2| = 1$. By Claim 4, there is q_i such that $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$. Since *i*'s cycling phase consists of q_i alone, we have $|X^i| = 1$. Suppose that $|X^j| > 1$. Then the cycling phase of *j* contains more than one element, say $\{q_j^1, \ldots, q_j^K\}$, K > 1. Thus for all q_j there is $s_i \in S_i$ such that $\tau_j(s_i, q_j) \neq q_j$. By Rule D, $a^i \in f(\theta^i)$ becomes implemented. But dropping all but one state in the cycling phase and choosing in the remaining state $s^j = (\theta^i, a^i, n^j, a^i)$ such that $n^j = 0$ if $n^i = 0$ and $n^j > n^i$ if $n^i > 0$ would not affect the implementation of a^i , by Rules A or C. Thus *j* would be better off by dropping the states. A contradiction.

Claim 6 Rule C of Γ^M does not apply.

Proof By Claim 3, $g^M(m)$ maximizes *i*'s payoff. By unanimity condition (i), $g^M(m)$ cannot maximize *j*'s payoff, $j \neq i$. By Claim 2, there is q_i in the cycling phase such that $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$. If rule C applies, then, since *j*'s actions are unconditional on *i*'s actions, and $\bar{n}^j > 0$, *j* can profitably deviate by choosing in all periods (n^i, b^i) such that $n^i > \bar{n}^j$ and b^i maximizes his payoff.

Claim 7 Rule B of Γ^M does not apply.

Proof Let *i* be defined as in Claim 4. If Rule B applies, then either (1) $\bar{n}^j > \bar{n}^i = 0$ or (2) $\bar{n}^i > \bar{n}^j = 0$. Since *i*'s actions are unconditional on *j*'s choices, *j* could deviate without consequences. Because of this, since in case (2) *j* does not move the game under the range of Rule C, $g^M(m)$ must maximize *j*'s payoff. By the unanimity condition (i), $g^M(m)$ cannot maximize *i*'s payoff, $j \neq i$. By Claim 2, there is q_j in *j*'s cycling phase such that $\tau_j(s_i, q_j) = q_j$, for all $s_i \in S_i$, and such that $n^j = 0$ for $(a^j, \theta^j, n^j, b^j) = \sigma_j(q_j)$. Since *j*'s actions are unconditional, *i* is free to choose either $b^i \in C_i(a^j, \theta^j)$ or $e(\theta^1, a^1, \theta^2, a^2) \in C_i(a^j, \theta^j)$. Thus $g^M(m)$ must maximize *i*'s payoff in $C_i(a^j, \theta^j)$. Since $g^M(m)$ already maximizes *j*'s payoff in *A*, by condition (ii), $g^M(m) \in f(\theta)$, a contradiction.

Hence (1) applies. Then any $b^j \in C_j(a^i, \theta^i)$ or $e(\theta^1, a^1, \theta^2, a^2) \in C_j(a^i, \theta^i)$ could become implemented. Since *i*'s actions are unconditional on *j*'s choices, $g^M(m)$ has to maximize *j*'s payoff in $C_j(a^i, \theta^i)$. By Claim 3, $g^M(m)$ also maximizes some player's payoff in *A*. If that player is *i*, then by condition (ii), $g^M(m) \in f(\theta)$, a contradiction. Thus $g^M(m)$ maximizes *j*'s payoff in *A*. To not violate condition (i), $g^M(m)$ cannot maximize *i*'s payoff in *A*. Hence, by Claim 2, there is q_j in *j*'s cycling phase such that $\tau_j(s_i, q_j) = q_j$, for all $s_i \in S_i$, and such that $n^j > 0$ for $(a^j, \theta^j, n^j, b^j) = \sigma_j(q_j)$. But since now *j*'s actions are unconditional on *i*'s choices, and *i* does not move the game under the range of Rule C, $g^M(m)$ must also maximize *i*'s payoff which contradicts condition (i).

Claim 8 Rule A of Γ^M does not apply.

Proof Suppose Rule A applies with (a, θ) . Then, since $a = g^M(m)$ maximizes *i*'s payoff by Claim 3, and cannot maximize *j*'s payoff since then it would violate condition (i), it must be that there is q_i in the cycling phase such that $\tau_i(s_j, q_i) = q_i$, for all $s_j \in S_j$, and such that $(a, \theta) = (a^i, \theta^i)$ for $(a^i, \theta^i, n^i, b^i) = \sigma_i(q_i)$. But then *j* could move the game under the range of Rule B, and induce any $b^j \in C_j(a, \theta)$. Thus $g^M(m)$ must maximize *j*'s payoff in $C_j(a, \theta)$.

Thus there is $b^j \in C_j(a^i, \theta^i)$ that improves j's payoff relative to a^i . Since i's strategy is unconditional, it becomes profitable for j to deviate by moving the game under the range of Rule B and choosing b^j .

Since one of Rules A,B,C,D, or E must apply, Claims 5, 6, 7, and 8 imply that $g^{M}(m) \in f(\theta)$.

References

- Abreu D, Rubinstein A (1988) The structure of Nash equilibrium in repeated games with finite automata. Econometrica 56:1256–1283
- Aumann R (1997) Rationality and bounded rationality. Games Econ Behav 21:2-14
- Aumann R, Hart S (2003) Long cheap talk. Econometrica 71:1619-1660
- Aumann R, Shapley L (1976) Solutions for continuously competitive situations. Unpublished manuscript
- Baliga S, Corchon L, Sjostrom T (1997) The theory of implementation when the planner is a player. Games Econ Behav 77:15–33
- Banks J, Sundaram R (1990) Repeated games, finite automata and complexity. Games Econ Behav 2:97– 117
- Chatterjee K, Sabourian H (2000) Multiperson bargaining and strategic complexity. Econometrica 68:1491– 1509
- Dutta B, Sen A (1991) A necessary and sufficient condition for two-person Nash imple-mentation. Rev Econ Stud 58:121–128
- Jackson M, Palfrey T (1998) Efficient and voluntary implementation in markets with repeted pairwise bargaining. Econometrica 66:1353–1389
- Jackson M, Palfrey T (2001) Voluntary implementation. J Econ Theory 98:1-25
- Kalai E, Ledyard J (1998) Repeated implementation. J Econ Theory 83:308-317
- Kalai E, Stanford W (1988) Finite rationality and interpersonal complexity in repeated games. Econometrica 56:397–410
- Ma A, Moore J, Turnbull S (1988) Stopping agents from cheating. J Econ Theory 46:355-376
- Maskin E (1999) Nash equilibrium and welfare optimality. Rev Econ Stud 66:23–38
- Maskin E, Moore J (1999) Implementation and renegotiation. Rev Econ Stud 66:39-57

Moore J, Repullo R (1990) Nash implementation: a full characterization. Econometrica 58:1083-1099

Osborne M, Rubinstein A (1995) A course in game theory. MIT Press, Cambridge

Rubinstein A (1986) Finite automata play the repeated prisoners dilemma. J Econ Theory 39:83-97

Tennenholz M (2004) Program equilibrium. Games Econ Behav 49:363–374

Vartiainen H (2003) Bargaining without commitment. Manuscript, Yrjö Jahnsson Foundation