Theoretical Economics 5 (2010), 51–71 1555-7561/20100051

Nash implementation with little communication

Ilya R. Segal
Department of Economics, Stanford University

The paper considers the communication complexity (measured in bits or real numbers) of Nash implementation of social choice rules. A key distinction is whether we restrict to the traditional one-stage mechanisms or allow multistage mechanisms. For one-stage mechanisms, the paper shows that for a large and important subclass of monotonic choice rules—called *intersection monotonic*—the total message space size needed for one-stage Nash implementation is essentially the same as that needed for "verification" (with honest agents who are privately informed about their preferences). According to Segal (2007), the latter is the size of the space of minimally informative budget equilibria verifying the choice rule. However, multistage mechanisms allow a drastic reduction in communication complexity. Namely, for an important subclass of intersection-monotonic choice rules (which includes rules based on coalitional blocking such as exact or approximate Pareto efficiency, stability, and envy-free allocations), we propose a two-stage Nash implementation mechanism in which at most 5 alternatives plus $4N \log_2 N$ bits are announced in any play. Such two-stage mechanisms bring about an exponential reduction in the communication complexity of Nash implementation for discrete communication measured in bits or a reduction from infinite- to low-dimensional continuous communication.

Keywords. Monotonic social choice rules, Nash implementation, communication complexity, verification, realization, budget sets, price equilibria, message space dimension.

JEL classification. D71, D82, D83.

1. Introduction

This paper considers the problem of Nash implementation of social choice rules, i.e., designing a mechanism whose set of Nash equilibria equals the set of socially desirable alternatives. As shown by Maskin (1999), any Nash implementable choice rule must satisfy the property of "monotonicity," which, together with the "no veto power" (NVP) property, also proves sufficient for Nash implementation with $N \geq 3$ agents. The sufficiency part is shown by constructing a "canonical" mechanism to implement the

Ilya R. Segal: ilya.segal@stanford.edu

I am grateful to Steven Williams, Jeffrey Ely (the co-editor), two anonymous referees, and participants of the 2007 NBER Conference on General Equilibrium for their comments, and to Kester Tong for excellent research assistance. I also gratefully acknowledge financial support from the NSF via Grant ITR 0427770 and from the Toulouse Network for Information Technology.

DOI: 10.3982/TE576

choice rule. The canonical mechanism has been criticized for its enormous communication burden: Indeed, it requires each agent to describe the preferences of all the agents (along with an integer), which is impractical in most settings. A number of papers have demonstrated that Nash implementation can be achieved with simpler mechanisms, even much simpler in some special settings.[1] However, these papers have not considered the problem of minimizing the communication cost of Nash implementation except in several special settings (such as Pareto, Walrasian, or Lindahl correspondences in classical economies with convex preferences).[2,3]

The present paper offers two contributions to this literature: (1) a mechanism for one-stage Nash implementation at a close to minimal communication cost and (2) a three-stage mechanism for Nash implementation with a drastically lower communication cost. The construction does not work for all implementable choice rules, but for a large class of them (which includes all specific monotonic choice rules examined previously). Our approach follows the program suggested by Williams (1986), who relates the message space needed for Nash implementation to that needed to verify the desirability of an alternative when agents know their preferences privately but can be trusted to report honestly. At first glance, the two problems appear quite different: under Nash implementation, selfish agents with symmetric information send messages, while under verification, honest agents with private information respond to a message announced by a hypothetical omniscient oracle. Yet, as observed in Reichelstein (1984) and Williams (1986), Nash implementation can be viewed a special case of verification, since each agent's acceptance of (lack of profitable deviations from) a candidate Nash equilibrium depends only on his own preferences. This observation implies that the communication cost of Nash implementation is bounded below by that of verification.[4]

This paper further exploits the relation between Nash implementation and verification, using concepts and results developed in Segal (2007). The latter paper focuses on a large and important subclass of monotonic choice rules, called *intersection monotonic*

[1] See Chakravorti (1991), Dubey (1982), Duggan (2003), Dutta et al. (1995), Hurwicz (1979), Li et al. (1995), McKelvey (1989), Peleg (1996), Reichelstein and Reiter (1988), Saijo (1988), Saijo et al. (1996, 1999), Schmeidler (1980), Sjostrom (1996), Suzuki (2009), Svensson (1991), Tian (1989, 1992, 1994), and Yoshihara (1999).

[2] Note that if the agents were honest, then under symmetric information we could simply ask one agent to report a socially desirable outcome, which would entail a low communication cost. Any additional communication cost of Nash implementation can thus be interpreted as the "communication cost of selfishness." Fadel and Segal (2009) examine the communication cost of selfishness for (partial) Bayesian–Nash and ex post implementation.

[3] The canonical mechanism has also been criticized for its use of integer or modulo games to eliminate undesirable equilibrium outcomes (see, e.g., Jackson 1992). While in a number of settings such tricks can be avoided, in this paper we do not examine this issue, because the communication cost of modulo games is fairly low.

[4] Williams (1986) does not attempt a reverse comparison of the communication costs of implementation and verification: while he "embeds" a verification protocol into an implementation mechanism under some conditions, he admits that "the strategy space in our construction is rather large, relative to the size of the message space [used for verification]. Clearly, if the goal is to devise games with small strategy spaces, then the embedding itself is a key step [...]. Within the context of economic theory, this issue has not yet been studied." Such an embedding with small strategy spaces is constructed in our Mechanism 1 below.

(IM), and shows that such rules are verified with minimal communication by announcing a "minimally informative verifying budget equilibrium." Such an equilibrium describes a proposed alternative and offers each agent a budget set—an appropriately restricted subset of alternatives. The fact that the proposed equilibrium is indeed an equilibrium in a given state—i.e., that each agent cannot improve upon the proposed alternative within his budget set—must verify the social desirability of the proposed alternative in this state. The budget sets must be chosen carefully: on the one hand, they must be large enough for the equilibrium to achieve verification; on the other hand, they must not be too large so that the equilibrium does not reveal more than necessary about the agents' preferences. Segal (2007) gives an algorithm for constructing such "minimally informative verifying budget equilibria" for any given IM choice rule.

To apply these ideas to Nash implementation, observe that a Nash equilibrium of a mechanism describes for each agent a "budget set" consisting of the alternatives he could achieve by unilateral deviations and that the described budget equilibrium must verify that the alternative is socially desirable. The only difference from the verification scenario is that an agent's budget set must be described by the other agents rather than by the hypothetical omniscient oracle. This observation leads us to construct a Nash implementation mechanism with small strategy spaces, in which exactly two agents—say, agents 1 and 2—describe a minimally informative budget equilibrium verifying a choice rule. In addition, each agent announces an alternative and an integer between 1 and N. When all the agents agree on an alternative, and agents 1 and 2 also agree on a budget equilibrium that supports it, the alternative is implemented. When one agent deviates from such unanimous agreement and proposes another alternative, his proposal is implemented if and only if it lies in his budget set as described by another agent. Thus, unanimous agreement is a Nash equilibrium in a given state if and only if agents 1 and 2 announce a budget equilibrium for this state, and since only budget equilibria verifying the choice rule can be announced, unanimous-agreement Nash equilibria yield desirable alternatives. To ensure that non-unanimous Nash equilibria do not yield any undesirable alternatives, we use the integers announced by the agents to induce a "modulo game" when more than one agent disagrees with others, and we make use of the NVP property, just as it is done in the canonical mechanism.

This construction yields our *Mechanism 1*, which implements any choice rule that is IM and NVP with $N \geq 3$ agents.[5] In this mechanism, two agents send a minimal message needed for verification (which takes the form of describing a minimally informative verifying budget equilibrium), and, in addition, each of the N agents describes an alternative and sends $\lceil \log_2 N \rceil$ bits.[6]

The proposed mechanism is particularly useful in conjunction with Segal's (2007) algorithm for constructing minimally informative verifying budget equilibria. When using such budget equilibria, Mechanism 1 gives us a "close-to-minimal" Nash implementation mechanism. In many important settings, this mechanism proves to have a much

[5] Mechanism 1 can be used both for the "weak" versions of the implementation and verification problems, in which it suffices to implement/verify a nonemtpy subset of desirable outcomes in any given state, and for the "full" version, in which *all* desirable alternatives must be implementable/verifiable.

[6] The notation $\lceil z \rceil$ denotes the smallest integer greater than or equal to z.

smaller strategy space than what full description of agents' preferences or even just their lower contour sets at a given alternative would require. For example, this is true for the problem of implementing interior Pareto efficient allocations in classical convex economies with private and public goods, in which the minimally informative verifying budget equilibria take the familiar form of Walrasian and Lindahl equilibria, respectively (Segal 2007).[7]

The second observation of the paper is that while the total size of strategy spaces describes the communication complexity of a one-stage mechanism, it may severely overstate the communication complexity of a multistage mechanism. This is because describing an agent's (complete contingent) strategy in a multistage game may take a lot longer than simply playing the game. In fact, we show that multistage mechanisms allow a huge reduction in the communication complexity of Nash implementation. We do this for a subclass of IM choice rules, called *coalitionally unblocked* (CU) rules (Segal 2007), which still includes all the specific monotonic choice rules that have been considered in economics (such as the Pareto rule, approximate Pareto, the core, stable matching, and envy-free rules). We implement such rules with a three-stage mechanism in which after an alternative is announced (by three agents, in Stage 1), it can be challenged by any agent proposing another alternative (in Stage 2). Then (in Stage 3) other agents (at most three of them) are asked to say which agents' budget sets allow the challenge. Thus, while the agents' complete contingent strategies in the mechanism describe all the budget sets, any single play of the mechanism describes only the placement of a single alternative into the budget sets.[8]

This construction yields our *Mechanism 2*, which implements any choice rule that is CU and NVP with $N \geq 3$ agents. The communication in any play of this mechanism (both in and out of equilibrium) is bounded by describing no more than five alternatives plus $4N\lceil \log_2 N \rceil$ bits.

To see the potential communication reduction allowed by Mechanism 2, recall from Segal (2007) that in some social choice problems, the minimally informative verifying budget equilibria must use all possible subsets of alternatives as budget sets. For example, to verify Pareto efficiency on the universal preference domain over X alternatives, any partition of the alternatives among the N agents must be used as a verifying budget equilibrium. Describing such a partition requires sending roughly $X \log_2 N$ bits (to

[7]Low-communication one-stage mechanisms for Nash implementation of general monotonic social choice rules also are proposed by McKelvey (1989, Section 5) and Hurwicz and Reiter (2006, Section 3.9.2). Without the benefit of Segal's (2007) budget-equilibrium characterization of verification of IM choice rules, these works could not tightly relate the communication requirements of Nash implementation to that of verification. McKelvey (1989) does not formalize the problem of minimizing communication, while Hurwicz and Reiter (2006) consider communication *in equilibrium*, rather than the total size of message spaces (which must include many message profiles that are never sent in equilibrium). Despite these differences, the two constructions are related to our Mechanism 1, as we will point out in greater detail below.

[8]This is not a complete description of the mechanism: we also need to make sure that agents' reports about budget sets are consistent with verifying the choice rule and that in any equilibrium, the strategies of different agents describe the same budget sets. We achieve both goals without raising the communication cost. We also use the modulo game to take care of equilibria that involve one or more challenges.

allocate each alternative into one of the agents' budget sets).[9] Any one-stage Nash implementation mechanism has to have at least as much communication. On the other hand, Mechanism 2 uses no more than $5\lceil\log_2 X\rceil + 4N \cdot \lceil\log_2 N\rceil$ bits, yielding an exponential reduction in communication complexity from one-stage mechanisms when the number of alternatives is large.

In a model with continuous communication, multistage mechanisms allow even more drastic reduction in communication. For example, consider the problem of allocating a divisible good among the agents, compensating them with unlimited transfers of the numeraire. Agents have utilities that are quasilinear in the numeraire and nondecreasing in their consumption of the good. The goal is to find a Pareto efficient (i.e., surplus-maximizing) allocation. As shown by Calsamiglia (1977), verifying this goal requires infinite-dimensional communication. In Nisan and Segal (2006), this result is derived from the observation that verifying efficiency requires describing a nonlinear personalized pricing function $[0,1] \to \mathbb{R}$ (for divisible good consumption in terms of the numeraire) for all agents but one, which is infinite dimensional (even if arbitrary smoothness is assumed). One-stage Nash implementation is at least as hard. In contrast, Mechanism 2 achieves Nash implementation while transmitting only $10(N-1)$ real numbers (and $4N \cdot \lceil\log_2 N\rceil$ bits).

We conclude with a philosophical discussion. While current mainstream economic thinking justifies price mechanisms by the need to provide incentives to selfish agents, we showed in Segal (2007) that supporting prices (more generally, budget sets, which could be described by personalized nonlinear prices) must be communicated to attain many important social goals even if agents are honest but their preference information is private and must be aggregated to find a socially desirable outcome. An intuition for this, based on Hayek (1945), is that to achieve social goals that are "sufficiently congruent" with private goals, communication is minimized by asking individuals to maximize their own preferences within certain "budget sets," which must be carefully outlined to coordinate their choices and attain the social goals.

The analysis of one-stage Nash implementation suggests another possible justification for prices: they must be used to create incentives even when information is symmetric. This justification is not valid, however, once multistage mechanisms are allowed. Multistage Nash mechanisms with symmetric information need not reveal supporting prices in any play, and so can have very low communication complexity. In contrast, multistage mechanisms with private information, even when agents are honest, cannot do any better than verification mechanisms (Kushilevitz and Nisan 1997, Chapter 2), and therefore according to Segal (2007) must still communicate supporting prices, which bounds below their communication complexity. In brief, once multistage mechanisms are allowed, price revelation becomes unnecessary when information is symmetric (even if agents are selfish), but is still necessary when preference information is private (even if agents are honest). Thus, we conclude that price revelation must arise due to the need to aggregate distributed preference information, rather than due to the need to create incentives for selfish agents.

[9] For other problems, the minimally informative budget equilibria may be described more succinctly. For example, for Pareto efficiency in smooth convex exchange economies, the budget sets are Walrasian and so can be described with linear anonymous prices (Segal 2007).

2. Setup

2.1 *The social choice problem*

Let X be the set of social alternatives and let N be a finite set of agents, numbered 1 to N. (With a slight abuse of notation, the same letter will denote a set and its cardinality when this causes no confusion.) Let $\mathcal{P}$ denote the set of all preference relations over set X. The set of agent i's possible preference relations is denoted by $\mathcal{R}_i \subset \mathcal{P}$. A *state* is a preference profile $R = (R_1, \ldots, R_N) \in \mathcal{R}_1 \times \cdots \times \mathcal{R}_N \equiv \mathcal{R}$, where $\mathcal{R}$ is the *state space*, also called *preference domain*. The goal is to realize a *choice rule*, which is a correspondence $F : \mathcal{R} \twoheadrightarrow X$. For every state $R \in \mathcal{R}$, the rule specifies the set $F(R)$ of "desirable" alternatives in this state.

The following two properties of choice rules, introduced in Maskin (1999), play a prominent role in Nash implementation. (We use the standard notation $L(x, R_i) = \{y \in X : x \, R_i \, y\}$—the *lower contour set* of agent i's preference relation R_i at alternative $x \in X$.)

Definition 1. Choice rule F is *monotonic* if $\forall R \in \mathcal{R}$, $\forall x \in F(R)$, and $\forall R' \in \mathcal{R}$ such that $L(x, R_i) \subset L(x, R'_i)$ $\forall i \in N$, we have $x \in F(R')$.

Definition 2. Choice rule F has *no veto power* (NVP) if $\forall i \in N$, $\forall R \in \mathcal{R}$, and $\forall x \in X$ such that $L(x, R_j) = X$ $\forall j \in N \setminus \{i\}$, we have $x \in F(R)$.

The next two properties are introduced in Segal (2007):

Definition 3. Choice rule F is *intersection-monotonic* (IM) if $\forall \widetilde{\mathcal{R}} = \widetilde{\mathcal{R}}_1 \times \cdots \times \widetilde{\mathcal{R}}_N \subset \mathcal{R}$, $\forall x \in \bigcap_{R \in \widetilde{\mathcal{R}}} F(R)$, and $\forall R' \in \mathcal{R}$ such that $\bigcap_{R_i \in \widetilde{\mathcal{R}}_i} L(x, R_i) \subset L(x, R'_i)$ $\forall i \in N$, we have $x \in F(R')$.

Note that this property implies monotonicity by taking $\widetilde{\mathcal{R}}$ to be a singleton. In addition to monotonicity, it requires, in particular, that if the desirability of alternative x is preserved by making an agent strictly prefer alternative y or alternative z to x (holding the other preferences fixed), then it should also be preserved by making the agent strictly prefer *both* y and z to x (assuming all the relevant preference profiles are feasible). (In fact, IM is characterized by this requirement and monotonicity when $\mathcal{R} = \mathcal{P}^N$ and X is finite.)

Definition 4. Choice rule F is a *coalitionally unblocked* (CU) choice rule if there exists a *blocking correspondence* $\beta : X \times 2^N \twoheadrightarrow X$ for which

$$F(R) = \left\{ x \in X : \beta(x, S) \subset \bigcup_{i \in S} L(x, R_i) \; \forall S \subset N \right\} \quad \forall R \in \mathcal{R}.$$

We can interpret $\beta(x, S)$ as the set of alternatives that coalition $S \subset N$ is allowed to use to block candidate alternative $x \in X$. The choice rule consists of candidate alternatives for which no coalition can find a strictly Pareto improving blocking. Note that

CU implies, in particular, that if the desirability of alternative x is preserved by either making agent i strictly prefer either alternative y over x or making agent j strictly prefer alternative $z \neq y$ over x (holding the other preferences fixed), then it should also be preserved by implementing both of these preference reversals at once (assuming all the relevant preference profiles are feasible). (In fact, CU is characterized by this requirement and monotonicity when $\mathcal{R} = \mathcal{P}^N$ and X is finite.) See Segal (2007) for further analysis, which formally establishes that every CU rule is IM, every IM rule is monotonic, and both inclusions are strict. The class of CU rules is still large enough to include all specific monotonic rules that have been considered, such as exact or approximate Pareto efficiency, the core, stable matchings, or envy-free rules.

2.2 *Nash implementation*

In the Nash implementation problem, all agents observe the state of the world and play a Nash equilibrium of the mechanism offered to them.

DEFINITION 5. A *mechanism* ("game form") $G = \langle M_1, \ldots, M_N, h \rangle$ describes a strategy space M_i for each agent and an outcome function $h : \prod_i M_i \to X$. The Nash equilibrium correspondence of the mechanism is given by

$$\nu_G(R) = \{m \in M : g(m)\, R_i\, g(m_i', m_{-i})\ \forall i \in N\ \forall m_i' \in M_i\}.$$

Mechanism G *fully implements* choice rule F if $h(\nu_G(R)) = F(R)\ \forall R \in \mathcal{R}$; G *weakly implements* F if $\varnothing \neq h(\nu_G(R)) \subset F(R)\ \forall R \in \mathcal{R}$.[10]

Note that we can also allow multistage mechanisms, whose normal form can still be described by mechanism G above. Since Nash equilibrium is defined on the normal form, allowing multistage mechanisms does not affect the implementability of choice rules. The usefulness to us of multistage mechanisms stems from their substantially lower communication costs than their normal-form representations.

2.3 *Verification*

Now we consider the communication problem in which each agent i observes only his own "type"—in our case, preference relation R_i—but can be prescribed to follow an arbitrary strategy, rather than being selfish. Furthermore, we focus on a special kind of communication, called *verification* (or *nondeterministic communication* in computer science). In the verification problem, an omniscient oracle knows the true state R and, consequently, knows the desirable alternatives. However, he needs to prove to an ignorant outsider that alternative $x \in F(R)$ is indeed desirable. He does this by publicly announcing a message $m \in M$. Each agent i either accepts or rejects the message, doing so on the basis of his own type R_i. The acceptance of message m by all agents must prove to the outsider that alternative x is desirable.

Formally, verification is defined as follows.

[10] We use the standard notation for the image of a set: $h(A) = \bigcup_{m \in A} h(m)$ (Aliprantis and Border 1999, p. 3).

DEFINITION 6. A *verification protocol* is a triple $\Gamma = \langle M, \mu, h \rangle$, where

- M is the message space
- $\mu : \mathcal{R} \twoheadrightarrow M$ is the message correspondence, which must satisfy *privacy preservation*,

$$\mu(R) = \bigcap_{i \in N} \mu_i(R_i) \quad \forall R \in \mathcal{R}, \text{ where } \mu_i : \mathcal{R}_i \twoheadrightarrow M \ \forall i \in N$$

- $h : M \to X$ is the outcome function.

Verification protocol Γ *fully verifies* choice rule F if $h(\mu(R)) = F(R)\ \forall R \in \mathcal{R}$; Γ *weakly verifies* F if $\varnothing \neq h(\mu(R)) \subset F(R)\ \forall R \in \mathcal{R}$.

While the verification scenario is patently unrealistic, it still proves quite useful. The key reason to consider verification is that it offers a lower bound on the communication requirements of any multistage communication protocol (Kushilevitz and Nisan 1997, Chapter 2). Formally, a multistage communication protocol consists of (i) an extensive-form game in which agents' moves are their messages, (ii) agents' strategies in the game that are contingent on their types as well as observed histories, and (iii) a function assigning alternatives to the game's terminal nodes. Observe that any multistage communication protocol can be represented as a weak verification protocol by letting all the messages be sent by the oracle instead of the agents, and by having each agent accept the message sequence if and only if all the messages the oracle sent in his stead are consistent with his strategy given his type. The oracle's message space M thus consists of the protocol's possible message sequences (terminal nodes). Therefore, the communication cost of weak verification bounds below the communication cost of computing an alternative in the choice rule. The lower bound is tight in some cases but weak in some other cases, where communication requires a lot more than verification.

In addition, note that verification gives a lower bound on the one-stage Nash implementation problem, since any Nash implementation protocol $G = \langle M_1, \ldots, M_N, h \rangle$ can be viewed as a verification protocol $\Gamma = \langle M_1 \times \cdots \times M_N, \nu_G, h \rangle$. Indeed, note that the Nash equilibrium correspondence ν_G by construction satisfies privacy preservation: $\nu_G(R) = \bigcap_i \nu_G^i(R_i)$, where

$$\nu_G^i(R_i) = \{m \in M : g(m)\ R_i\ g(m_i', m_{-i})\ \forall m_i' \in M_i\}$$

is the best-response correspondence of agent i, which depends only on this agent's preferences R_i. Thus, the oracle can announce a candidate Nash equilibrium strategy profile, and each agent accepts the announcement if and only if he cannot find a profitable unilateral deviation from this profile.

We show in this paper that the relation between verification and Nash implementation is quite tight (unlike the relation between verification and communication). Intuitively, this is because in the implementation problem, each agent has full information and so can send the oracle's message by himself as long as he does not have an incentive to misrepresent it.

2.4 *Measures of communication cost*

In the case of discrete communication, the communication cost is naturally defined as "communication complexity," which is the (worst-case) number of bits needed to encode the messages (Kushilevitz and Nisan 1997).[11] In the case of verification, the oracle needs $\lceil \log_2 M \rceil$ bits to encode his message from M. The minimal communication complexity of a verification protocol offers a lower bound on the communication complexity without an oracle (Kushilevitz and Nisan 1997).

For continuous communication, the communication cost can be defined naturally as the total dimension of the messages sent. However, for a meaningful concept of dimension, we need to rule out "smuggling" a multidimensional message in a single dimension with a 1-to-1 function such as the inverse Peano function. The economic literature on message space dimension suggests various topological concepts of dimension that prevent such smuggling. In particular, Segal (2007) defines the topological dimension of the message space using a topology on messages defined based on their "meaning," i.e., the set of states in which they are sent. For the sake of brevity, we do not repeat the definitions of Segal (2007) in this paper; instead we use the concept of dimension in its intuitive sense.

These concepts of communication cost can be applied naturally to Nash implementation mechanisms. First consider one-stage mechanisms. In a discrete one-stage mechanism, each agent i needs $\lceil \log_2 M_i \rceil$ bits to encode his strategy from M_i, so the total communication complexity of the game can be defined as $\sum_i \lceil \log_2 M_i \rceil$. For a continuous one-stage mechanism Γ, we can define the total dimension of the mechanism as the sum of the dimensions of the individual agents' message spaces.[12]

Now consider multistage mechanisms and note that the communication cost of a mechanism whose normal form is G could be drastically lower than that of the one-stage mechanism G. For example, consider the discrete case in which the agents' moves in the multistage extensive-form game can be represented as announcing bits. Suppose that the maximum number of bits sent in the game is d. This game can have up to $2^d - 1$ decision nodes, and to describe agents' contingent strategies in it requires 1 bit per decision node and so up to $2^d - 1$ bits in total. Thus, the communication complexity of describing strategies in a multistage game can be exponentially higher than that of playing the game. In a continuous mechanism, the increase can be even more drastic: even a very simple multistage mechanism can have an infinite-dimensional strategy space. For example, consider the two-stage mechanism in which first agent 1 announces $x_1 \in [0, 1]$

[11] Using bits is merely a normalization, because an elementary message in *any* finite alphabet could be coded with a fixed number of bits. What *is* important for the definition is that the coding and the communication protocol can be selected optimally for the problem at hand: if instead agents could only communicate using messages with preexisting meanings, this might raise the communication cost substantially.

[12] Hurwicz and Reiter (2006, Section 3.9.2) instead consider the "in-equilibrium" communication cost, which only counts message profiles that can ever arise in equilibrium. Since the message profiles that never arise in equilibrium must still be allowed to eliminate undesirable Nash equilibria, we do count such messages profiles, following the spirit of worst-case communication cost measures (even if the worst case has probability zero of arising in equilibrium). As we will point out, the in-equilibrium communication cost of mechanisms proposed in this paper will be even lower than the worst-case cost that we define.

and then agent 2 announces $x_2 \in [0,1]$. Agent 2's strategy in this mechanism is an arbitrary function $[0,1] \to [0,1]$ and so it is infinite-dimensional. In contrast, only two numbers are communicated in any play of the mechanism. While these examples are abstract, we will construct examples of similar reduction in the communication complexity of Nash implementation.

2.5 *Role of budget equilibria*

A famous economic example of verification is Walrasian equilibrium. The role of the oracle is played by the "Walrasian auctioneer," who announces the equilibrium prices and allocations. Each agent accepts the announcement if and only if his announced allocation constitutes his optimal choice from the budget set given by the announced prices. This concept can be generalized to that of a "budget equilibrium," in which the oracle's message consists of a proposed alternative $x \in X$ and a *budget set* $B_i \subset X$ for each agent i. Each agent $i \in N$ accepts message $(B_1, \dots, B_N, x)$ if and only if there is no alternative in his budget set B_i that he strictly prefers to the proposed alternative x. Message $(B_1, \dots, B_N, x)$ is a *budget equilibrium in state* $R \in \mathcal{R}$ if it is accepted by all agents in this state. Formally, the budget equilibrium correspondence $E : \mathcal{R} \twoheadrightarrow 2^{XN} \times X$ is described as

$$E(R) = \{(B,x) \in 2^{XN} \times X : B_i \subset L(x, R_i)\ \forall i \in N\}.$$

The correspondence E satisfies privacy preservation because each agent's acceptance depends only on his own preferences.

The oracle's message space M in a budget protocol is a collection of budget equilibria that he is allowed to announce, and the outcome function simply implements the proposed alternative.

DEFINITION 7. Protocol $\langle M, \mu, h\rangle$ is a *budget protocol* if $M \subset 2^{XN} \times X$, $\mu(R) = E(R) \cap M$ $\forall R \in \mathcal{R}$, and $h(B,x) = x\ \forall (B,x) \in M$.

Clearly, the space M of budget equilibria used is important for whether the protocol verifies F. In particular, for the protocol to verify F (either fully or weakly), it must use only budget equilibria of the following kind.

DEFINITION 8. The message $(B,x) \in 2^{XN} \times X$ is a *budget equilibrium verifying* F if $\mu^{-1}(B_1, \dots, B_N, x) \subset F^{-1}(x)$.

However, the message space need not include *all* the budget equilibria verifying F. In fact, it turns out that for IM choice rules, the size of the message space can be reduced while restricting attention to the following budget equilibria.

DEFINITION 9. Suppose IM choice rule F on preference domain $\mathcal{R}$ extends to an IM choice rule on the universal preference domain $\mathcal{P}^N$. Then $(B,x) \in 2^{XN} \times X$ is a *minimally informative budget equilibrium verifying* F if for some $R \in \mathcal{P}^N$,

$$B_i = L(x, R_i) = \bigcap_{R_i' \in \mathcal{R}_i : x \in F(R_i', R_{-i})} L(x, R_i') \quad \forall i \in N. \tag{1}$$

In Segal (2007), this concept is not postulated, but is derived by constructing messages that verify F while revealing minimal information about the state of the world. It is shown that when F is IM on $\mathcal{P}^N$, these messages can be characterized as budget equilibrium messages of the form (1).[13] Furthermore, Segal (2007) offers an algorithm for constructing these minimally informative budget equilibria for a given social choice problem. Application of this algorithm yields such familiar budget equilibria as Walrasian and Lindahl equilibria for the problem of verifying interior Pareto efficient allocations in smooth convex economies with private and public goods, respectively; price equilibria with nonlinear personalized prices for efficient allocation problems with general nonconvex utilities or in combinatorial auctions; or "prematchings" for stable-matching problems, in which an agent's budget set is a subset of his potential matching partners.

Letting $\mathcal{E}_F$ be the space of all minimally informative budget equilibria verifying F, the following proposition follows from Segal (2007).

PROPOSITION 1. *The minimal message space size for fully or weakly verifying an IM choice rule F is achieved with a budget equilibrium protocol whose message space is a subset of $\mathcal{E}_F$.*

3. ONE-STAGE MECHANISMS

Recall that a Nash implementation protocol can be viewed as a verification protocol. (Furthermore, it can be viewed as a budget protocol, with message space $\{B_1(m), \ldots, B_N(m), g(m)\}_{m\in M_1\times\cdots\times M_N}$, where $B_i(m) = \{g(m_i', m_{-i}) : m_i' \in M_i\}$.) Thus, we have the following lemma.

LEMMA 1. *The minimal total size of strategy spaces required for full/weak Nash implementation is at least as high as the minimal size of message space for full/weak verification.*

We also provide an upper bound for the communication cost of Nash implementation relative to that of verification by starting with a budget equilibrium protocol with message space $\mathcal{E} \subset 2^{XN} \times X$ and constructing a mechanism in which two agents announce a budget equilibrium from $\mathcal{E}$, so that each agent's budget set is described by another agent.

MECHANISM 1. The strategy spaces are $M_1 = M_2 = \mathcal{E} \times X \times N$ and $M_i = X \times N$ for $i \geq 3$. When the messages are $m_i = (E_i, y_i, l_i) \in M_i$, for $i = 1, 2$ and $m_i = (y_i, l_i) \in M_i$ for $i \geq 3$, the outcome function h is specified as follows.

(a) If $\exists E = (B_1, \ldots, B_N, x) \in \mathcal{E}$ such that $(y_i, l_i) = (x, 1)\ \forall i \in N$ and $E_1 = E_2 = E$, then $h(m) = x$.

[13] In general, a preference profile $R \in \mathcal{P}^N$ that satisfies (1) need not be a feasible state in $\mathcal{R}$. When it is, then it is an "F-minimal state" as defined by McKelvey (1989). McKelvey's definition applies for general monotonic choice rules. However, if F is not IM, the F-minimal states are not characterized by (1) and do not generate minimally informative messages verifying F (indeed, such messages are no longer equivalent to announcing a supporting budget equilibrium).

(b) If not case (a) but $\exists i \in N$ and $\exists E = (B_1, \ldots, B_N, x) \in \mathcal{E}$ such that $E_j = E$ $\forall j \in \{1,2\} \setminus \{i\}$ and $(y_j, l_j) = (x, 1)$ $\forall j \in N \setminus \{i\}$, then

(i) if $y_i \in B_i$, then $h(m) = y_i$

(ii) if $y_i \notin B_i$, then $h(m) = x$.

(c) Otherwise $h(m) = y_i$ for $i = (\sum_j l_j \bmod N) + 1$.

PROPOSITION 2. *Suppose that choice rule F is NVP and $N \geq 3$. If the budget equilibrium protocol with equilibrium space $\mathcal{E}$ weakly/fully verifies F, then Mechanism 1 weakly/fully implements F.*

PROOF. We begin by demonstrating some Nash equilibria (NE) of the mechanism.

CLAIM 1. *If $E = (B_1, \ldots, B_N, x) \in \mathcal{E}$ is a budget equilibrium in state $R \in \mathcal{R}$, then the message profile with $E_1 = E_2 = E$ and $(y_i, l_i) = (x, 1)$ is a NE of the mechanism.*

PROOF. The outcome for this message profile is x. An agent i can unilaterally affect the outcome only by deviating to case (b)(i) and implementing an alternative $y_i \in B_i$, but such a deviation would not be profitable since $(B_1, \ldots, B_N, x)$ is a budget equilibrium in state R. ◁

In the case of full/weak verification, in any state $R \in \mathcal{R}$, for any/some $x \in F(R)$ there exist budget sets $B_1, \ldots, B_N \subset X$ such that $(B_1, \ldots, B_N, x) \in \mathcal{E}$ is a budget equilibrium in state R; hence by Claim 1, any/some $x \in F(R)$ can arise in a NE of the mechanism.

It remains to show that any NE outcome of the mechanism in state R is in $F(R)$, which is done in the following two claims:

CLAIM 2. *Any case (a) NE outcome x in a state $R \in \mathcal{R}$ is in $F(R)$.*

PROOF. Since each agent i can unilaterally deviate to case (b)(i) to implement any alternative $y_i \in B_i$, for this to be a NE, E must be a budget equilibrium in state R. Now, by the (full or weak) verification assumption, E verifies F, hence $x \in F(R)$. ◁

CLAIM 3. *Any case (b) or case (c) NE outcome x in a state $R \in \mathcal{R}$ is in $F(R)$.*

PROOF. From any case (b) or case (c) message profile, each agent i except possibly one can deviate to attain any alternative $y_i \in X$ in case (c) by choosing l_i, hence for this message profile to be a NE we must have $L(x, R_i) = X$. By NVP, this implies that $x \in F(R)$. ◁

This completes the proof of the proposition. □

By Proposition 1, for an IM choice rule F we can choose the space $\mathcal{E}$ to be a minimal subspace of minimally informative budget equilibria needed to weakly/fully verify F. Thus, using Proposition 2 and examining the size of strategy spaces in Mechanism 1 yields the following result.

COROLLARY 1. *Suppose choice rule F is IM and NVP, and $N \geq 3$. Then for discrete communication, using Mechanism 1 we can fully/weakly Nash implement F with a one-stage mechanism whose communication complexity is twice the communication complexity of full/weak verification plus $N \cdot (\lceil \log_2 X \rceil + \lceil \log_2 N \rceil)$ bits. For continuous communication, we can fully/weakly Nash implement F with a one-stage mechanism in which the total dimensionality of the agents' strategy spaces is twice the minimal message space dimension needed for full/weak verification plus N times the dimension of X.*

The corollary gives an upper bound on the communication cost of one-stage Nash implementation that is pretty close to the verification lower bound of Lemma 1. The upper bound can be tightened a bit more in some typical cases. For example, instead of asking each agent to announce the whole social outcome (which might be costly if the number of agents is large), it suffices to have each agent announce only the part of the outcome that his preferences are concerned with (e.g., his own consumption of goods). Also, the duplication of budget set descriptions can be improved upon in some settings, as long as we still ensure that each agent's budget set is described by the other agents' reports. For example, Reichelstein and Reiter (1988) show that for Nash implementation of Walrasian allocations in classical convex exchange economies with L goods, the additional cost relative to verification is roughly $L/(N-1)$ real numbers, which is enough to ensure price taking by each agent, while duplicate announcement of a Walrasian price vector as required by Mechanism 1 would require $L-1$ numbers. However, such additional improvements over Mechanism 1 appear to be possible only in special settings, so we do not pursue them here.

We could alternatively consider the communication cost of only "in-equilibrium" communication, as proposed by Hurwicz and Reiter (2006, Section 3.9.2). For simplicity, restrict attention to environments in which there does not exist an alternative in X that is simultaneously optimal for $N-1$ agents, and so the NVP property holds vacuously. (This includes any environment with a private good that all agents desire.) In such environments, Mechanism 1 has only case (a) Nash equilibria, the description of which is the same as describing a budget equilibrium from $\mathcal{E}$. Since the verification lower bound of Lemma 1 also applies to in-equilibrium communication (the oracle can replicate the mechanism using only the message profiles that are potential Nash equilibria), we see that the in-equilibrium communication cost of one-stage Nash implementation of such choice rules exactly equals the communication cost of verification.[14] In contrast, the minimal total size of the agents' message spaces required for Nash implementation may strictly exceed the verification lower bound, as demonstrated by Reichelstein and Reiter (1988).

4. MULTISTAGE MECHANISMS

Considering multistage games allows substantial savings in communication. The idea is that while the agents' complete contingent strategies in the extensive-form mecha-

[14] Note that this conclusion does not hold for monotonic choice rules that are not IM: the communication cost of verification of such rules may be minimized using messages that do not correspond to describing supporting budget sets, whereas Nash implementation is always a budget equilibrium protocol.

nism must still describe supporting budget sets (as is the case in any verification mechanism), these strategies need not be revealed in any single play of the mechanism. Thus, the communication cost of a multistage mechanism can be substantially lower than the cost of its normal-form representation, i.e., the cost of describing the agent's complete contingent strategies in the mechanism.[15]

Applying this idea to the revelation of budget sets, we see that a multistage mechanism need not reveal the budget sets in any single play. Instead, it suffices that whenever a candidate equilibrium alternative x is described and then challenged by a single agent proposing another alternative x', other agents are asked to "approve" the challenge (thus confirming that x' is in the challenger's budget set) or "disapprove" it (protest that x' is outside his budget set). Then, while the complete approval strategies contingent on all possible challenges x' describe all the budget sets, these strategies and the corresponding budget sets are not revealed in any single play of the mechanism (either in or out of equilibrium).

The tricky part of the construction is restricting the agents to approve sufficiently many challenges so that the corresponding budget sets are large enough to verify the choice rule. (Recall that if the budget sets are too small, the budget equilibrium would not verify the choice rule and so it would not be implemented by the mechanism. For an extreme example, if agents' strategies disapprove any challenge of a candidate equilibrium outcome x, this yields budget sets $B_i = \{x\}$ for all i, and x is sustained in equilibrium in any state of the world, regardless of whether it is socially desirable.) This is not straightforward to ensure because the entire approval strategies and the corresponding budget sets are not revealed in any play of the mechanism (and so, in general, we cannot verify that the budget sets described by the strategies satisfy characterization (1)). We are able to accomplish this for the class of CU choice rules defined in Section 2, for which the following observation holds:

LEMMA 2. *The CU choice rule given by a blocking rule β is fully verified with the budget protocol whose message space consists of budget equilibria $(B_1, \dots, B_N, x) \in 2^{NX} \times X$ that satisfy*[16]

$$\beta(x, T) \subset \bigcup_{i \in T} B_i \quad \text{for all } T \subset N. \tag{2}$$

PROOF. If $(B_1, \dots, B_N, x)$ is a budget equilibrium in state $R \in \mathcal{R}$ and satisfies (2), we must have $\beta(x, T) \subset \bigcup_{i \in T} B_i \subset \bigcup_{i \in T} L(x, R_i)$ for all $T \subset N$ and, therefore, $x \in F(R)$. Hence, any budget equilibrium that satisfies (2) verifies F. Furthermore, in any state

[15]Multistage implementation mechanisms were previously considered by Moore and Repullo (1988), who used the subgame-perfection refinement to implement choice rules that are not Nash implementable. Our goal is quite different since we still consider Nash implementation and use multistage mechanisms to reduce the communication complexity of implementing those choice rules that *are* Nash implementable. Note also that our Mechanism 2 has imperfect information and no proper subgames, and so subgame perfection has no bite in it.

[16]Budget equilibria of the form (2) are typically not minimally informative budget equilibria verifying F (and so do not satisfy (1)). We are not concerned about this as we can still use them to construct a low-communication multistage mechanism.

$R \in \mathcal{R}$ for any $x \in F(R)$, $(L(x,R_1),\dots,L(x,R_N),x)$ is a budget equilibrium that satisfies (2). □

The advantage of condition (2) is that it can be checked one alternative $x' \in X$ at a time—namely, by checking that each alternative x' belongs to the budget sets of "sufficiently many" agents so that each coalition $T \subset N$ that satisfies $x' \in \beta(x,T)$ contains an agent whose budget set contains x'. We can impose this restriction on an agent's approval strategy by showing him only the proposed challenge x' but not the identity of the challenger and by restricting him to list "sufficiently many" agents for whom this challenge should be approved.

EXAMPLE 1. The weak Pareto efficient choice rule is described by the blocking rule $\beta(x,S) = X$ if $S = N$ and $= \varnothing$ otherwise. For this choice rule, condition (2) takes the form $X \subset \bigcup_i B_i$, which can be checked by verifying that each alternative in X belongs to some agent's budget set. ◊

Just as in the previous section, we do not need the budget sets to be described by *all* the agents, as long as each agent's budget set is described by other agents. Here it will suffice to have just three agents describe all the budget sets, i.e., to approve deviations. We need to ensure that in equilibrium the three agents describe exactly the same budget sets, i.e., use the same approval strategies. For this purpose, we reward an agent who challenges with an alternative x' on whose approval other agents disagree by letting him implement any alternative. Similarly, it suffices for just three agents to describe a candidate equilibrium alternative, as long as any agent is allowed to challenge it. Finally, just as in the previous section, we use the modulo game to make sure that potential equilibria that involve a challenge do not yield undesirable outcomes.

Formally, consider the following three-stage mechanism, which has imperfect information: the only information the agents observe about each other's previous messages are the public "revelations" by the mechanism.[17]

MECHANISM 2.

Stage 1. Each agent $i \in \{1,2,3\}$ announces $x_i \in X$ and $l_i \in N$. Each agent $i \geq 4$ announces $l_i \in N$.

(a) If $l_i = 1\ \forall i \in N$ and $\exists x \in X$ such that $x_i = x\ \forall i \in \{1,2,3\}$, implement x.

(b) If not case (a) but $\exists i \in N$ ("the challenger") and $\exists x \in X$ such that $l_j = 1\ \forall j \in N \setminus \{i\}$ and $x_j = x\ \forall j \in \{1,2,3\} \setminus \{i\}$, then reveal "b," reveal x, and continue.

Stage 2b. Agent i (the challenger) announces $x', y \in X$. Only x' is revealed.

[17]The mechanism can be converted into a two-stage mechanism by letting the agents in Stage 1 also announce their contingent strategies in Stage 2, which would require describing $2N+3$ alternatives instead of just five. This conversion would preserve the huge communication savings exemplified below, which are due to economizing on budget set descriptions using Stage 3 strategies rather than on descriptions of alternatives.

Stage 3b. Each agent $j \in \{1,2,3\}/\{i\}$ announces $S_j \in \Sigma(x,x')$, where

$$\Sigma(x,x') = \{S \subset N : \forall T \subset N,\ x' \in \beta(x,T) \Rightarrow S \cap T \neq \varnothing\}.$$

(i) If $S_j = S$ for all $j \in \{1,2,3\}/\{i\}$ and $i \in S$, implement x'.

(ii) If $S_j = S$ for all $j \in \{1,2,3\}/\{i\}$ and $i \notin S$, implement x.

(iii) Otherwise implement y.

(c) Otherwise reveal "c" and continue.

Stage 2c. Agent $i = (\sum_j l_j) \bmod N + 1$ announces $y \in X$, which is implemented.

PROPOSITION 3. *If F is a CU choice rule described by the blocking rule β, F satisfies NVP, and $N \geq 3$, then Mechanism 2 fully Nash implements F.*

PROOF. We start by describing the agents' complete contingent strategies in the mechanism. For simplicity, we restrict attention to strategies that do not condition on the agent's own earlier actions. (We can do it because any strategy with such conditioning is equivalent in the normal-form representation of the game to one without it.)

The strategy of each agent $i \geq 4$ can then be described as $\langle l_i, x_i', y_i^b, y_i^c \rangle$, where $l_i \in N$ is the integer he announces in Stage 1, $x_i', y_i^b \in X$ are the alternatives x', y he announces in Stage 2b when he is the challenger, and $y_i^c \in X$ is the alternative he announces in Stage 2c.

As for an agent $i \in \{1,2,3\}$, his strategy in addition describes his Stage 1 announcement $x_i \in X$ and also a function $\sigma_i : X \to 2^N$ satisfying $\sigma_i(x') \in \Sigma(x_i,x')\ \forall x' \in X$ that gives his announcement $S_i = \sigma_i(x')$ in Stage 3b when someone else challenged x_i with an alternative x' in Stage 1. We interpret $\sigma_i(x')$ as the set of agents whose challenges x' are approved by agent i. The function σ_i can be equivalently described by defining for each agent j the budget set $B_j^i = \{x' \in X : j \in \sigma_i(x')\}$—the set of agent j's challenges that are approved by agent i. The restriction $\sigma_i(x') \in \Sigma(x_i,x')$ for all $x' \in X$ is then equivalent to requiring that $(B_1^i, \ldots, B_N^i, x^i)$ satisfy (2). Thus, we describe a feasible strategy of agent $i \in \{1,2,3\}$ as $\langle l_i, x_i, B_1^i, \ldots, B_N^i, x_i', y_i^b, y_i^c \rangle$ satisfying (2) (from which we can deduce for each x', $\sigma_i(x') = \{j \in N : x' \in B_j^i\} \in \Sigma(x_i,x')$).

Now the result is proved with the following three claims.

CLAIM 1. *If $x \in F(R)$ in state $R \in \mathcal{R}$, then x is a case (a) NE outcome in state R.*

PROOF. Consider the strategy profile given by

$$\langle l_i, x_i, B_1^i, \ldots, B_N^i, x_i', y_i^b, y_i^c \rangle = \langle 1, x, L(x,R_1), \ldots, L(x,R_N), x, x, x \rangle \text{ for all } i \in \{1,2,3\}$$
$$\langle l_i, x_i', y_i^b, y_i^c \rangle = \langle 1, x, x, x \rangle \text{ for all } i \geq 4.$$

The strategies of agents $i \in \{1,2,3\}$ are feasible because the described budget sets satisfy (2) due to the fact that $x \in F(R)$. These strategies result in case (a) and yield outcome x. To see that these strategies form a NE, note that each agent $i \in N$ can unilaterally change the outcome only by challenging it and going to case (b)(i), in which he can only attain an outcome $x' \in L(x, R_i)$. ◁

Claim 2. *Each case (a) NE in state $R \in \mathcal{R}$ yields an outcome $x \in F(R)$.*

Proof. (i) If $B_i^1 = B_i^2 = B_i^3 = B_i$ for each $i \in N$, then each agent $i \in N$ can deviate to implement any alternative $x' \in B_i$ by announcing $l_i > 1$ and inducing case (b)(i). Thus, $(B_1, \ldots, B_N, x)$ must be a budget equilibrium in state R, and since it satisfies (2), $x \in F(R)$. (ii) If, on the contrary, there exist $x' \in X$, an agent $i \in N$, and agents $j, k \in \{1,2,3\}$ such that $x' \in B_i^j \setminus B_i^k$, then each agent r except possibly one (if the other agent in $\{1,2,3\}$ describes the same budget sets as j or k) can deviate to attain any alternative $y_r^b \in X$ by announcing $l_r > 1$ and $x_r' = x'$, inducing case (b)(iii). Hence, to have a NE, we must have $L(x, R_r) = X$, thus, by NVP, $x \in F(R)$. ◁

Claim 3. *Each case (b) or case (c) NE in state $R \in \mathcal{R}$ yields an outcome $x \in F(R)$.*

Proof. Each agent i except one possibly one (the challenger in case (b)) can deviate to attain any alternative $y_i^c \in X$ in case (c) by choosing l_i; hence, to have a NE, we must have $L(x, R_i) = X$, and thus by NVP, $x \in F(R)$. ◁

This completes the proof of the proposition. □

Observe that in any play of Mechanism 2, agents describe at most five alternatives and, in addition, send no more than $N \cdot \lceil \log_2 N \rceil + 3N \leq 4N \cdot \lceil \log_2 N \rceil$ bits (the longest communication takes place in case (b)). This offers a potentially huge reduction in communication relative to one-stage implementation mechanisms, which, as we know, must describe budget sets—subsets of alternatives. Below we offer two examples of such reduction—for discrete and for continuous communication problems. (Note also that if we were just interested in in-equilibrium communication, and restricted our attention to economic environments in which the NVP property holds vacuously, any Nash equilibrium of Mechanism 2 would be a case (a) equilibrium, describing which amounts to describing a single alternative.)

4.1 *Discrete communication: Exponential reduction*

It is known in the communication complexity literature that going from one-stage to two-stage communication protocols sometimes allows an exponential reduction in the communication complexity measured in bits (Kushilevitz and Nisan 1997, Section 4.2). Using Mechanism 2, we can see that such exponential reduction can also be achieved for the Nash implementation problem. (Note that multistage mechanisms cannot generate a more than exponential reduction in communication complexity, because every

extensive-form game can be converted into a one-stage normal form with at most an exponential increase in communication by the argument given in Section 2.4.)

For example, take the Pareto efficient choice rule

$$F(R) = \left\{ x \in X : X = \bigcup_{i \in N} L(x, R_i) \right\} \quad \forall R \in \mathcal{R}$$

and consider the universal preference domain $\mathcal{R} = \mathcal{P}^N$. For this domain, the minimally informative verifying budget equilibria (1) are the *partitional equilibria* (B, x) that support x, i.e., those in which $\bigcup_i B_i = X$ and $B_i \cap B_j = \{x\}$ for all $i, j \in N$. (In words, each alternative in $X \setminus \{x\}$ must belong to *exactly one* agent's budget set.) Furthermore, any such partitional equilibrium must be used for full verification of F. Indeed, for every partitional equilibrium (B, x), we can find a state $R \in \mathcal{P}^N$ in which $L(x, R_i) = B_i$ for all i and thus $x \in F(R)$. Then (B, x) is a unique partitional equilibrium verifying the desirability of alternative x in state R.

There are XN^{X-1} partitional equilibria (choose $x \in X$ and allocate each of the alternatives in $X \setminus \{x\}$ to a budget set). Describing such an equilibrium thus requires $\lceil \log_2(XN^{X-1}) \rceil = \lceil \log_2 X + (X-1)\log_2 N \rceil$ bits. As X grows large, this communication cost is asymptotically proportional to X, which is exponentially larger than that of simply naming an alternative (which takes $\lceil \log_2 X \rceil$ bits). In fact, the communication cost is comparable to that of full revelation of an agent's preferences, which is asymptotically equivalent to $\log_2 X! \sim X \log_2 X$ bits as $X \to \infty$. By Lemma 1, this communication cost also bounds below the communication complexity of a one-stage mechanism that fully Nash implements F.

Compare this with the multistage Mechanism 2, whose communication complexity is at most $5\lceil \log_2 X \rceil + 4N \cdot \lceil \log_2 N \rceil$ bits—exponentially lower as the number X of alternatives grows. Intuitively, the exponential savings arise because instead of describing budget sets, we simply allocate a given alternative to a budget set in any play of the mechanism.

4.2 *Continuous communication: From infinite to finite dimensional*

Consider the problem of implementing Pareto efficiency with quasilinear preferences in which a unit of a divisible good is to be allocated among the agents, along with the transfers of numeraire. Thus,

$$X = \left\{ (q, t) \in \mathbb{R}^N_+ \times \mathbb{R}^N : \sum_i q_i = 1, \sum_i t_i = 0 \right\},$$

where $q_i \geq 0$ is agent i's allocation of the nonmonetary good and t_i is his consumption of numeraire. Thus, X is a $2(N-1)$-dimensional space.

Each agent i's preferences are described by a quasilinear utility function of the form $u_i(q_i) + t_i$, where u_i can be an arbitrary nondecreasing function. Note that the space of such utility functions is infinite-dimensional (even if we impose arbitrary smoothness restrictions on the functions). The Pareto efficient allocations $(q, t) \in X$ in this setting are characterized by maximizing the total surplus $\sum_i u_i(q_i)$.

Calsamiglia (1977) shows that the problem of verifying efficiency in this setting requires infinite-dimensional communication. Segal (2007) rederives the result using the fact that any verification protocol, even with two agents, must reveal an infinite-dimensional nonlinear price function $[0,1] \to \mathbb{R}$ for consumption of the good for one agent in terms of the numeraire. This result implies that the one-stage Nash implementation problem also requires infinite-dimensional communication.

However, for multistage Nash implementation, we can use Mechanism 2, in which only 5 alternatives are described in any play, using a total of $10(N-1)$ real numbers (sending bits is "free" relative to the real numbers). Intuitively, instead of describing numeraire prices for all possible allocations of the divisible good, in any (off-equilibrium) play of Mechanism 2 the agents' approval strategies only need to describe, for one proposed challenge, which agents can afford this challenge. Thus, we learn only about the prices of at most one nonmoneraty allocation q instead of the prices of all possible allocations, which makes the communication finite- instead of infinite-dimensional.

References

Aliprantis, Charalambos D. and Kim C. Border (1999), *Infinite Dimensional Analysis*, second edition. Springer, Berlin. [57]

Calsamiglia, Xavier (1977), "Decentralized resource allocation and increasing returns." *Journal of Economic Theory*, 14, 263–283. [55, 68, 69]

Chakravorti, Bhaskar (1991), "Strategy space reduction for feasible implementation of Walrasian performance." *Social Choice and Welfare*, 8, 235–245. [52]

Dubey, Pradeep (1982), "Price-quantity strategic market games." *Econometrica*, 50, 111–126. [52]

Duggan, John (2003), "Nash implementation with a private good." *Economic Theory*, 21, 117–131. [52]

Dutta, Bhaskar, Arunava Sen, and Rajiv Vohra (1995), "Nash implementation through elementary mechanisms in economic environments." *Economic Design*, 1, 173–203. [52]

Fadel, Ronald and Ilya Segal (2009), "The communication cost of selfishness." *Journal of Economic Theory*, 144, 1895–1920. [52]

Hayek, Friedrich A. von (1945), "The use of knowledge in society." *American Economic Review*, 35, 519–530. [55]

Hurwicz, Leonid (1979), "Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points." *Review of Economic Studies*, 46, 217–225. [52]

Hurwicz, Leonid and Stanley Reiter (2006), *Designing Economic Mechanisms*. Cambridge University Press, Cambridge. [54, 59, 63]

Jackson, Matthew O. (1992), "Implementation in undominated strategies: A look at bounded mechanisms." *Review of Economic Studies*, 59, 757–775. [52]

Kushilevitz, Eyan and Noam Nisan (1997), *Communication Complexity*. Cambridge University Press, Cambridge. [55, 58, 59, 67]

Li, Qi, Shinsuke Nakamura, and Guogiang Tian (1995), "Nash-implementation of the Lindahl correspondence with decreasing returns to scale technologies." *International Economic Review*, 36, 37–52. [52]

Maskin, Eric (1999), "Nash equilibrium and welfare optimality." *Review of Economic Studies*, 66, 23–38. [51, 56]

McKelvey, Richard D. (1989), "Game forms for Nash implementation of general social choice correspondences." *Social Choice and Welfare*, 6, 139–156. [52, 54, 61]

Moore, John and Rafael Repullo (1988), "Subgame perfect implementation." *Econometrica*, 56, 1191–1220. [64]

Nisan, Noam and Ilya Segal (2006), "The communication requirements of efficient allocations and supporting prices." *Journal of Economic Theory*, 129, 192–224. [55]

Peleg, Bezalel (1996), "A continuous double implementation of the constrained Walras equilibrium." *Economic Design*, 2, 89–97. [52]

Reichelstein, Stefan (1984), *Information and Incentives in Economic Organization*. Ph.D. thesis, Northwestern University. [52]

Reichelstein, Stefan and Stanley Reiter (1988), "Game forms with minimal message spaces." *Econometrica*, 56, 661–692. [52, 63]

Saijo, Tatsuyoshi (1988), "Strategy space reduction in Maskin's theorem: Sufficient conditions for Nash implementation." *Econometrica*, 56, 693–700. [52]

Saijo, Tatsuyoshi, Yoshikatsu Tatamitani, and Takehiko Yamato (1996), "Toward natural implementation." *International Economic Review*, 37, 949–980. [52]

Saijo, Tatsuyoshi, Yoshikatsu Tatamitani, and Takehiko Yamato (1999), "Characterizing natural implementability: The fair and Walrasian correspondences." *Games and Economic Behavior*, 28, 271–293. [52]

Schmeidler, David (1980), "Walrasian analysis via strategic outcome functions." *Econometrica*, 48, 1585–1593. [52]

Segal, Ilya (2007), "The communication requirements of social choice rules and supporting budget sets." *Journal of Economic Theory*, 136, 341–378. [52, 53, 54, 55, 56, 57, 59, 61, 69]

Sjostrom, Tomas (1996), "Implementation by demand mechanisms." *Economic Design*, 1, 343–354. [52]

Suzuki, Takeshi (2009), "Natural implementation in public goods economies." *Social Choice and Welfare*, 33, 647–664. [52]

Svensson, Lars-Gunnar (1991), "Nash implementation of competitive equilibria in a model with indivisible goods." *Econometrica*, 59, 869–877. [52]

Tian, Guoqiang (1989), "Implementation of the Lindahl correspondence by a single-valued, feasible, and continuous mechanism." *Review of Economic Studies*, 56, 613–621. [52]

Tian, Guoqiang (1992), "Implementation of the Walrasian correspondence without continuous, convex, and ordered preferences." *Social Choice and Welfare*, 9, 117–130. [52]

Tian, Guoqiang (1994), "Implementing Lindahl allocations by a withholding mechanism." *Journal of Mathematical Economics*, 23, 323–337. [52]

Williams, Steven R. (1986), "Realization and Nash implementation: Two aspects of mechanism design." *Econometrica*, 54, 139–151. [52]

Yoshihara, Naoki (1999), "Natural and double implementation of public ownership solutions in differentiable production economies." *Review of Economic Design*, 4, 127–151. [52]

Submitted 2009-6-12. Final version accepted 2009-11-11. Available online 2009-9-29.