Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 136 (2007) 341–378
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

The communication requirements of social choice rules and supporting budget sets

Ilya Segal*

Department of Economics, Stanford University, Stanford, CA 94305-6072, USA

Received 31 March 2005; final version received 5 September 2006
Available online 27 March 2007

Abstract

The paper examines the communication requirements of social choice rules when the (sincere) agents privately know their preferences. It shows that for a large class of choice rules, any minimally informative way to verify that a given alternative is in the choice rule is by verifying a "budget equilibrium", i.e., that the alternative is optimal to each agent within a "budget set" given to him. Therefore, *any* communication mechanism realizing the choice rule must find a supporting budget equilibrium. We characterize the class of choice rules that have this property. Furthermore, for any rule from the class, we characterize the minimally informative messages (budget equilibria) verifying it. This characterization is used to identify the amount of communication needed to realize a choice rule, measured with the number of transmitted bits or real variables. Applications include efficiency in convex economies, exact or approximate surplus maximization in combinatorial auctions, the core in indivisible-good economies, and stable many-to-one matchings.

JEL classification: D71; D82; D83

Keywords: Communication complexity; Message space dimension; Informational efficiency; Nondeterministic communication (verification); Price mechanisms; Social choice rules; Pareto efficiency; Stability; Envy-free allocations; Approximation; Convex economies; Indivisible goods; Combinatorial auctions; Stable many-to-one matching

1. Introduction

This paper considers the problem of finding allocations that satisfy certain social goals when economic agents have private information regarding their preferences. This problem has received renewed interest in the literature on "market design"—in particular, in two-sided matching (e.g., [44]) and combinatorial auctions (e.g., [5]). The goals of market design include exact or

* Fax: +1 650 725 5702.
E-mail address: ilya.segal@stanford.edu.

0022-0531/$ - see front matter
doi:10.1016/j.jet.2006.09.011

approximate efficiency, voluntary participation, stability to group deviations, and some notions of fairness. A key theme in the literature is that incentives alone do not determine the choice of the mechanism. Indeed, if incentive compatibility were the only concern, it could be ensured with a direct revelation mechanism. However, full revelation of agents' preferences is often impractical, for several reasons. First, sometimes full revelation would require a prohibitive amount of communication: for example, in a combinatorial auction, a bidder would have to communicate his valuations for all possible bundles of objects, and the number of such bundles grows exponentially with the number of objects. Second, agents may have an "evaluation cost" of learning their own preferences. Finally, agents may desire to keep some preference information private (e.g., because they are worried about possible abuse of this information by the designer or other agents in future interactions).

For all these reasons, the "market design" literature has proposed a variety of mechanisms that achieve the desired goals without fully revealing agents' preferences. This raises the question: What is the minimal information that must be elicited by the designer in order to achieve the goals (even if agents are sincere)?

The problem of communication in economic mechanisms was first discussed by Hayek [19], who called attention to the "problem of the utilization of knowledge that is not given to anyone in its totality," when "practically every individual... possesses unique information of which beneficial use might be made." Hayek argued that "we cannot expect that this problem will be solved by first communicating all this knowledge to a central board which, after integrating all knowledge, issues its orders." Instead, "the ultimate decisions must be left to the people who are familiar with the... particular circumstances of time and place." At the same time, the decisions must be guided by prices, which summarize the information needed "to co-ordinate the separate actions of different people." While Hayek did not discuss allocation mechanisms other than the price mechanism and central planning (full revelation), he noted that "nobody has yet succeeded in designing an alternative system" that would fully utilize individual knowledge.

While price mechanisms have received extensive scrutiny since Hayek, existing research has failed to answer the following questions:

(1) Is it ever *necessary* to find some supporting prices in order to achieve social goals?
(2) For which *preference domains* is it necessary to find supporting prices?
(3) For which *social goals* is it necessary to find supporting prices?
(4) *What kind of prices* verify a given social goal on a given preference domain while revealing the minimal necessary information?

Economists have often justified the use of price mechanisms with the Fundamental Theorems of Welfare Economics. However, these theorems fail to address even question (1). Indeed, the First Welfare Theorem says that supporting prices are sufficient to verify Pareto efficiency, but does not establish their necessity. The Second Welfare Theorem only says that supporting prices can be constructed for a given Pareto efficient allocation once all the information about the economy is available. However, once all the information is available, the desired allocation can be imposed directly, without using prices. The theorems have nothing to say about possible efficient nonprice mechanisms in an economy with distributed knowledge of preferences.

Similarly, economists have emphasized the interpretation of prices as the dual variables (Lagrange multipliers) for an optimization program. (Just as the Second Welfare Theorem, this interpretation is based on the Separating Hyperplane Theorem.) However, there are many optimization algorithms that do not use dual variables and do not find their equilibrium values (e.g., the simplex method and the ellipsoid method for linear programs [27]). Thus, question (1) cannot be answered in the affirmative based on purely computational considerations.

A better understanding of the role of prices is offered by the economic literature on the "informational efficiency" of price equilibrium. In contrast to all the previously mentioned approaches, the literature followed Hayek's insight in emphasizing that the allocation mechanism must operate in an economy with *decentralized knowledge of preferences*. The seminal papers of Hurwicz [20] and Mount and Reiter [38] showed that in convex exchange economies in which preference information is decentralized, the Walrasian price mechanism uses the least-dimensional message space among all Pareto efficient verification mechanisms satisfying a continuity restriction. Jordan [24] strengthened this result by showing that the Walrasian mechanism is a *unique* individually rational mechanism with this property. These results were later extended to convex economies with public goods and externalities [46,52]. While providing an important inspiration for the present paper, this literature still comes short of answering questions (1)–(4). Indeed, it does not answer (1), because it focuses on dimensionally minimal continuous mechanisms, and does not rule out the possibility that either discontinuous or slightly more complex continuous mechanisms could achieve efficiency without revealing supporting prices. It does not answer (2), because it only considers settings in which agents have convex preferences over divisible allocations. [1] In fact, the typical continuity restriction in the literature rules out the communication of discrete allocations, and so makes it inapplicable to most market design settings. The literature does not answer (3), because it restricts attention to the goal of Pareto efficiency. As noted by Nisan and Segal [40, Subsection 7.2], this restriction may overstate the hardness of the problem, because in some settings (notably that of Calsamiglia [4]) permitting a very small inefficiency allows a dramatic reduction in the communication cost. In other settings (such as matching without side transfers), efficiency may be achieved trivially, and the designer may be interested in other objectives, such as voluntary participation, stability to group deviations, or some notions of fairness. The literature does not answer (4), because of its *ad hoc* focus on linear-price equilibria, which fail to exist in many important social choice problems. [2]

The present paper answers questions (1)–(4). It examines communication protocols realizing a social choice rule when the (sincere) agents privately know their preferences. A general communication protocol can be viewed as a multi-stage (extensive-form) message game. However, a simple lower bound on this problem is offered by an omniscient oracle's problem of *verifying* the desirability of an alternative. This problem is known as the "verification problem" in the informational efficiency literature and as the "nondeterministic problem" in computer science.

In one special class of verification protocols, the oracle proposes an alternative and gives each agent a *budget set*—a subset of social alternatives (which could in general be delineated by personalized and nonlinear prices). Each agent is asked to verify that the proposed alternative is optimal to him within his budget set. A choice rule can be verified with such a "budget protocol"

[1] Calsamiglia [4] considered the communication burden with nonconvex preferences over divisible goods, but failed to note the role of prices in this setting.

[2] Another related result is obtained by Parkes [41]. He considers the combinatorial auction problem with quasilinear preferences and shows the necessity of revealing supporting prices by those communication languages that reveal so-called "outcome-independent" information and implement surplus-maximizing allocations. This result still does not answer questions (1)–(4), because it considers a restricted set of communication mechanisms, a specific allocation setting, and only the goal of surplus maximization. Parkes's proof uses the duality theory for optimization problems, and thus could not be easily extended to social choice rules that cannot be described as solutions to a maximization problem (including the Pareto rule in the presence of wealth effects).

if and only if it is monotonic (in the sense of Maskin [33]). This result generalizes the classical welfare theorems, by characterizing the choice rules which can be verified with price mechanisms. However, just like the classical welfare theorems, this result does not preclude the possibility that the same choice rule could be realized with a completely different mechanism, that does not reveal price information and perhaps requires much less communication.

Thus, we proceed to characterize the choice rules for which price revelation is necessary, i.e., those satisfying *the Budget Equilibrium Revelation Property* (BERP): *any* protocol verifying the choice rule must reveal enough information to construct supporting budget sets verifying the rule. We show that the choice rules satisfying BERP are characterized by the property of *intersection monotonicity*, which is a strengthening of monotonicity, and which is satisfied by such important rules as Pareto, approximate Pareto, the core, stable matching, and no-envy rules. For all these choice rules, any verification protocol (and therefore any communication) must reveal supporting budget sets.

What appears striking about this result is that even in a social choice problem with sincere agents, as long as preference information is decentralized, a minimally informative verification mechanism asks the agents to pursue their individual objectives independently within their budget sets. Our intuition for this result is that intersection monotonicity postulates certain congruence between the agents' individual preferences and the social goals. Given this congruence, as suggested by Hayek [19], communication can be minimized by giving agents some freedom to utilize their individual knowledge, while coordinating their choices by a careful design of budget sets.

This brings us to question (3): how to design supporting budget sets to verify a given choice rule in a minimally informative way (in a partial informativeness order). A key observation here is that if an agent is offered a larger budget set, the fact that the proposed alternative is optimal to him within this set reveals more information about his preferences. On the one hand, the agents' budget sets must be large enough so as to verify that the proposed alternative is in the choice rule. On the other hand, supporting budget sets that are too large reveal more information about preferences than necessary for the verification. We propose an algorithm that constructs the *minimally informative* budget equilibria verifying that a given alternative is desirable. (When there are many equally informative budget equilibria, we select among them the ones with the largest budget sets.) Under BERP, such budget equilibria exhaust all the minimally informative verifying messages. Application of the algorithm to several well-known social choice problems yields the following characterizations:

- The minimally informative messages verifying Pareto efficiency in an exchange economy with smooth convex preferences are equivalent to *Walrasian equilibria*, in which the budget sets are delineated by linear and anonymous prices.
- The minimally informative messages verifying Pareto efficiency in a social choice problem with numeraire are equivalent to the *valuation equilibria* of Mas-Colell [31], in which the budget sets are delineated by nonlinear personalized prices whose sum across agents is independent of the public decision.
- The minimally informative messages verifying the approximation of Pareto efficiency in a social choice problem with numeraire within some $\delta > 0$ (as measured by the compensating variation in terms of numeraire) are equivalent to *δ-valuation equilibria*, in which the sum of the nonlinear personalized prices across agents for any off-equilibrium public decision exceeds by δ that for the equilibrium decision.
- The minimally informative messages verifying Pareto efficiency and individual rationality on the universal preference domain are equivalent to *partitional equilibria*, in which the

agents' budget sets include the status-quo alternative and partition all the other off-equilibrium alternatives.

- The minimally informative messages verifying the stability of a many-to-one two-sided matching are equivalent to *match-partitional equilibria*, in which each off-equilibrium match is allocated to either partner's budget set (but not both).

While these results are interesting in themselves, their practical importance stems from their application to identifying the communication cost of the relevant choice rule. In this paper, we focus on the traditional measures of communication cost as the number of announced bits as in the discrete "communication complexity" literature surveyed by Kushilevitz and Nisan [29], or as the number of announced real variables as in the economic literature on "dimension of message spaces." While these two literatures have examined the communication costs of a number of important problems, the present paper offers a new systematic approach to analyzing the communication costs for a large and important class of problems. The new approach is based on the economic idea that it is necessary to discover supporting prices, and on an algorithm for constructing the form of prices (more generally, budget sets) that must be discovered to solve a given problem. Formally, the communication cost of intersection-monotonic choice rules is identified as that of finding a minimally informative verifying budget equilibrium, and the space of such equilibria is characterized for any given choice rule.

Calculation of the communication cost is complicated by the fact that a verification protocol need not use *all* minimally informative verifying budget equilibria, since it only needs to verify *one* desirable alternative in a given state rather than all of them. For example, in a convex exchange economy, we can realize Pareto efficiency using only those Walrasian budget sets that contain an (arbitrarily fixed) endowment allocation, which reduces the dimensionality of the message space. In general, the nondeterministic communication cost of an intersection-monotonic choice rule F is determined by a minimal collection $\mathcal{E}$ of minimally informative budget equilibria verifying F that ensures the existence of an equilibrium from $\mathcal{E}$. Namely, the communication cost of F is exactly that of communicating an equilibrium from $\mathcal{E}$, which requires $\dim \mathcal{E}$ real variables for continuous communication, or $\log_2 |\mathcal{E}|$ bits for discrete communication. This number also bounds below the cost of *deterministic* communication, i.e., *finding* a desirable alternative. (While in some settings there exists a known deterministic communication protocol coming close to achieving this lower bound, the general problem of identifying the deterministic communication cost of a social choice rule appears to be much harder and is not tackled here.)

We illustrate the general approach by identifying (or bounding) the communication cost of several well-known social choice problems. In some of these problems, the communication costs have been previously identified, and so we are merely re-deriving the existing results in a more elegant and systematic way. These include the problems of Pareto efficiency in convex economies (where the communication cost was obtained by [20,38]) and in quasilinear settings (where the cost was by obtained by [40]).[3] In other problems, such as Pareto efficient individually rational allocations in economies with indivisible goods and stable many-to-one matchings, the communication costs have not been known, and so our results are completely new.

To interpret the practical significance of the results, we can think of a problem as "hard" if its communication cost is of the same order of magnitude as full revelation of agents' preferences,

[3] Nisan and Segal [40] also report a result on the necessity to reveal prices in order to verify surplus maximization in a quasilinear economy. However, [40] does not consider social choice problems with other goals or preference domains, nor does it construct minimally informative verifying price equilibria for the problem it does consider.

and “easy” if it much lower. We show that some problems are “hard” in this sense—in particular, efficient of even approximately efficient combinatorial allocations, and efficient individually rational allocations in economies of indivisible goods. At the same time, some problems prove “easy”—in particular, Pareto efficiency in convex exchange economies, and stability in many-to-one matchings. The “easiness” of nondeterministic communication in these problems stems from the fact that the space of minimally informative verifying budget equilibria proves to be much smaller than the preference domain. In the stable matching problem, we find that *deterministic* communication is also “easy” when preferences are substitutable: the Gale–Shapley deferred acceptance algorithm finds a stable matching using only slightly more communication than that necessary for verification, and exponentially less than full revelation of preferences.

2. Social choice and communication

2.1. The social choice problem

Let N be a finite set of agents, and X be a set of social alternatives. (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.[4] For any preference relation $R \in \mathcal{P}$ and any alternative $x \in X$, it is convenient to define the relation's *lower contour set* at x, $L(x, R) = \{y \in X : xRy\}$.

Each agent i's preference relation is assumed to be his privately observed *type*, and the set of his possible types is denoted by $\mathcal{R}_i \subset \mathcal{P}$.[5] 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.

2.2. Communication

We now describe the communication procedures solving the social choice problem, using the notation and terminology introduced by Hurwicz [20] and Mount and Reiter [38], as well as in the communication complexity literature [29].

It is well known that the communication cost can be reduced by letting agents send messages sequentially rather than simultaneously. For example, if we want to find a Pareto efficient alternative, agents need not report their preferences between alternatives x and y if it is clear from the preceding messages that y is dominated by z for all of them. Therefore, we must consider multi-round communication protocols.

In the language of game theory, a multi-round communication protocol specifies an extensive-form message game and each agent's strategy in this game (complete message plan contingent on his type and the observed history). Instead of payoffs, the game assigns alternatives to terminal nodes (and so is more properly called a “mechanism”). Agents are assumed to follow the prescribed strategies (but see Section 9 for a discussion of incentive compatibility). Such communication protocols are known in computer science as “deterministic,” because the message sent by an agent at a given information set is fully determined by his type and the preceding messages. A

[4] A preference relation R over set X is a binary relation over X, with xRy interpreted as “x is weakly preferred to y.” It is common to restrict attention to preference relations that are *rational*, i.e., complete and transitive. Rationality will play no role in the general analysis, but it will be assumed in all the applications.

[5] Thus, we focus on “private-value” environments. It would be interesting to extend the analysis to “interdependent-value” environments, in which an agent's preferences may depend on other agents' private information.

protocol *realizes* choice rule F if in every state R it achieves a terminal node to which an alternative from $F(R)$ is assigned. [6]

Characterizing all deterministic communication protocols is a tall order. Analysis is drastically simplified by generalizing the notion of communication to allow what is called "nondeterministic communication" in computer science and "the verification scenario" in economics: imagine an omniscient oracle who knows the true state R and consequently 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 this 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. [7]

While nondeterministic communication is patently unrealistic, it is considered for the following reasons:

1. Any deterministic communication protocol can be represented as nondeterministic by letting all the messages be sent by the oracle instead of the agents, and having each agent accept the message sequence if and only if all the messages sent in his stead are consistent with his strategy given his type. The oracle's message space M is thus identified with the set of the protocol's possible message sequences (terminal nodes). Therefore, any statement about nondeterministic protocols will apply to deterministic protocols as a particular case (this is explained more thoroughly in [29, Chapter 2]).
2. A famous economic example of nondeterministic communication 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. A generalization of this nondeterministic communication is described in the next section.
3. A nondeterministic protocol realizing choice rule F may be viewed as a steady state of an iterative *deterministic* protocol realizing or approximating F. At each stage of the iteration, a message $m \in M$ is announced, and each agent reports a direction in which the message should be adjusted to become "more acceptable" to him. Examples of such adjustment processes include "tatonnement" for finding Walrasian equilibria, "deferred acceptance algorithms" for finding stable matchings, and ascending-bid auctions for finding efficient combinatorial allocations.

Formally, nondeterministic communication can be defined as follows:

Definition 1. A (*nondeterministic communication*) *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 satisfying Privacy Preservation:

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

- $h : M \twoheadrightarrow X$ is the outcome correspondence.

[6] Note that only nonempty-valued choice rules can be realized. Nonempty-valuedness could be ensured by thinking of states $R \in \mathcal{R}$ in which $F(R) = \emptyset$ as "illegal," and allowing any alternative in such states (i.e., redefining $F(R) = X$).

[7] This communication is called "nondeterministic" in computer science because the oracle has to "guess" a message that is acceptable to all agents (and there may be more than one such message in a given state).

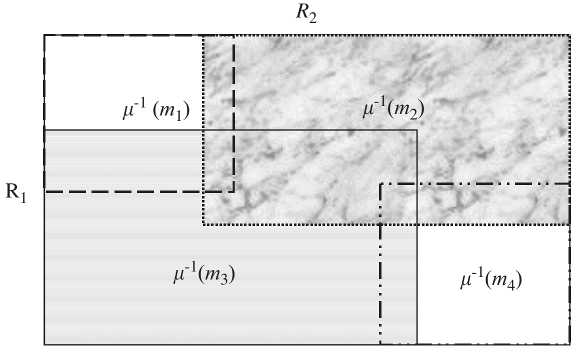

Fig. 1. Nondeterministic communication.

Γ realizes choice rule F if $\emptyset \neq h(\mu(R)) \subset F(R)\ \forall R \in \mathcal{R}$.[8]
Γ fully realizes F if $h(\mu(R)) = F(R)\ \forall R \in \mathcal{R}$.

Privacy Preservation captures the fact that each agent does not observe other agents' types, thus the set of messages acceptable to him is a function $\mu_i(R_i)$ of his own type R_i only.[9] Realization means the set of alternatives generated by the protocol's acceptable messages in state R is a nonempty subset of the set of desirable alternatives $F(R)$, while full realization means that it is exactly $F(R)$. We are ultimately interested in realization, but the concept of full realization will allow comparisons with some existing literature.

Definition 2. Message $m \in M$ in protocol $\Gamma = \langle M, \mu, h\rangle$ *verifies* alternative $x \in X$ in choice rule F if $\mu^{-1}(m) \subset F^{-1}(x)$. ($\Gamma$ and F will be omitted when clear from the context.)

If we are interested in whether a given message correspondence μ can be used to realize choice rule F, without loss we can define the outcome correspondence $h(m)$ to be the set of alternatives verified by message m. Then realization means that, in any state R, *some* alternative is verified by an acceptable message, while full realization means that *any* alternative in $F(R)$ is verified by some acceptable message.

The above concepts have a graphical illustration, discussed in [29], and depicted in Fig. 1. Namely, each $\mu^{-1}(m)$ is the subset of the state space $\mathcal{R}$ on which message $m \in M$ is acceptable. *Privacy Preservation* requires each such subset to be a product set $\mu_1^{-1}(m) \times \cdots \times \mu_N^{-1}(m)$—a "rectangle" in the computer science parlance. Message m verifies alternative x if the corresponding rectangle $\mu^{-1}(m)$ is contained in the set $F^{-1}(x)$ on which x is desirable—in the computer science parlance, the rectangle is "monochromatic". Realization requires that the whole state space be covered by monochromatic rectangles, while full realization requires that each set $F^{-1}(x)$ for $x \in X$ be exactly covered by some set of rectangles.

[8] We use the standard notation for the image of a set: $h(A) = \cup_{m\in A} h(m)$ [1, p. 3].

[9] This is an established term in the economic literature on "informational efficiency," but it differs from the layman's concept of "privacy" in that it places no constraints on the revelation of information in the course of communication.

2.3. Communication costs and informativeness order

As mentioned in the introduction, there could be many different measures of communication costs: the number of bits or real numbers transmitted, the agents' costs of evaluating their preferences, or their loss of privacy. Here we offer a new way to measure the informativeness of communication, which is only a partial rather than complete ordering, but which proves useful for estimating many different scalar measures of communication cost:

Definition 3. Message $m \in M$ in protocol $\langle M, \mu, h \rangle$ *is less informative than* (or *is verified by*) message $\tilde{m} \in \tilde{M}$ in protocol $\left\langle \tilde{M}, \tilde{\mu}, \tilde{h} \right\rangle$ if $\tilde{\mu}^{-1}(\tilde{m}) \subset \mu^{-1}(m)$. Messages m and $\tilde{m}$ are *equivalent* if they are equally informative, i.e., $\mu^{-1}(m) = \tilde{\mu}^{-1}(\tilde{m})$. Message m is a *minimally informative* message verifying alternative $x \in X$ if it verifies x, and any less informative message verifying x is equivalent to m.

In words, message m is less informative than message $\tilde{m}$ if m is accepted in a larger set of states than $\tilde{m}$. In the graphical interpretation, this means that the "rectangle" on which message m is accepted includes the "rectangle" on which message $\tilde{m}$ is accepted. Also, m is a minimally informative message verifying x if it corresponds to a maximal "rectangle" contained in the set $F^{-1}(x)$ of states in which x is desirable.

We will examine how a given choice rule can be realized using messages that are less informative, and possibly minimally informative. Intuitively, this will show *what information* must be revealed in order to realize a given social choice rule F. In turn, this will allow to measure the minimal communication cost needed to realize F. Namely, we can restrict attention to protocols that use only minimally informative verifying messages, without increasing the communication cost. To give one example, if the communication cost of a protocol is defined as the worst-case number of bits it transmits (as in the communication complexity literature), then it can be calculated as the binary logarithm of the size of its message space M (see Section 6 for more detail). Thus, starting with a protocol realizing choice rule F and replacing one message in it with a less informative message that still verifies the same alternative would produce another protocol realizing F, with the same size of message space, and therefore the same communication cost. (But this replacement could also allow us to eliminate some of the messages while still covering the state space with the corresponding "rectangles.") Therefore, when looking for the cheapest nondeterministic protocol realizing F, without loss one can restrict attention to protocols using only minimally informative verifying messages. The same observation will apply to many other measures of communication costs.

3. Budget equilibria and their revelation

Consider a special class of nondeterministic protocols, 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 $(B_1, \ldots, B_N, x)$ if and only if there is no alternative in his budget set B_i that he strictly prefers to the proposed alternative x. $(B_1, \ldots, B_N, x)$ is a *budget equilibrium in state* $R \in \mathcal{R}$ if it is accepted by all agents in this state. [10] Formally, the budget equilibrium correspondence

[10] A number of related concepts have been suggested, including "social equilibrium" [7], "social situations" [13], "effectivity functions" [37], "effectivity forms" [36], "opportunity equilibrium" [25], "attainable sets" [9], and "interactive choice sets" [48]. However, all these papers have motivated the concept by incentive compatibility, rather than deriving it from communication among sincere agents (see Section 9 for a more detailed comparison).

$E : \mathcal{R} \twoheadrightarrow 2^{XN} \times X$ is described as

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

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 correspondence simply implements the proposed alternative:

Definition 4. 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 $h(B, x) = \{x\}\ \forall (B, x) \in M$.

The informativeness of a budget equilibrium message depends on how large the agents' budget sets are. Formally, consider:

Definition 5. For two budget equilibria $(B, x), (B', x') \in 2^{XN} \times X$, (B', x') is *larger than* (B, x) if $x = x'$ and $B_i \subset B_i'\ \forall i \in N$.

It is clear that if budget equilibrium (B', x') is larger than budget equilibrium (B, x), then (B', x') is more informative than (B, x), and so is more likely to verify alternative x. [11]

Which choice rules can be realized by a budget protocol? Classical Welfare Theorems say that any Pareto efficient allocation in a convex exchange economy can be verified with a budget equilibrium (specifically, a Walrasian equilibrium). The theorems have been extended to some "nonclassical" social choice problems. [12] These results can be generalized as follows:

Definition 6 (*Maskin [33]*). 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 $x \in F(R')$.

Theorem 1. *A choice rule F is fully realized by a budget protocol if and only if it is monotonic.* [13]

Proof. That F is fully realized with a budget protocol means that $\forall R \in \mathcal{R}\ \forall x \in F(R)\ \exists B \in 2^{XN}$ such that (a) $(B, x) \in E(R)$, i.e., $B_i \subset L(x, R_i)\ \forall i \in N$, and (b) (B, x) verifies x, i.e., $E^{-1}(B, x) \subset F^{-1}(x)$ Since a larger budget equilibrium is more informative and so more likely to verify x, this is equivalent to checking that the *largest* budget equilibrium (B, x) satisfying (a), which has $B_i = L(x, R_i)\ \forall i \in N$, verifies x. This is in turn equivalent to the monotonicity of F. □

Theorem 1 is not novel: analogous results are stated in [53, Theorem 2], [13, Theorem 10.1.2], [36, Theorem 1], and [25]. The key deficiency of Theorem 1 is that, just like the classical Welfare Theorems, it does not say anything about nonbudget protocols realizing choice rule F, which could possibly reveal less information and have lower communication costs than any budget protocol

[11] For example, when $B_i = \{x\}$ for all i, budget equilibrium (B, x) is uninformative and does not verify x, unless it is always selected by the choice rule.

[12] Including the Pareto rule in public-good economies [35] and general economies with numeraire [31,2,3], and stable many-to-one matching problems with and without transfers [28,18].

[13] This implies that F is *realized* by a budget protocol if and only if has a nonempty-valued monotonic subcorrespondence.

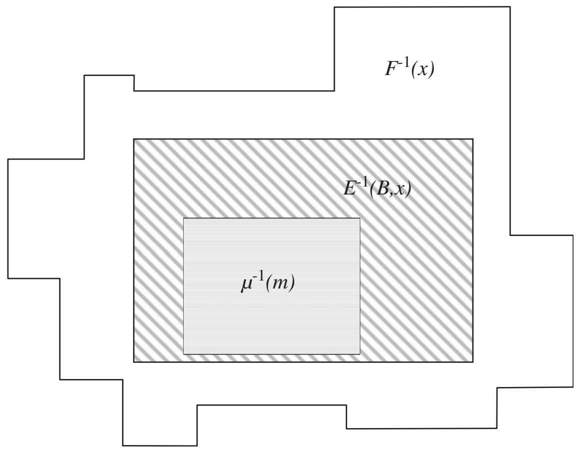

Fig. 2. Budget equilibrium revelation property.

realizing F. We remedy this deficiency by characterizing choice rules that satisfy the following property:

Definition 7. Choice rule F satisfies the *Budget Equilibrium Revelation Property* (*BERP*) if for any message verifying an alternative $x \in X$ in F there exists a less informative budget equilibrium (B, x) that verifies that x is in F.

BERP is illustrated in Fig. 2. When applied to messages m that fully reveal a state R (i.e., $\mu^{-1}(m) = \{R\}$, which would be represented by a single point in Fig. 2), BERP says that for any $x \in F(R)$ we can construct a budget equilibrium (B, x) in state R that verifies x. Thus, BERP implies the classical welfare theorems, and so by Theorem 1 it implies the monotonicity F. However, BERP is stronger, since it requires a budget equilibrium verifying x to be constructed without knowing the exact state, upon observing *any* communication verifying x. This strengthening indeed eliminates some monotonic choice rules, as the following examples demonstrate:

Example 1. Let $N = 1$, $X = \{x, y, z\}$, and $\mathcal{R} = \mathcal{P}$. Take the choice rule

$$F(R_1) = \begin{cases} \{x, y, z\} & \text{if } x R_1 y \text{ or } x R_1 z, \\ \{y, z\} & \text{otherwise} \end{cases} \quad \forall R_1 \in \mathcal{P}.$$

It is easy to see that F is monotonic, hence by Theorem 1 it can be fully realized with a budget protocol. Namely, note that y and z are verified with any budget set, while in all states R_1 in which $x \in F(R_1)$, it can be verified with the budget equilibrium $\left(L\left(x, R_1\right), x\right)$. More generally, budget equilibrium (B_1, x) verifies x if and only if $B_1 \backslash \{x\} \neq \emptyset$.

Now consider the communication protocol in which agent 1 announces "yes" if $L(x, R_1) \neq \{x\}$, in which case x is implemented, and announces"no" otherwise, in which case y is implemented. Message "yes" verifies x, but does not reveal any other alternative in $L(x, R_1)$, thus it does not reveal a budget equilibrium (B_1, x) that would verify x. Therefore, F does not satisfy BERP.

Example 2. Let $N = 1$, let $X = \left\{x \in \mathbb{R}_+^K : \sum_k x_k = 1\right\}$ represent the set of lotteries over a finite set K, and let $\mathcal{R}$ be the set of von Neumann–Morgenstern preferences relations over the lotteries.

That is, $\mathcal{R} = \{R(u) : u \in \mathbb{R}_+^K\}$, where preference relation $R(u)$ is described by $x R(u) y$ if and only if $\sum_k x_k u_k \geqslant \sum_k y_k u_k$. Then any social choice rule F that selects only *strict* lotteries x (i.e., those with $x_k > 0$ for all $k \in K$) is trivially monotonic. (Indeed, if x is a strict lottery and R, R' are two von Neumann–Morgenstern preference relations such that $L(x, R) \subset L(x, R')$, then we must have $R' = R$.) On the other hand, F is not intersection-monotonic if for some strict lottery x, the set $\{u \in \mathbb{R}_+^K : x \in F(R(u))\}$ is not a convex cone. Indeed, if $u'' = \alpha u + \alpha' u'$ with $\alpha, \alpha' \geqslant 0$, then $L(x, R(u)) \cap L(x, R(u')) \subset L(x, R(u''))$, and so for an intersection-monotonic F, $x \in F(R(u)) \cap F(R(u'))$ must imply that $x \in F(R(u''))$.

The choice rules that *do* satisfy BERP are characterized as follows:

Definition 8. Choice rule F is *Intersection-Monotonic* (*IM*) if $\forall \widetilde{\mathcal{R}} = \widetilde{\mathcal{R}}_1 \times \cdots \times \widetilde{\mathcal{R}}_N \subset \mathcal{R}$, $\forall x \in \cap_{R \in \widetilde{\mathcal{R}}} F(R)$, and $\forall R' \in \mathcal{R}$ such that $\cap_{R_i \in \widetilde{\mathcal{R}}_i} L(x, R_i) \subset L(x, R_i')$ $\forall i \in N$ we have $x \in F(R')$.[14]

Theorem 2. *Choice rule F satisfies BERP if and only if it is Intersection-Monotonic.*

Proof. Observe that for any $\widetilde{\mathcal{R}} \subset \mathcal{R}$, $\widetilde{\mathcal{R}} = \mu^{-1}(m)$ for some message $m \in M$ in some protocol $\Gamma = \langle M, \mu, h \rangle$ if and only if $\widetilde{\mathcal{R}}$ is a product set. Thus, that F satisfies BERP means that $\forall \widetilde{\mathcal{R}} = \widetilde{\mathcal{R}}_1 \times \cdots \times \widetilde{\mathcal{R}}_N \subset \mathcal{R}$ $\forall x \in \cap_{R \in \widetilde{\mathcal{R}}} F(R)$ $\exists B \in 2^{XN}$ such that (a) $\widetilde{\mathcal{R}} \subset E^{-1}(B, x)$, i.e., $B_i \subset L(x, R_i)$ $\forall i \in N$ $\forall R_i \in \widetilde{\mathcal{R}}_i$ and (b) (B, x) verifies x, i.e., $E^{-1}(B, x) \subset F^{-1}(x)$. Since a larger budget equilibrium is more informative and so more likely to verify x, this is equivalent to checking that the *largest* budget equilibrium (B, x) satisfying (a), which has $B_i = \cap_{R_i \in \widetilde{\mathcal{R}}_i} L(x, R_i)$ $\forall i \in N$, verifies x. This is in turn equivalent to the intersection monotonicity of F. $\square$

To see directly that IM is a strengthening of monotonicity, take $\widetilde{\mathcal{R}} = \{R\}$ in the definition. Note also that IM is fairly easy to verify: just as with monotonicity, it suffices to check changes in one agent i's preferences holding all other agents' preferences fixed (i.e., letting $\widetilde{\mathcal{R}}_j = \{R_j'\}$ for $j \neq i$)—the full definition would then follow by iterating over agents. Thus, Theorem 2 offers a simple way to check whether a given choice rule satisfies BERP, i.e., whether the revelation of prices is necessary for realizing it. We proceed to identify a large and economically important class of choice rules which indeed have this property.

4. A class of intersection-monotonic choice rules

Definition 9. 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

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

[14] A property with the same name is defined by Miyagawa [36], but he intersects the lower contour sets of different agents, and uses the property for an apparently different purpose. IM is also related to Sjostrom's [50] Condition W, but the latter is much stronger in that it allows to construct supporting budget sets verifying alternative x using no information other than the desirability of x. Therefore, Condition W allows F to be fully realized with a "proposed action" protocol (Ishikida and Marschak [23]), which announces only the alternative to be implemented, and does not announce any supporting budget sets. This condition is not satisfied in any of the applications considered in Section 7.

In words, a CU choice rule is described by specifying for any coalition $S \subset N$ and any candidate alternative $x \in X$ a "blocking set" $\beta(x, S) \subset X$—the set of alternatives that S can use to block x. Alternative x is "unblocked" by coalition S if it is weakly Pareto efficient for its members within its blocking set $\beta(x, S)$, i.e., if it is not possible to make all members of S strictly better off within $\beta(x, S)$. [15] $x \in F(R)$ if it is unblocked by all coalitions. [16] Note that a CU choice rule defined on any preference domain is extendable to the universal preference domain $\mathcal{R} = \mathcal{P}^N$ using the same blocking correspondence.

Now we describe several important examples of CU choice rules. Note that according to Definition 9, the empty coalition $S = \emptyset$ will block in any state, hence F can only include alternatives in the set $\bar{X} = \{x \in X : \beta(x, \emptyset) = \emptyset\}$, which we interpret as the set of *feasible* alternatives. [17] With this notation, a CU choice rule will include those feasible alternatives that are not blocked by nonempty coalitions.

- *The Pareto rule:* $\beta(x, S) = \bar{X}$ if $S = N$, $\emptyset$ if $S \notin \{N, \emptyset\}$. That is, the grand coalition can block any alternative with any feasible alternative, and no other nonempty coalition has any blocking power. [18]
- *The approximate Pareto rule:* Let $\beta(x, S) = X^{\delta}$ if $S = N$, $\emptyset$ if $S \notin \{N, \emptyset\}$, where $X^{\delta} \subset \bar{X}$ is interpreted as the set of alternatives that waste a given amount $\delta > 0$ of resources. In words, a feasible alternative x is desirable if it is impossible to make everyone strictly better off while wasting amount δ of resources. Thus, δ is the "compensating variation" measure of inefficiency—the amount of resources that could be extracted from the agents while compensating all of them for the change. There are many ways to define X^{δ} in an economy with multiple divisible goods. For example, letting X^{δ} consist of allocations that waste proportion δ of the economy's aggregate endowment results in F choosing allocations whose "coefficient of resource utilization" [6] is at least $1 - \delta$. Alternatively, if X^{δ} consists of allocations that waste amount δ of a specific good—"numeraire," and if preferences are quasilinear in numeraire, then F chooses allocations that achieve within δ of the maximum possible surplus.
- *The core:* For all $S \neq \emptyset$, $\beta(x, S) = \varepsilon(S)$—the "effectivity set" of coalition S. Pareto efficiency is imposed by letting $\varepsilon(N) = \bar{X}$. Individual rationality (i.e., voluntary participation) is imposed by letting $\varepsilon(\{i\}) = \{x_0\}$ for all $i \in N$, where $x_0 \in X$ is the "status-quo" alternative. Specification of effectivity sets for intermediate coalitions reflects the coali-

[15] We use weak Pareto efficiency because the strong Pareto rule is not even monotonic, let alone IM. Note, however, that the weak and strong Pareto criteria coincide for preferences that are strictly monotonic and nonsatiated in some divisible economic good.

[16] Such choice rules are also known as "respecting group rights," with $y \in \beta(x, S)$ interpreted as the "one-way right" of coalition S to replace alternative x with alternative y [17, Section 5]. The "rights" literature, initiated by Sen [47], is concerned with the problem that individual and group rights may be incompatible with each other on the universal preference domain, i.e., that "group rights-respecting" choice rules may be empty-valued. In the applications considered in Section 7, the preference domains and coalitional rights will be defined to ensure nonempty-valuedness.

[17] For example, the empty coalition may be responsible for the satisfaction of resource constraints. We permit X to be larger than $\bar{X}$ to allow budget sets that include infeasible allocations, as they may in the Walrasian protocol. If X consisted only of feasible allocations in a convex exchange economy, the Walrasian choice rule would not be monotonic [21], hence it could not be fully realized with a budget protocol.

[18] If any preference $R_i \in \mathcal{R}_i$ of agent i has a maximal alternative in the feasible set $\bar{X}$, the Pareto rule could be realized simply by letting the agent choose this alternative. To rule out this dictatorial solution, the literature on the communication requirements of the Pareto rule has either considered settings in which the feasible set is noncompact, or introduced additional restrictions on the alternatives.

tions' powers. For example, the majority voting (Condorcet) choice rule is described by $\varepsilon(S) = \bar{X}$ if $|S| \geqslant N/2$, $\emptyset$ otherwise. In an exchange economy, $\varepsilon(S)$ is usually defined by allowing the members of S to reallocate resources among each other. We can also define *the approximate core (quasi-core, epsilon-core)* of an exchange economy, by letting $\varepsilon(S)$ consist of allocations that destroy at least amount $\delta_S > 0$ of resources available to the coalition. [19]

In the above examples the blocking sets $\beta(x, S)$ did not depend on the candidate alternative x, but in other examples this dependence is important:

- *Stable network*: Let $X = 2^{N\times N}$—i.e., an alternative $x \in X$ is a binary relation on X (a list of ordered pairs of agents). $(i, j) \in x$ is interpreted as the directed link from agent $i \in N$ to agent $j \in N$ in network $x \in X$. The blocking sets are described by

$$\beta(x, S) = \{y \in X : y\backslash(S \times N) = x\backslash(S \times N)\}.$$

  In words, members of coalition S can change only their outgoing links. A *stable matching* problem (such as the one studied by Roth and Sotomayor [44]) obtains as particular case by defining the matching relation as the symmetric part of x (i.e., a match is a bidirectional link). The blocking sets described above allow a coalition to break matches with outsiders but not create new matches with them, which corresponds to the concept of *setwise stability* [51].
- *The envy-free rule*: Let $X = X_1 \times \cdots \times X_N$, where $x_i \in X_i$ is interpreted as agent i's component of the allocation. Let

$$\beta(x, \{i\}) = \left\{y \in X : \left(y_i, y_j, y_{-i-j}\right) = \left(x_j, x_i, x_{-i-j}\right), j \in N\right\} \quad \forall i \in N,$$
$$\beta(x, S) = \emptyset \quad \text{for } |S| > 1.$$

  In words, any individual agent can block an alternative by "trading places" with another agent. We can also define the *approximate envy-free (or bounded-envy) rule* [30], by letting $\beta(x, \{i\})$ consist of allocations in which agent i pays a numeraire penalty $\delta > 0$ for trading places with another agent.
- *Any combination of the above goals:* For any family $\{F_k\}_{k\in K}$ of CU rules, the intersection rule, given by $F(R) = \cap_{k\in K} F(R)$, is also CU. Indeed, it is described by the blocking correspondence $\beta(x, S) = \cup_{k\in K}\beta_k(x, S)$, where β_k is the blocking correspondence describing F_k.

Lemma 1. *Any CU choice rule is IM.*

Proof. Suppose in negation that a CU choice rule F described by blocking correspondence β is not IM, i.e., $\exists\widetilde{\mathcal{R}} = \widetilde{\mathcal{R}}_1 \times \cdots \times \widetilde{\mathcal{R}}_N \subset \mathcal{R}\ \exists R' \in \mathcal{R}\ \exists x \in X$ such that (a) $x \in F(R)$ $\forall R \in \widetilde{\mathcal{R}}$, (b) $\cap_{R_i\in\widetilde{\mathcal{R}}_i} L(x, R_i) \subset L\left(x, R_i'\right) \forall i \in N$, but (c) $x \notin F\left(R'\right)$. (c) means that $\exists S \subset N\ \exists y \in \beta(x, S)$ such that $y \notin L\left(x, R_i'\right) \forall i \in S$. By (b), this implies that $\forall i \in S\ \exists R_i^* \in \widetilde{\mathcal{R}}_i : y \notin L\left(x, R_i^*\right)$. Choosing such $R_i^* \in \widetilde{\mathcal{R}}_i$ for all $i \in S$ and arbitrary $R_i^* \in \widetilde{\mathcal{R}}_i$ for all

[19] In particular, [49] requires the destruction of amount δ_S of numeraire, [26] requires the destruction of amount δ_S of each good, and [34, Subsection 3.3] requires the destruction of share δ_S of a given commodity bundle.

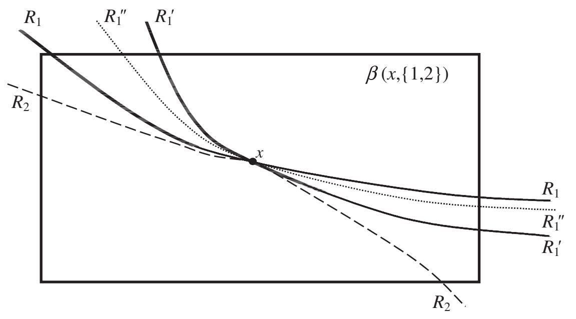

Fig. 3. Intersection-monotonicity of coalitionally unblocked rules.

$i \in N \backslash S$, we obtain $R^* \in \widetilde{\mathcal{R}}$ such that $\beta(x, S) \nsubseteq \cup_{i \in S} L\left(x, R_i^*\right)$, and therefore $x \notin F(R^*)$, contradicting (a). $\square$

To illustrate the proof of Lemma 1, take a CU choice rule F, and suppose that for two different preferences R_1, R_1' of agent 1 and some preference profile R_{-1} of other agents, we have $x \in F(R_1, R_{-1})$ and $x \in F\left(R_1', R_{-1}\right)$. This means that in each state, x is Pareto efficient for each coalition within its blocking set. For example, the situation for coalition $\{1, 2\}$ is illustrated in Fig. 3, in which the box represents the coalition's blocking set $\beta(x, \{1, 2\})$, agent 1's preferences are increasing in the top-right direction, and agent 2's preferences are increasing in the bottom-down direction (as in the traditional Edgeworth box). The Pareto efficiency of x for coalition $\{1, 2\}$ within the box in states (R_1, R_{-1}) and $\left(R_1', R_{-1}\right)$ means that the indifference curves representing R_1 and R_1' passing through x both lie above the indifference curve representing R_2 passing through x. Now take a third preference R_1'' for agent 1 such that $L(x, R_1) \cap L\left(x, R_1'\right) \subset L\left(x, R_1''\right)$. In Fig. 3 this means that the indifference curve representing R_1'' passing through x lies above the lower envelope of the curves representing R_1 and R_1'. But this implies that the curve representing R_1'' still lies above that representing R_2, and therefore in state $\left(R_1'', R_{-1}\right)$, x remains Pareto efficient for coalition $\{1, 2\}$ within its blocking set. Since the same argument works for all coalitions, we see that x remains unblocked in state $\left(R_1'', R_{-1}\right)$, hence $x \in F\left(R_1'', R_{-1}\right)$. Iterating the argument by sequentially changing the preferences of agents 2, 3, etc., we can see that F is IM.

The converse to Lemma 1 is not true:

Example 3. Let $N = 2$, $X = \{x, y, z\}$, and $R = \mathcal{P}^2$. Take the choice rule

$$F(R) = \begin{cases} \{x, y, z\} & \text{if } x R_1 y \text{ or } x R_2 z \\ \{y, z\} & \text{otherwise} \end{cases} \quad \forall R \in \mathcal{P}^2.$$

It is easy to verify that F is IM. On the other hand, if F were a CU choice rule described by blocking correspondence β, we would have $y, z \notin \beta(x, S)\ \forall S \subset N$ (since $x \in F(R)$ in the states

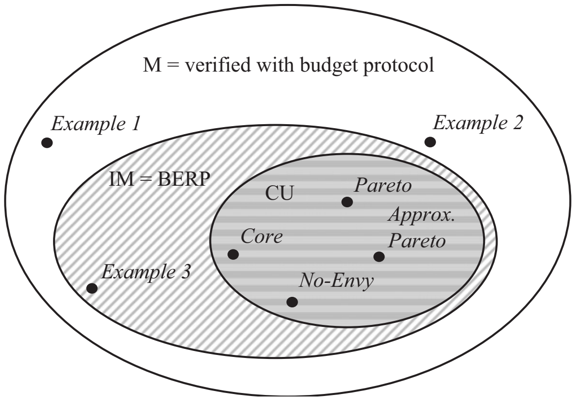

Fig. 4. Venn diagram for choice rules.

R in which the agents $i = 1, 2$ share a strict preference ordering $yR_i xR_i z$ or $zR_i xR_i y$), but then we would have $x \in F(R)\ \forall R \in \mathcal{P}^2$, which is not true.

A Venn diagram for choice rules summarizing the above results is drawn in Fig. 4.[20]

5. Minimally informative verifying equilibria

We next address the question of *which* supporting budget equilibria must be revealed to verify a given choice rule. We do it by characterizing the minimally informative messages verifying a given choice rule, which, under BERP, are all equivalent to budget equilibrium messages. Recall that a budget equilibrium is more informative the larger its budget sets are, thus the minimally verifying informative budget equilibria must have large enough budget sets to verify the choice rule, but not any larger.

First we justify the focus on minimally informative verifying messages by showing that any message m verifying alternative x verifies a minimally informative message $\tilde{m}$ verifying x. When the state space $\mathcal{R}$ is finite, $\tilde{m}$ can be constructed by starting with m and finding progressively strictly less informative messages verifying x while this is possible (the procedure terminates since the number of possible nonequivalent messages is finite). For an infinite state space, we

[20] An alternative way to understand this classification of choice rules is by describing them with boolean formulas (as proposed by Rubinstein [45]). Namely, we can describe a rule F by giving, for each $x \in X$, a boolean formula to calculate the truth value of $x \in F(R)$ from the truth values of the atoms $\{yR_i z\}_{i\in N;\, y,z\in X}$ that describe the preference profile. It can then be seen that F is monotonic if and only if the formula can be taken to depend only on the atoms in $\{xR_i y\}_{i\in N;\, y\in X}$ and to be monotone in these atoms. Any such formula can be written in a "Monotone Conjunctive Normal Form" (MCNF), i.e., as a conjunction of disjunctions of individual atoms. F is IM if and only it we can use an MCNF in which no disjunctive clause contains two atoms $xR_i y$ and $xR_i z$ with the same i but $y \neq z$ (which is violated by Example 1). F is CU if and only if we can use an MCNF in which no disjunctive clause contains two atoms $xR_i y$ and $xR_j z$ with $y \neq z$ for some i, j (which is violated by Example 3).

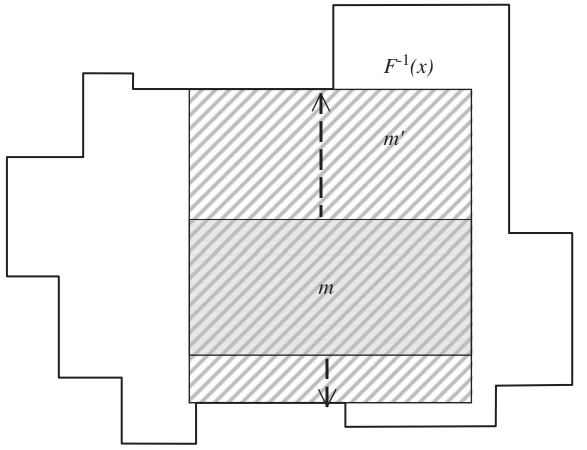

Fig. 5. Agent 1-wise stretching.

need a different algorithm to construct $\tilde{m}$. We propose such an algorithm and use it to characterize minimally informative messages verifying a given choice rule. [21]

It is notationally convenient to identify each message with its content by focusing on *direct* protocols $\langle M, \mu, h\rangle$, in which $M \subset 2^{\mathcal{R}}$ and $\mu^{-1}(m) = m$ for all $m \in M$. A *direct message* is a message in a direct protocol, and by *Privacy Preservation* it must be a product set $m_1 \times \cdots \times m_N \subset \mathcal{R}_1 \times \cdots \times \mathcal{R}_N$. Direct message m is more informative than direct message $\tilde{m}$ if and only if $m \subset \tilde{m}$. Direct message m verifies alternative x if $m \subset F^{-1}(x)$.

Definition 10. For $i \in N$, $x \in X$, the agent i-wise x-stretch of a direct message $m \subset \mathcal{R}$ is the direct message

$$\bigcup_{m'_i \subset \mathcal{R}_i : m'_i \times m_{-i} \subset F^{-1}(x)} m'_i \times m_{-i}.$$

For an illustration with $N = 2$ agents, consider Fig. 5, where direct message m' is the agent 1-wise x-stretch of direct message m.

Lemma 2. (a) *Any direct message* [22] $m \in 2^{\mathcal{R}} \setminus \{\emptyset\}$ *verifying alternative x verifies a minimally informative message verifying x, which can be constructed by sequentially agent i-wise x-stretching message m, for* $i = 1, \ldots, N$.

(b) *A direct message* $m \in 2^{\mathcal{R}} \setminus \{\emptyset\}$ *is a minimally informative message verifying alternative x if and only if it is invariant to any agent-wise x-stretching.*

[21] The same approach to constructing minimally informative messages is independently proposed by Hurwicz and Reiter [22], who call it the “rectangle method”. However, our application of the algorithm to the special case of intersection-monotonic choice rules allows to restrict attention to budget equilibrium messages, and interpret the stretching of such messages as the shrinking of the agents’ budget set.

[22] The most informative direct message $m = \emptyset$ is never accepted and so it is not useful for realization.

Proof. (a) Let $m^0 = m$, and for each $i = 1, \ldots, N$, let message m^i be the agent i-wise x-stretch of message m^{i-1}. Note that $m^i = m_1^N \times \cdots \times m_i^N \times m_{i+1} \times \cdots \times m_N$ for all $i \in N$.

By construction, $m^i \subset F^{-1}(x)$ for any $i = 0, \ldots, N$. This in turn implies that by construction, $m^i \supset m^{i-1}$ for all $i \in N$, and therefore $m^N \supset m^0 = m$, i.e., m verifies m^N.

Suppose now that $m^N \subset \hat{m}_1 \times \cdots \times \hat{m}_N \subset F^{-1}(x)$. Then for any $i \in N$,

$$\hat{m}_i \times m_{-i}^{i-1} \subset \hat{m}_i \times m_{-i}^N \subset \hat{m} \subset F^{-1}(x),$$

and therefore by construction, $m_i^N = m_i^i \supset \hat{m}_i$. Hence, $m^N = \hat{m}$, and therefore m^N is a minimally informative message verifying x.

(b) “Only if” holds by the definition of a minimally informative message. “If” follows from part (a), since sequential agent-wise x stretching of m yields m itself. □

Under BERP, any minimally informative message verifying x verifies, and is thus equivalent to, a budget equilibrium message verifying x. We would like to characterize the verifying budget equilibria that are minimally informative. First note that different budget equilibria may be informationally equivalent. For example, in exchange economies with monotone preferences, a Walrasian budget equilibrium, in which the budget sets are half-spaces, is equivalent to the one in which the half-spaces are replaced with their boundary hyperplanes (i.e., waste is not allowed). It is convenient to focus on the largest equivalent budget equilibria[23]:

Lemma 3. *The largest budget equilibrium* $\left(\hat{B}, x\right)$ *equivalent to a given budget equilibrium* (B, x) *exists and has the budget sets*

$$\hat{B}_i = \bigcap_{R_i \in \mathcal{R}_i : B_i \subset L(x, R_i)} L(x, R_i) \quad \forall i \in N.$$

Proof. Budget equilibrium $\left(\hat{B}, x\right)$ satisfies the following two properties by construction: (i) it is less informative than budget equilibrium (B, x), and (ii) it is larger than any budget equilibrium (B', x) $(B' \in 2^{XN})$ that is equivalent to (B, x). (ii) implies that $\left(\hat{B}, x\right)$ is more informative than (B, x), which, together with (i), implies that $\left(\hat{B}, x\right)$ is equivalent to (B, x). Then (ii) implies the statement of the lemma. □

Lemma 3 allows us to focus on the largest equivalent budget equilibria, which we do from now on. The lemma also implies some useful properties of such budget equilibria in specific settings. In particular, when all feasible lower contour sets satisfy a property that is invariant to set intersections, the largest equivalent budget sets must also satisfy this property. Examples of such properties include: (i) free disposal of some good when preferences are monotone in this good, (ii) closedness in some good when preferences are continuous in this good, (iii) budget sets take the “private” form $B_i = \tilde{B}_i \times X_{-i}$ when the alternative space is $X = X_1 \times \cdots X_N$ and agent i’s preferences over allocations $(x_1, \ldots, x_N) \in X$ depend only on his own allocation x_i.

[23] One reason for this focus is that, as shown below, such an equilibrium always exists (in contrast to, say, a smallest equivalent budget equilibrium). One might also argue on normative grounds for giving agents as much freedom as possible while sustaining the socially desirable alternative.

For realizing an intersection-monotonic choice rule, Lemmas 2 and 3 together with BERP allow to restrict attention to the largest budget equilibria that are minimally informative verifying messages. The lemmas also allow to characterize such budget equilibria: namely, by BERP, in agent-wise stretching we can restrict attention to the largest equivalent budget equilibria verifying a given alternative x. Then agent-wise stretching corresponds to shrinking the agent's budget set by intersecting all of his lower contour sets for which x is still verified given the revealed information about the other agents' preferences. Formalizing this intuition yields a characterization of minimally informative verifying budget equilibria. To simplify the characterization, we assume that the choice rule is extendable to an IM rule on the universal domain $\mathcal{P}^N$, which allows us to write any budget set as $L(x, R_i)$ for some $R_i \in \mathcal{P}^N$ (the assumption is true, in particular, for CU choice rules).

Theorem 3. *Suppose that F is an intersection-monotonic choice rule defined on $\mathcal{P}^N$, and the preference domain is $\mathcal{R} \subset \mathcal{P}^N$. Then*

(a) *Budget equilibrium $(B, x) \in 2^{XN} \times X$ is a largest minimally informative budget equilibrium verifying alternative $x \in X$ if and only if for some $R \in \mathcal{P}^N$,*

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

(b) *If (1) holds for $R \in \mathcal{R}$, then (B, x) is a* unique *largest equivalent budget equilibrium verifying alternative x in state R.*

Proof. (a) A largest equivalent budget equilibrium (B, x) must have $x \in B_i \ \forall i \in N$, hence we can write $(B, x) = (L(x, R_1), \ldots, L(x, R_N), x)$ for some $R \in \mathcal{P}^N$. Lemma 3 and the intersection monotonicity of F on $\mathcal{P}^N$ imply that any largest equivalent budget equilibrium of this form that verifies x must have $x \in F(R)$ (and by monotonicity of F, any such budget equilibrium with $x \in F(R)$ verifies x). Thus, we can restrict attention to such budget equilibrium messages. By the same token, in agent i-wise stretching of such a message, we can restrict attention to budget equilibria $\left(L\left(x, \tilde{R}_i\right), B_{-i}, x\right)$ for $\tilde{R}_i \in \mathcal{P}$ such that $x \in F\left(\tilde{R}_i, R_{-i}\right)$. Thus, the stretching includes all preferences $R_i' \in \mathcal{R}_i$ such that $x \in F\left(\tilde{R}_i, R_{-i}\right)$ for some $\tilde{R}_i \in \mathcal{P}$ satisfying $L\left(x, \tilde{R}_i\right) \subset L\left(x, R_i'\right)$, which by the monotonicity of F is equivalent to $x \in F\left(R_i', R_{-i}\right)$. By Lemma 3, (B, x) is a largest equivalent equilibrium invariant to such stretching if and only if (1) holds.

(b) As noted in the proof of part (a), any largest equivalent budget equilibrium verifying x takes the form $\left(L\left(x, R_1'\right), \ldots, L\left(x, R_N'\right), x\right)$ for some $R' \in \mathcal{P}^N$ such that $x \in F\left(R'\right)$. If it is an equilibrium in state R, then $L\left(x, R_i'\right) \subset L(x, R_i)$ for all i. By monotonicity of F, this implies $x \in F\left(R_i', R_{-i}\right)$ for each i. But then by (1) we have $L(x, R_i) \subset L\left(x, R_i'\right)$, and therefore $L(x, R_i) = L\left(x, R_i'\right)$. $\square$

In words, Theorem 3(a) establishes that the largest minimally informative budget equilibria are those in which each agent's budget set is the intersection of all his feasible lower contour sets for which x is desirable given the information about the others' preferences. Furthermore, Theorem 3(b) says that if the budget sets in such an equilibrium happen to coincide with the lower contour sets in some feasible state R, then it is a *unique* (up to equivalence) budget equilibrium verifying alternative x in state R.

Intuitively, intersection monotonicity implies that alternative x is desirable when it is high enough in the agents' preference rankings. Then (1) means that x is so low in the preference rankings that any further drop in any agent's preferences would render it undesirable. In other words, (1) describes the "boundary" of the states in which x is desirable, and this boundary describes a trade-off between the ranking of x in different agents' preferences. In any state R satisfying (1), there is a unique (up to equivalence) budget equilibrium verifying x, whose budget sets are the agents' lower contour sets at R. By BERP, this budget equilibrium must be a unique (up to equivalence) minimally informative message verifying x in state R.

Finally, observe that if (1) holds in state $R \in \mathcal{R}$, then it also holds when the domain $\mathcal{R}$ is replaced with a smaller domain $\widetilde{\mathcal{R}} \subset \mathcal{R}$ such that $R \in \widetilde{\mathcal{R}}$. Thus, (B, x) remains a unique largest equivalent budget equilibrium verifying alternative x in state R on domain $\widetilde{\mathcal{R}}$. This observation can be used to identify some minimally informative budget equilibria on a reduced domain.

6. Implications for the communication cost

This section discusses the implications of our characterization of minimally informative messages for the communication cost of intersection-monotonic choice rules. The (deterministic/nondeterministic) communication cost of a choice rule is defined as the minimal communication cost of a (deterministic/nondeterministic) protocol realizing it. In this paper, we focus on the traditional measures of communication cost as the length of the realized message sequence, i.e., the number of messages sent in the course of the protocol. Since this number may differ across states, here we focus on the "worst-case" communication cost—the maximum length of the message sequence over all states. For this measure to be interesting, the amount of information conveyed with each message must be bounded, so that all messages are encoded with "elementary" messages.

The computer science literature on "communication complexity" [29] considers discrete communication, and elementary messages that are binary, i.e., convey a bit of information.[24] The nondeterministic communication cost is then the number of bits needed to encode the oracle's message from set M, which is $\log_2 |M|$. In the economic literature on continuous communication, the elementary messages are real-valued. The nondeterministic communication cost is then identified with the number of real numbers needed to encode the oracle's message from space M, i.e., the dimension of M. The discrete and continuous cases have some similarities and some differences, so we discuss them in turn.

6.1. Discrete communication

Starting with any protocol realizing F, we can replace any message verifying alternative x with a less informative minimally informative message verifying x. Doing such replacement for all messages, we obtain a new protocol realizing F using the same number of message, but which uses only minimally informative verifying messages. Thus, in minimizing the communication cost, we can restrict attention to protocols that use minimally informative verifying messages, which are exactly the budget equilibrium messages characterized in Theorem 3(a).

[24] This is just a normalization, because an elementary message from any other finite set (alphabet) could be recoded with a fixed number of bits.

This observation allows us to bound above the nondeterministic communication cost of F by counting all the budget equilibria of the form (1) and taking the binary logarithm. However, we are more interested in having a *lower* bound on the nondeterministic communication cost of F, which would then also serve as a lower bound on the *deterministic* cost of F. Such a lower bound can be obtained using Theorem 3(b), which says that any budget equilibrium of the form (1) for some state $R \in \mathcal{R}$ and alternative $x \in F(R)$ is indispensable for verifying alternative x in state R. However, realization (as opposed to full realization) only requires to verify *one* desirable alternative in any state R. Thus, F may be realized using only a subset the budget equilibria of the form (1).

Nevertheless, in applications considered below, the nondeterministic communication cost of realization is shown to be not much smaller than that of full realization. In some applications, a good lower bound on the nondeterministic cost of realization is obtained by counting only the budget equilibria of the form (1) with states $R \in \mathcal{R}$ in which $F(R)$ is single-valued (and so by Theorem 3(b), each such budget equilibrium is indispensable for realization). In other applications, in which single-valuedness of $F(R)$ cannot be ensured, the following technique proves useful: say that $\mathcal{R}^{\mathrm{f}} \subset \mathcal{R}$ is a *k-degree fooling set* for choice rule F if at most k distinct states from $\mathcal{R}^{\mathrm{f}}$ can share a message verifying an alternative in F. Then the cardinality of the message space in any protocol realizing F is bounded below by $\left|\mathcal{R}^{\mathrm{f}}\right|/k$, and the communication cost of F is bounded below by the binary logarithm of this number. [25] This paper's results allow to show that $\mathcal{R}^{\mathrm{f}}$ is a k-degree fooling set by showing that at most k distinct states from $\mathcal{R}^{\mathrm{f}}$ can share a budget equilibrium of the form (1).

6.2. Continuous communication

The study of continuous communication requires a metric $\rho_{\mathcal{R}}$ on the state space $\mathcal{R}$. Following a suggestion of Debreu [8], we use the Hausdorff metric on the agents' preference relations derived from a given metric ρ_X on the underlying alternative space X. [26]

We would like to define the continuous communication cost as the (worst-case) number of real-valued elementary messages sent in the course of the protocol. We also want to allow finite-valued messages, e.g., to announce discrete allocations, but not count such messages towards the communication cost. In the nondeterministic case, we can identify the communication cost with the dimension of the oracle's message space M.

A well-known problem in measuring continuous communication is the possibility of "smuggling" multi-dimensional information in a one-dimensional message space (e.g., using the inverse Peano function). Note, however that with such smuggling, an arbitrarily small error in the message could yield a large error in its meaning. This suggests that smuggling is prevented when the topology on the message space must be based on their meaning rather than chosen *ad hoc*. Thus, we define the distance between two messages m and m' in protocol $\Gamma = \langle M, \mu, h\rangle$ as the Hausdorff

[25] This is known as the "rectangle-counting" method in the computer science literature [29]. In the case of $k = 1$, $\mathcal{R}^{\mathrm{f}}$ is simply called a "fooling set" in the computer science literature, and "a set with the uniqueness property" in the economic literature.

[26] Formally, $\rho_{\mathcal{R}}\left(R, R'\right) = \max_{i \in N} \max\left\{d_{\mathcal{R}}\left(R_i, R_i'\right), d_{\mathcal{R}}\left(R_i', R_i\right)\right\}$,
with $d_{\mathcal{R}}\left(R_i, R_i'\right) = \sup_{x,y \in X: x R_i y} \inf_{x',y' \in X: x' R_i' y'} \left[\rho_X\left(x, x'\right) + \rho_X\left(y, y'\right)\right]$, where ρ_X is the given metric on X.

distance between the events $\mu^{-1}(m)$ and $\mu^{-1}(m')$ in which they occur. Formally,

$$\rho_M(m, m') = \max\left\{d_M\left(\mu^{-1}(m), \mu^{-1}(m')\right), d_M\left(\mu^{-1}(m'), \mu^{-1}(m)\right)\right\} \quad \text{where}$$

$$d_M(A, B) = \sup_{R \in A} \inf_{R' \in B} \rho_{\mathcal{R}}(R, R') \quad \text{for } A, B \subset \mathcal{R}.$$

Given this metric ρ_M, we use the Hausdorff dimension of M (e.g., [11]) as the measure of continuous communication cost.[27, 28] With this definition of dim M, if messages are coded with d real numbers with a coding whose inverse is Lipschitz continuous (so that small errors in the transmission of the code do not result in large distortion of the state), then we must use $d \geqslant \dim M$ real variables [11, Exercise 6.1.9(1)]. Also, if M is metrically equivalent to a set in $\mathbb{R}^d$ that contains an open set, we must have $d = \dim M$ [11, Exercise 6.2.6]. Thus, the proposed dimensionality measure of M is the relevant measure of communication cost if the communication must be robust to using a channel that is subject to small errors, due either to analog noise or to discretization ("quantization").[29]

Thus defined continuous communication cost can be bounded above using a fooling set technique:

Definition 11. $\mathcal{R}^{\mathrm{f}} \subset \mathcal{R}$ is a fooling set for choice rule F if $\exists C > 0$ such that $\forall R, R' \in \mathcal{R}^{\mathrm{f}}$ and any direct message m verifying any alternative in state R we have

$$\inf_{R'' \in m} \rho_{\mathcal{R}}\left(R'', R'\right) \geqslant C \rho_{\mathcal{R}}(R, R').$$

This definition strengthens the (1-degree) fooling set defined in the previous subsection. (The two definitions coincide when the state space $\mathcal{R}$ is finite, since we can then take $C = \frac{\min_{R, R' \in \mathcal{R}: R' \neq R} \rho_{\mathcal{R}}(R, R')}{\max_{R, R' \in \mathcal{R}} \rho_{\mathcal{R}}(R, R')} > 0$.)

Lemma 4. *If* $\mathcal{R}^{\mathrm{f}}$ *is a fooling set for choice rule* F, *then the continuous communication cost of* F *is at least* dim $\mathcal{R}^{\mathrm{f}}$.

Proof. Take any protocol $\Gamma = \langle M, \mu, h \rangle$, and any selection γ from the message correspondence μ on domain $\mathcal{R}^{\mathrm{f}}$. We must have $\forall R, R' \in \mathcal{R}^{\mathrm{f}}$,

$$\rho_M(\gamma(R), \gamma(R')) \geqslant \inf_{R'' \in \gamma(R)} \rho_{\mathcal{R}}\left(R'', R'\right) \geqslant C \rho_{\mathcal{R}}(R, R'),$$

[27] Alternatively, we could use other metric dimension measures of M, such as the box dimension or the packing index. In most practical cases, the different dimensions would coincide, provided that M is bounded.

[28] This definition of the continuous communication burden stands in contrast to the existing economic literature on message space dimension, in which the message space comes endowed with a Hausdorff topology, its dimension is defined in a topological way, and a "regularity" restriction is imposed on the communication protocol to prevent dimension smuggling. The typical regularity restriction is that the message correspondence μ be "locally threaded"—i.e., have a continuous selection on a neighborhood of any point [38]. This restriction rules out *a priori* some useful communication protocols: for example, in problems with continuous preferences and discrete (e.g., combinatorial) allocations, it prevents the communication of discrete allocations (any selection from μ is discontinuous at a point at which the optimal discrete allocation switches).

[29] A formal result about robust discretization is stated in [40, Proposition 4].

where the first inequality is by definition of metric ρ_M as the Hausdorff metric, and the second inequality is because $\gamma(R)$ verifies an alternative in state R and the definition of a fooling set. Therefore, $\gamma : \mathcal{R}^{\mathrm{f}} \to M$ has a Lipschitz continuous inverse, hence $\dim M \geqslant \dim \mathcal{R}^{\mathrm{f}}$ [11, Exercise 6.1.9(1)]. □

Note that it suffices to check Definition 11 only for *minimally informative* verifying messages m, since for them the inequality is the least likely to hold. Thus, just as for discrete communication, characterization (1) of minimally informative verifying messages (budget equilibria) facilitates the calculation of the continuous communication cost for intersection-monotonic choice rules.

7. Applications

7.1. Pareto efficiency in convex economies

Consider *smooth convex exchange economies*, in which the alternatives represent the consumption of L divisible goods by the N agents, hence $X = \mathbb{R}_+^{NL}$. Each agent i's preference domain consists of the convex preferences described by differentiable utility functions of his own consumption $x_i \in \mathbb{R}_+^L$ with a nonnegative nonzero gradient everywhere. The feasible set consists of allocations of a given aggregate endowment $\bar{x} \in \mathbb{R}_{++}^L$: $\bar{X} = \left\{x \in X : \sum_i x_i = \bar{x}\right\}$. Recall that the Pareto rule is described by

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

We use the stretching algorithm described in Section 5 to derive minimally informative messages verifying the Pareto efficiency of an allocation $x \in \bar{X}$ with $x \gg 0$.[30] The derivation can be illustrated in the standard Edgeworth box depicted in Fig. 6. Start with a state R in which x is Pareto efficient, which means that agent 1's indifference curve passing through x is below agent 2's indifference curve passing through x. Note that given smoothness, the two curves must be tangent at x, and let p denote the agents' common marginal rate of substitution at x. Now, for agent 1-wise stretching, we shrink agent 1's lower contour set as much as possible while preserving the Pareto efficiency of x and keeping agent 1's preferences convex. This shrinking is illustrated with the left-down arrows in the Figure. The furthest we can shrink agent 1's lower contour set is to that of linear preferences—a hyperspace with gradient p. This yields a Walrasian budget set for agent 1 described by the commodity price vector p. Next, for agent 2-wise stretching, we shrink agent 2's lower contour set as illustrated with the right-up arrows, yielding for him a Walrasian budget set with the same commodity price vector p. Thus the stretching algorithm yields a Walrasian equilibrium. Furthermore, any Walrasian equilibrium is invariant to stretching—i.e., satisfies (1). A formalization of this argument yields[31]

[30] We restrict attention to $x \gg 0$ to avoid the problem of nonexistence of supporting Walrasian prices [32, Figure 16.D.2].

[31] If nonsmooth preferences are allowed, the Walrasian equilibria still satisfy (1), but other minimally informative messages verifying Pareto efficiency emerge. For example, let $N = L = 2$ and $\bar{x} = (2, 2)$, and consider the budget equilibrium (B_1, B_2, x) with $x = (1, 1, 1, 1)$, $B_1 = \{x \in X : \min\{x_{11}, x_{12}\} \leqslant 1\}$, and $B_2 = \{x \in X : x_{21}, x_{22} \leqslant 1\}$. This is a budget equilibrium in state $R \in \mathcal{R}$ if and only if $L(x, R_1) = B_1$. This is a minimally informative message verifying the Pareto efficiency of x, but it is not equivalent to a Walrasian equilibrium.

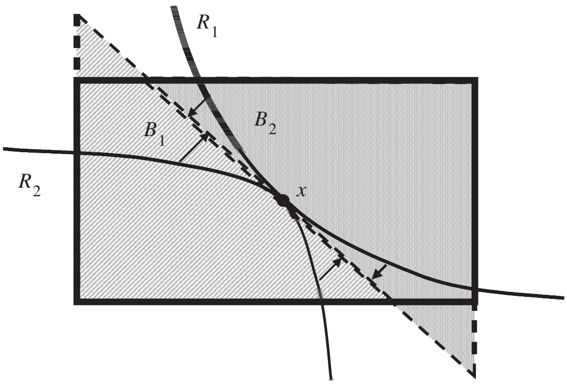

Fig. 6. The stretching algorithm in convex economies.

Proposition 1. *A message is a minimally informative message verifying the Pareto efficiency of allocation $x \in \bar{X}$ with $x \gg 0$ in a smooth convex exchange economy if and only it is equivalent to a* Walrasian equilibrium *supporting x, i.e., a budget equilibrium (B, x) with*

$$B_i = \{y \in X : p \cdot y_i \leqslant p \cdot x_i\} \quad \forall i \in N \tag{2}$$

for some price vector $p \in \mathbb{R}^L_+$ such that $\|p\| = 1$. Any such equilibrium is a unique Walrasian equilibrium supporting allocation x in any state in which it is an equilibrium.[32]

Proof. (B, x) verifies the Pareto efficiency of x if and only if the normalized gradients of *all* agents' utility functions at x in *all* states in $E^{-1}(B, x)$ equal some $p \in \mathbb{R}^L_+$. By Lemmas 2 and 3, (B, x) is a largest minimally informative budget equilibrium verifying x if and only if for each $i \in N$, B_i is the intersection of all lower contour sets at x of agent i's utility functions with gradient p at x. This means that B_i is given by (2). Furthermore, in any state in which such (B, x) is an equilibrium, the normalized gradients of all agents' utilities at x equal p, which implies that in this state (B, x) is a unique Walrasian equilibrium supporting x. □

The proposition implies that the minimal message space required for verifying any Pareto efficient allocation in any convex economy is the space of Walrasian equilibria. Since a feasible allocation $x \in \bar{X}$ is described with $(N-1)L$ real variables, and a normalized price vector p is described with $L-1$ real variables, the space of Walrasian equilibria has dimension $(L-1) + (N-1)L = NL - 1$.

However, *realizing* Pareto efficiency only requires to verify *one* efficient allocation in each state. In fact it is possible to realize the Pareto rule without any communication, e.g., by giving all the aggregate endowment to agent 1. To rule this out, we restrict attention to allocations satisfying a "subsistence" requirement that $\|x\| \geqslant \sigma$, for a given $\sigma < \frac{1}{N} \min_l \bar{x}_l$.[33] Note that the subsistence

[32] Note that the last statement is stronger than that in Theorem 3(b): in this particular setting, the minimally informative messages verifying *x partition* $F^{-1}(x)$.

[33] The "informational efficiency" literature only ruled out the corners of the feasible set $\overline{X}$, but we need to rule out neighboring allocations as well, because we do not impose any "regularity" restriction on protocols and use a metric measure of dimensionality. Intuitively, if only the corners of $\overline{X}$ were ruled out, Pareto efficiency could still be approximated arbitrarily closely without any communication, by giving nearly all the aggregate endowment $\bar{x}$ to one agent.

Pareto rule can be realized by fixing an "endowment allocation" $\omega \in \bar{X}$ with $\omega \geqslant (\sigma, \ldots, \sigma)$ and announcing a Walrasian equilibrium (B, x) such that $\omega \ni B_i$ for all i, which exists in any convex economy [32, Section 17.BB]. Since such equilibria satisfy the additional "budget constraints" $\sum_l p_l \omega_{il} = \sum_l p_i x_{il}$ for all i, they can be communicated using $(L-1) + (N-1)(L-1) = N(L-1)$ real numbers.

It is in fact impossible to realize subsistence Pareto efficiency using less communication. This can be shown using the fooling set consisting of the *Cobb–Douglas economies*, in which each agent i's preferences are described by a utility function of the form $u_i(x_i) = \prod_l x_{il}^{\alpha_{il}}$ with the normalization $\sum_l \alpha_{il} = 1$. Indeed, all subsistence Pareto efficient allocations in a Cobb–Douglas economy with parameters $\alpha \gg 0$ are interior, and the first-order equilibrium conditions imply that no two distinct Cobb–Douglas economies share a Walrasian equilibrium sustaining an interior allocation. [34] Therefore, we must use a subspace of Walrasian equilibria whose dimension is at least that of Cobb–Douglas economies, which is $N(L-1)$:

Corollary 1. *The nondeterministic communication cost of subsistence Pareto efficiency in the convex exchange economy is exactly* $N(L-1)$ *real numbers, and it is achieved by the Walrasian equilibrium protocol with a fixed endowment.*

This result was first obtained by the "informational efficiency" literature [20,38]. Unlike this literature, we have derived it from the purely set-theoretic Proposition 1, which does not require any topological restrictions on communication or any scalar measure of the communication cost.

7.2. *Pareto efficiency in economies with numeraire*

Consider *economies with numeraire*, in which the set of alternatives takes over the form $X = K \times \mathbb{R}^N$, where K is a finite set of *(nonmonetary) allocations*, and $\mathbb{R}^N$ describes the *transfers* of numeraire (money) to the agents. Each agent i's preference domain $\mathcal{R}_i$ consists of preference relations over $(k, t) \in X$ that are (i) independent of other agents' transfers t_{-i}, (ii) continuous and strictly increasing in his own transfer t_i, and (iii) allow compensation— i.e., for any $x \in X$ and any $k \in K$ there exists $t \in \mathbb{R}^N$ such that (k, t) is indifferent to x. The feasible set takes the form $\bar{X} = \{(k, t) \in X : \sum_i t_i = 0\}$, i.e., requires a balanced budget. We consider the problem of finding a Pareto efficient allocation within $\bar{X}$.

An important subclass of preference relations satisfying (i)–(iii) consists of *quasilinear* preferences, which are described by utility functions of the form $u_i(k, t) = v_i(k) + t_i$. For such preferences, Pareto efficiency is equivalent to maximizing the total surplus $\sum_i v_i(k)$.

We use the stretching algorithm of Section 5 to derive minimally informative messages verifying Pareto efficiency. We illustrate this algorithm in an Edgeworth box depicted in Fig. 7, in which the vertical dimension represents allocations of numeraire between the agents, and the horizontal dimension represents the nonmonetary allocations $k \in K$ (arranged in no particular order). Start with a state R in which x is Pareto efficient, which means that the indifference curve of agent 1 passing through x is above the indifference curve of agent 2 passing through x. For agent 1-wise stretching, shrink the lower contour set of agent 1 as much as possible while preserving the Pareto efficiency of x (as illustrated with the downward arrows in the figure). The furthest we

[34] Furthermore, we can also show that Definition 11 holds: the minimal distance between a Cobb–Douglas economy with parameters α and any economy that shares a subsistence Walrasian equilibrium with the Cobb–Douglas economy with parameters α' is at least $C\|\alpha - \alpha'\|$, provided that $\alpha, \alpha' \geqslant (\delta, .., \delta)$ for a fixed $\delta > 0$.

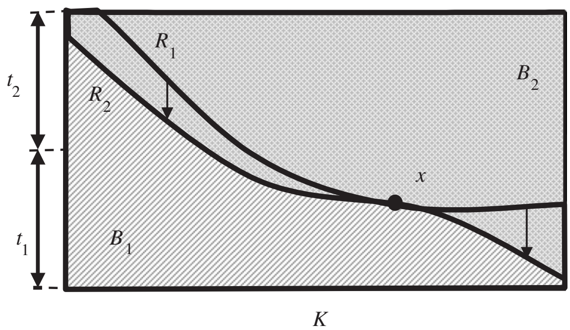

Fig. 7. The stretching algorithm with nonconvexities and numeraire.

can shrink it is until agent 2's indifference curve (unlike in the previous subsection, there is no convexity restriction to hold us back). Once this shrinking is completed, agent 2–wise stretching is impossible—agent 2's lower contour set cannot be shrunk without violating the Pareto efficiency of x. The obtained budget sets for the two agents can be delineated by general nonlinear and personalized prices $p_i(k)$ $(i = 1, 2, k \in K)$, specifying the cost of allocation k to agent i in terms of numeraire. The fact that the two budget sets' boundaries coincide means that the sum of the prices, $p_1(k) + p_2(k)$, must be the same for all allocations $k \in K$. The budget equilibria described in this way are the only budget equilibria that are invariant to the budget-shrinking procedure, i.e., satisfy (1). The argument extends to any number of agents, yielding the following result:

Proposition 2. *A message is a minimally informative message verifying the Pareto efficiency of allocation* $x = (k, t) \in \bar{X}$ *in an economy with numeraire if and only if it is equivalent to a* valuation equilibrium *supporting* x*, i.e., a budget equilibrium* (B, x) *in which for each* $i \in N$,

$$B_i = \left\{\left(k', t'\right) \in X : p_i\left(k'\right) + t'_i \leqslant p_i(k) + t_i\right\} \quad \textit{for some } p_i \in \mathbb{R}^K, \tag{3}$$

and such that

$$\sum_i p_i\left(k'\right) = \sum_i p_i(k) \quad \textit{for all } k' \in K. \tag{4}$$

Any such equilibrium is a unique valuation equilibrium supporting allocation x in the states R in which $L(x, R_i) = B_i$ *for all i.*

Proof. Observe first that for each agent i and any $B_i \subset X$, $B_i = L(x, R_i)$ for some $R_i \in \mathcal{R}$ if and only if B_i takes the form (3). (For the "if" part, take preference relation R_i described by the quasilinear utility function $p_i\left(k'\right) + t_i$; for the "only if" part, take $p_i\left(k'\right)$ for each $k' \in K$ such that $\left(k', -p_i\left(k'\right)\right)$ is indifferent to x in R_i.) Since this form is preserved under set intersection, any budget equilibrium (B, x) satisfying (1) must have budget sets of this form (allowing, possibly, for $p_i(k) = +\infty$ for $k' \neq k$). Furthermore, $x \in F(R)$ if and only if it is impossible to extract numeraire while making all agents equally well off, i.e.,

$$\sum_i p_i\left(k'\right) \leqslant \sum_i p_i(k) \quad \text{for all } k' \in K.$$

(In particular, this implies that $p_i(k) < +\infty\ \forall k' \neq k \quad \forall i \in N$.) (1) means that the prices $p_i\left(k'\right)$ for all $k' \in K \setminus \{k\}$ are maximized subject to the inequality, which yields condition (4). Theorems 2 and 3 imply the proposition. □

The term "valuation equilibrium" was coined by Mas-Colell [31]; such equilibria were also studied by Bikhchandani and Mamer [2] and Bikhchandani and Ostroy [3]. These papers have established classical welfare theorems for such equilibria: an allocation is Pareto efficient if and only if it is supported by a valuation equilibrium. The contribution of Proposition 2 lies is in showing that valuation equilibria constitute the minimal information that must be revealed in order to verify the Pareto efficiency of an allocation.

Proposition 2 implies that the minimal message space required for verifying any efficient allocation in any economy with numeraire is the space of valuation equilibria. Normalizing the prices, e.g., so that $\sum_k p_i(k) = 0$ for each agent i, we can announce a price vector satisfying (4) using $(N-1)(K-1)$ real numbers. In addition, $K-1$ real numbers are needed to announce a transfer vector t adding up to zero (a discrete allocation k is zero-dimensional).

For *realizing* Pareto efficiency, we only need to verify *one* efficient allocation in each state, and so need not use all valuation equilibria. However, it turns out that *all* the possible normalized valuation prices satisfying (4) must still be used, even if we restrict attention to the domain of quasilinear preferences. Indeed, consider *diagonal* economies, in which the agents have quasilinear utility functions $u_i(k,t) = p_i(k) + t_i$ with $p \in \mathbb{R}^{NK}$ satisfying (4). In such an economy, all allocations $x \in \bar{X}$ are surplus-maximizing, but by the second part of Proposition 2, in the valuation equilibrium supporting any such allocation the agents' budget sets must be described by prices p. Thus, no two distinct diagonal economies share a valuation equilibrium, and so diagonal economies form a fooling set.[35] Therefore, realizing Pareto efficiency with quasilinear preferences requires the announcement of an $(N-1)(K-1)$ dimensional price vector. This amount of communication in fact allows a *deterministic* surplus-maximizing protocol, in which the first $N-1$ agents announce their normalized utilities and then the last agent chooses a surplus-maximizing allocation. Thus we have:

Corollary 2. *The communication cost* (*both deterministic and nondeterministic*) *of Pareto efficiency in a quasilinear economy is* $(N-1)(K-1)$ *real numbers.*

One class of quasilinear allocation problems with numeraire that has received a lot of attention recently is the "combinatorial allocation problem," in which there is a set L of objects to be allocated among the agents, thus $K = N^L$, and the preference domain includes those quasilinear preferences in which each agent i cares only about his own consumption bundle $k^{-1}(i)$ and his preference is monotonic in this bundle (in the set inclusion order). Consider the particular case of $N = 2$, and note that for any normalized price vector $p \in \mathbb{R}^{NK}$ satisfying (4) such that $p_1(k)$ is nondecreasing in $k^{-1}(1)$, we also have that $p_2(k)$ is nondecreasing in $k^{-1}(2)$. In the state in which the agents' preferences are described by utility functions $u_i(k,t) = p_i(k) + t_i$ $(i = 1, 2)$ for such prices, all allocations $x \in \bar{X}$ are surplus-maximizing by (4), but the normalized price vector in any valuation equilibrium must coincide with p by the second part of Proposition 2. Thus, any normalized monotonic price vector for an agent must be announced by an efficient protocol, which implies:

[35] Formally, to apply Lemma 4, we need diagonal economies to satisfy the stronger Definition 11 of a fooling set, which is shown in [40, Proposition 2].

Corollary 3. *The continuous communication cost (both deterministic and nondeterministic) of efficient combinatorial allocation of L objects between two agents is* $2^L - 1$.

To see that the deterministic communication cost coincides with the nondeterministic cost, consider the communication protocol in which agent 1 announces its utility function and then agent 2 chooses an efficient allocation. Corollary 3 is obtained by Nisan and Segal [40], who also examine the potential communication savings when agents' combinatorial valuations are *a priori* restricted to lie in certain important classes.

7.3. Approximate Pareto efficiency in economies with numeraire

Recall that the approximate Pareto rule is defined by

$$F(R) = \left\{x \in \bar{X} : X^\delta \subset \cup_{i \in N} L(x, R_i)\right\} \quad \forall R \in \mathcal{R},$$

where $X^\delta \subset \bar{X}$ denotes the set of alternatives in which at least amount $\delta > 0$ of resources is wasted. We consider the domain $\mathcal{R}$ with numeraire defined in the previous subsection, and let X^δ be the set of alternatives that waste at least amount δ of numeraire: $X^\delta = \left\{(k, t) : \sum_i t_i \leqslant -\delta\right\}$.

Proposition 3. *A message is a minimally informative message verifying δ-approximate Pareto efficiency of allocation* $x = (k, t) \in \bar{X}$ *in an economy with numeraire if and only if it is equivalent to a δ-valuation equilibrium supporting x, i.e., a budget equilibrium* (B, x) *with budget sets described by* (3) *with the price vector* $p \in \mathbb{R}^{NK}$ *satisfying*

$$\sum_i p_i\left(k'\right) = \sum_i p_i(k) + \delta \quad \textit{for all } k' \in K \setminus \{k\}. \tag{5}$$

Any such equilibrium is a unique δ-valuation equilibrium in the states R in which $L(x, R_i) = B_i$ *for all i.*

Proof. Recall from the proof of Proposition 2 that for each agent i and any $B_i \subset X$, $B_i = L(x, R_i)$ for some $R_i \in \mathcal{R}$ if and only if B_i takes the form (3), which implies that any budget equilibrium satisfying (1) takes this form (allowing, possibly, for $p_i(k) = +\infty$ for $k' \neq k$). $x \in F(R)$ if and only if it is impossible to extract more than amount δ of the numeraire while making all agents equally well off, i.e.,

$$\sum_i p_i\left(k'\right) \leqslant \sum_i p_i(k) + \delta \quad \text{for all } k' \in K.$$

(In particular, this implies that $p_i(k) < +\infty \; \forall k' \neq k \quad \forall i \in N$.) (1) means that the prices $p_i\left(k'\right)$ for all $k' \in K \setminus \{k\}$ are maximized subject to the inequality, which yields condition (5). Theorems 2 and 3 imply the proposition. □

We now focus on the domain of quasilinear preferences, for which $F(R)$ is the set of alternatives that approximate the maximum surplus in state R within δ. Furthermore, we restrict attention to bounded utility functions: $u_i(k) \in [0, 1]$ for all $k \in K$, $i \in N$. Then letting one agent choose an allocation to maximize his own utility approximates the maximum surplus within $\delta = N - 1$; we examine the communication cost of improving the approximation to some $\delta < N-1$. Observe that any approximation $\delta > 0$ can be achieved with finite communication in which agents announce their utilities discretized to multiples of a sufficiently small $\varepsilon > 0$. Thus, the communication cost of approximation should be measured with the number of bits.

We bound below the number of δ-valuation equilibria needed to ensure equilibrium existence on the subset $\widetilde{\mathcal{R}}$ of states in which for all $k \in K$, $u_i(k) \in \{0, 1\}$ for all i, and $\sum_i u_i(k) = 1$. Observe that $\left|\widetilde{\mathcal{R}}\right| = N^K$, since the value 1 for any allocation $k \in K$ can be assigned to any of the N agents. Now consider how many states from $\widetilde{\mathcal{R}}$ can share a given δ-valuation equilibrium $(B, (k, t))$ described by a price vector $p \in \mathbb{R}^{NK}$. We can assign value 1 for the proposed allocation k to one of the N agents. In all states in which $(B, (k, t))$ is an equilibrium, for any allocation $k' \neq k$, each agent i's utility must satisfy

$$u_i\left(k'\right) \leqslant \gamma_i\left(k, k'\right) \equiv u_i(k) + p_i\left(k'\right) - p_i(k).$$

On the other hand, (5) implies that in any state from $\widetilde{\mathcal{R}}$,

$$\sum_i \gamma_i\left(k, k'\right) = \sum_i u_i(k) + \delta = 1 + \delta < N.$$

Therefore, for some agent i we must have $\gamma_i\left(k, k'\right) < 1$, and so this agent cannot have value 1 for allocation k'. Thus, we are left with at most $N - 1$ possibilities to assign value 1 for allocation k' among the other agents. Since this holds for any $k' \neq k$, a given δ-valuation equilibrium can be an equilibrium in at most $N(N-1)^{K-1}$ states from $\widetilde{\mathcal{R}}$, i.e., $\widetilde{\mathcal{R}}$ is a $N(N-1)^{K-1}$-degree fooling set, as defined in Section 6. Thus, we need to use at least $\frac{\left|\widetilde{\mathcal{R}}\right|}{N(N-1)^{K-1}} = (1 + 1/(N-1))^{K-1}$ such equilibria to ensure equilibrium existence on $\widetilde{\mathcal{R}}$, and the communication cost of F is bounded below by the binary logarithm of this number:

Corollary 4. *When agents have quasilinear utilities in* [0,1], *the communication cost of approximating the maximum surplus within* $\delta < N - 1$ (*i.e., achieving a better approximation than by letting one agent choose an allocation*) *is at least* $(K-1)\log_2(1 + 1/(N-1))$ *bits.*

The corollary reproves Nisan's [39] Theorem 2 on the communication complexity of the "approximate disjointness problem" using BERP. It can also be used to prove Nisan and Segal's [40] result on the communication cost of approximately efficient combinatorial auctions. Namely, they construct a "large" subset K of allocations such that the agents can have arbitrary utilities in [0,1] for allocations from K, and in looking for approximately efficient allocations we can restrict attention to those from K. (The allocations from K correspond to partitions of objects with the "pairwise intersection" property.) Corollary 4 implies that achieving a better approximation than giving all objects to one agent requires communication proportional to $|K|$, which proves to be exponential in the number of objects.

7.4. Individually rational Pareto efficiency with universal preferences and in discrete economies

Let us require individual rationality along with Pareto efficiency, with $x^0 \in X$ being the status-quo alternative. Formally, F is defined by

$$F(R) = \left\{x \in X : x^0 \in L(x, R_i)\, \forall i \in N, X = \cup_{i \in N} L(x, R_i)\right\} \quad \forall R \in \mathcal{R}.$$

Let X be a finite set, which ensures that this choice rule is nonempty-valued (e.g., it includes agent 1's preferred alternative from those that are individually rational for the other agents). Consider first the universal domain:

Proposition 4. *A message is a minimally informative message verifying the Individually Rationality and Pareto efficiency of alternative $x \in X$ on the universal domain $\mathcal{R} = \mathcal{P}^N$ if and only if it is equivalent to a partitional equilibrium supporting x, i.e., a budget equilibrium (B, x) in which $x, x^0 \in B_i$ for all $i \in N$, and $(B_1, \ldots, B_N)$ forms a partition of $X \setminus \{x, x^0\}$. Furthermore, any such equilibrium is a unique* partitional equilibrium *supporting alternative x in any state $R \in \mathcal{P}^N$ in which $L(x, R_i) = B_i$ for all $i \in N$.*

Proof. (1) means that for each $i \in N$,

$$B_i = \bigcap_{R_i' \in \mathcal{P}: x \in F(R_i', R_{-i})} L\left(x, R_i'\right) = \bigcap_{Y \subset X: x, x^0 \in Y, x^0 \in B_j \forall j \in N \setminus \{i\}, Y \cup \left(\cup_{j \in N \setminus \{i\}} B_j\right) = X} Y.$$

This implies that $x, x^0 \in B_i \; \forall i \in N$, and then holds if and only $B_i = \{x, x^0\} \cup \left(X \setminus \{\cup_{j \in N \setminus \{i\}} B_j\}\right)$ $\forall i \in N$, i.e., (B, x) is a partitional equilibrium. Theorems 2 and 3 imply the proposition. □

Proposition 4 implies that the minimal message space required for verifying *any* Pareto efficient IR alternative with universal preferences is the space of partitional equilibria. *Realization* of the choice rule requires verifying only *one* desirable alternative in each state, which in principle may not require all possible partitional equilibria. However, 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 x is a *unique* desirable alternative. In this state, the status-quo alternative x^0 (if different from x) is the next-best alternative to x in each agent's preference ranking. This ensures that the only alternatives that are individually rational for all agents in state R are x and x^0, and Pareto efficiency dictates that $F(R) = \{x\}$. The second part of Proposition 4 then implies that (B, x) is a unique partitional equilibrium in state R. Hence, all partitional equilibria must be used for realizing the choice rule.

There are N^{X-1} partitional equilibria with $x = x^0$ (each of the alternatives in $X \setminus \{x^0\}$ can be allocated to any of the N agents' budget sets), and N^{X-2} such equilibria for any given $x \neq x^0$ (each of the alternatives in $X \setminus \{x, x^0\}$ can be allocated to any budget set). Adding up, we obtain $N^{X-1} + (X-1) N^{X-2}$ partitional budget equilibria. Taking the binary logarithm, we obtained the number of bits that must be communicated:

Corollary 5. *The nondeterministic communication cost of the individually rational Pareto rule on the universal preference domain is exactly* $(X-2)\log_2 N + \log_2(N + X - 1)$ *bits.*

When X is large, this cost is asymptotically proportional to X, which is exponentially larger than that of simply naming an alternative (which takes $\log_2 X$ bits). In fact, the 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 $X \to \infty$.[36]

One setting where the alternative space X is naturally large is the exchange economy with L indivisible goods, in which $X = N^L$ (note that unlike in the combinatorial allocation problem

[36] Since there are $X!$ *strict* preference orderings of X elements, by Stirling's formula, it takes $\log_2 X! \sim X \log_2 X$ bits to communicate such an ordering as $X \to \infty$. That allowing indifference does not raise the asymptotic communication burden follows from the approximation in [14].

described in Section 7.2, there is no divisible "numeraire" good). Suppose that each agent's preferences depend only on his own consumption of goods and are monotonic in it. While we no longer have a universal preference domain, we can focus on the case where $N = 2$, and on the subset $\tilde{X} \subset X$ of alternatives that give $L/2$ objects to each agent. If the status-quo allocation $x^0 \in \tilde{X}$, and if the agents' preferences are restricted to be such that they always strictly prefer to consume a larger number of objects, then all individually rational allocations must also lie in $\tilde{X}$. Furthermore, the restriction still allows the agents to have arbitrary preferences over $\tilde{X}$. Thus, we can restrict attention to the problem on the set $\tilde{X}$ with universal preferences, and Corollary 5 yields:

Corollary 6. *The communication cost of verifying an individually rational Pareto efficient allocation in an indivisible-good exchange economy with two agents and L objects is at least* $\tilde{X} - 1 = \binom{L}{L/2} - 1$ *bits.*

Thus, the communication cost is exponential in the number of objects. [37]

7.5. Stable many-to-one matching

Let the set N of agents be partitioned into the set F of firms and the set W of workers. A two-sided matching between firms and workers is described by a binary relation $x \subset F \times W$. With a slight abuse of notation, we also let x represent the correspondence $x : N \twoheadrightarrow N$ defined by

$$x(i) = \{j \in N : (i, j) \in x \text{ or } (j, i) \in x\} \quad \text{for } i \in N.$$

We restrict attention to *many-to-one* matching problems, in which a worker cannot match with more than one firm, and so the set of alternatives is

$$X = \{x \subset F \times W : |x(w)| \leqslant 1 \ \forall w \in W\}.$$

We focus on matching problems without externalities, i.e., those in which each agent i's preferences depend only on the set $x(i)$ of his matching partners.

The stable matching rule is a CU rule that is described with the following blocking sets:

$$\beta(x, S) = \{y \in X : y \setminus (S \times S) \subset x \setminus (S \times S)\} \quad \forall S \subset N, \ \forall x \in X.$$

In words, a coalition cannot create new matches involving outsiders, but can break any match and can create any match between its members. [38] This stable matching problem is studied by Roth and Sotomayor [44].

Proposition 5. *A message is a minimally informative message verifying the stability of a many-to-one matching x if and only if it is equivalent to a* match-partitional equilibrium *supporting x*,

[37] The setting can also be reinterpreted as bilateral bargaining over L binary attributes, where it is known that, other things equal, agent 1 prefers value 1 and agent 2 prefer value 0 for any attribute, but otherwise the agents can have arbitrary preferences over attribute profiles. The corollary implies that finding a Pareto efficient and individually rational attribute profile requires exponential communication in the number of attributes.

[38] We might also prevent a coalition from breaking matches between outsiders, but this is irrelevant when externalities in preferences are ruled out.

i.e., a budget equilibrium (B, x) *satisfying*

$$B_f = \{y \in X : y(f) \subset \omega(f)\} \quad \forall f \in F,$$
$$B_w = \{y \in X : y(w) \subset \phi(w)\} \quad \forall w \in W,$$

for some $\phi, \omega \subset F \times W$ *such that* $\phi \cap \omega = x$ *and* $\phi \cup \omega = F \times W$. *Furthermore, any such equilibrium is a unique match-partitional equilibrium supporting matching* x *in any state* $R \in \mathcal{R}$ *in which* $L(x, R_i) = B_i$ *for all* $i \in N$.

Proof. For any agent $i \in N$, $B_i = L(x, R_i)$ for some $R_i \in \mathcal{R}_i$ if and only if

$$B_i = \{y \in X : y(i) \in \Omega_i\}$$

for some $\Omega_i \subset 2^W$ for $i \in F$ or $\Omega_i \subset 2^F$ for $i \in W$. Since this form is preserved under set intersection, any budget equilibrium (B, x) satisfying (1) must take this form. Furthermore, $x \in F(R)$ if and only if

(i) each worker $w \in W$ prefers x to being unmatched, and
(ii) each firm $f \in F$ prefers x to matching with any subset consisting of some workers who strictly prefer f to their equilibrium match and some of those already matched with f.

(i) means that $\emptyset \in \Omega_w$ for each worker $w \in W$. Since the worker can match with at most one firm, and the set of his possible matching partners in B_w is $\phi(w) \equiv \{f \in F : \{f\} \in \Omega_w\}$, B_w is not affected by redefining $\Omega_w = 2^{\phi(w)}$. This allows to write the workers' budget sets in the desired form for some relation $\phi \subset F \times W$. Then (ii) means that for each firm $f \in F$,

$$2^{(W \backslash \phi(f)) \cup x(f)} \subset \Omega_f.$$

(1) means that each budget set B_i is the smallest possible given B_{-i} such that the above inclusion holds. For $i \in F$ (firms), this means that $\Omega_i = 2^{\omega(i)}$ for $\omega(i) = x(i) \cup (W \backslash \phi(i))$, thus the firm's budget sets take the desired form for the relation $\omega \subset F \times W$ such that ω and ϕ partition $(F \times W) \backslash x$. This also ensures the minimality of the budget set B_i of any worker $i \in W$ given B_{-i}. Theorems 2 and 3 imply the proposition. □

Intuitively, since a worker's preferences depend only on his matching partner, his (largest equivalent) budget sets can be described in terms of the available firms. On the other hand, since a firm has preferences over *groups* of workers, its (largest equivalent) budget sets can be described in terms of such available groups. A budget equilibrium with such budget sets verifies stability if and only if each firm *f*'s budget set includes all groups consisting of all the subsets workers who do not have *f* in their budget sets and some of those currently employed by *f*. Indeed, this ensures that no deviation can make firm *f* and all of its new employees strictly better off. Finally, minimally informative budget equilibria have the minimal budget sets necessary for verification; this means that each firm *f*'s budget set must include *exactly* all the subsets of *f*'s current employees and those workers who do not have *f* in their budget set. Thus, in a minimally informative budget equilibrium, the firms' budget sets are implied by the workers' budget sets, and they can be described by listing *individual* workers that are available to the firm rather than groups of workers.

The fact that combinatorial budget sets for firms need not be used brings about an exponential reduction in the communication cost. Indeed, the workers' budget sets are described by a relation $\phi \subset F \times W$, which is communicated with at most FW bits, the equilibrium matching x is communicated with $W \log_2(F+1)$ bits, and the firms' budget sets are implied. Thus, the cost of verifying a many-to-one stable matching is $O(FW)$ as $F, W \to \infty$. This is exponentially smaller than that of full revelation of a firm's preferences over subsets of workers, which asymptotically takes $\log_2\left(2^W!\right) \sim W \cdot 2^W$ bits as $W \to \infty$ (see footnote 36).

For *realizing* the choice rule, we only need to verify *one* stable matching in each state, and need not use all match-partitional equilibria. However, we can show that "almost" all such equilibria need to be used, and so the nondeterministic communication cost of stability is asymptotically *FW* bits. This is true even if the preference domain is restricted to include only preferences that are strict and *one-to-one*, i.e., each firm prefers being unmatched to matching with more than one worker. With such preferences, we can restrict attention to *one-to-one* matchings *x*, in which $|x(i)| \leqslant 1$ for all $i \in N$. We show that with such preferences, the uniqueness of a stable matching can be ensured by adding one agent on each side:

Lemma 5. *In the one-to-one matching problem with strict preferences, for any stable matching x in any state R, we can add a firm f^* and a worker w^* and complete the preferences in a way consistent with R so that $x \cup \{(f^*, w^*)\}$ is the unique stable matching.*

Proof. Let the new agents' preferences have $w R_{f^*} w^* R_{f^*} \{\emptyset\}$ and $f R_{w^*} f^* R_{w^*} \{\emptyset\}$ for all $f \in F, w \in W$, i.e., each new agent prefers all other partners to the other new agent, which he in turn prefers to being single. For the old agents, let every firm $f \in F$ rank w^* just below its current match $x(f)$, and let every worker $w \in W$ rank f^* just below his current match $x(w)$. Such completion of preferences guarantees that matching $x^* = x \cup \{(f^*, w^*)\}$ is stable. We show that x^* is a *unique* stable matching by contradiction: If it were not, then by the Lattice Theorem [44, Theorem 2.16], either the worker-pessimal stable matching x^w or the firm-pessimal stable matching x^f would differ from x^*. For definiteness let $x^w \neq x^*$. By [44, Theorem 2.22], the set of single agents is the same in x^w as in x^*. Therefore, worker w^* must still be matched in x^w, and since cannot be better off in than in x^*, we must have $x^w(w^*) = f^*$. But this implies that any worker $w \neq w^*$ who is strictly worse off in x^w than in x^* would have a strictly Pareto improving blocking by matching with firm f^*. It follows that all workers must be indifferent between x^w and x^*, which implies that $x^w = x^*$, yielding a contradiction. □

By the lemma and the second part of Proposition 5, for any match-partitional budget equilibrium (B, x) on the first $F-1$ firms and $W-1$ workers we can construct a state *R* in which the unique stable matching coincides with *x* and the unique supporting match-partitional budget sets coincide with *B* for the first $F-1$ firms and $W-1$ workers (firm *F* and worker *W* are matched with each other and their budget sets only include each other). Letting for definiteness $F \leqslant W$, and considering an allocation *x* in which all the firms are matched, we can let the budget set of any of the first $F-1$ firms include any of the first $W-1$ workers in addition to its current match (the workers' match-partitional budget sets are implied). Since any such budget equilibrium is a unique match-partitional equilibrium in some state, we have:

Corollary 7. *The communication cost of stable one-to-one matching with strict preferences between W workers and $F \leqslant W$ firms is at least $(F-1)(W-2)$ bits. The nondeterministic communication cost of stable many-to-one matching between W workers and F firms on any preference domain that includes strict one-to-one preferences and guarantees the existence of a stable matching is asymptotically equivalent to FW as $F, W \to \infty$.*

Corollary 7 generalizes quadratic lower bounds obtained by Gusfield and Irving [16] for finding a stable one-to-one matching with $F = W$ using particular querying languages. Specifically, they only allow queries of the form "which partner has rank *r* in your preference ranking" or "what rank partner i has in your preference ranking" [16, Theorems 1.5.1, 1.5.2]. Allowing general communication could in general reduce the communication cost, [39] but the corollary establishes that this is not the case.

The *deterministic* communication cost, i.e., that of actually of *finding* a stable matching, can in principle be substantially higher. However, for the preference domain on which the firms' preferences are strict and substitutable [44, Definition 6.2], a stable matching exists and can be found using only somewhat more communication. This can be done with a Gale–Shapley "deferred acceptance algorithm" [44, Theorems 6.7, 6.8], which takes at most $3FW$ steps, at each of which a match is proposed, accepted, or rejected. Since a match is described with at most $\log_2(FW)$ bits, we have a deterministic protocol that communicates at most $3FW\log_2(FW)$ bits. This only slightly exceeds the verification cost, and is exponentially less than full revelation of firms' preferences over combinations of workers. [40]

8. Deterministic communication

Of course, any practical protocol must be deterministic: it must find a desirable allocation without the benefit of an omniscient oracle. Such a protocol in general may need to reveal more information than needed for verification. In fact, deterministic realization of an IM choice rule sometimes require exponentially more communication than nondeterministic:

Example 4. Let $N = 2$ and $X = \{x \subset L : |x| = 2\}$, for some set L such that $|L| = 3m$. We interpret the agents as managers in a firms and L as a set of workers, and allocation $x \in X$ as choosing a pair of workers for a certain task. Manager 1 receives payoff 1 if the workers in x share a language, and payoff 0 otherwise. Manager 1 knows privately the language spoken by each worker. Publicly it is only known that each worker speaks one language, there are m languages spoken by a pair of workers, and m languages spoken by a single worker. Manager 2 receives payoff 1 if $x \subset y$ and payoff 0 otherwise, where $y \subset L$ is a particular group of $2m+1$ workers known privately to manager 2. The social goal is to give both managers a payoff of 1, which describes a choice rule that is CU (letting each manager's blocking set be X) and thus intersection-monotonic. Note that a socially desirable pair x always exists, and it can be verified simply by announcing it, which takes $2\log_2 L$ bits. However, the deterministic communication complexity of finding such a pair is asymptotically proportional to L, which follows from the problem's equivalence to the "Pair-Disjointness" problem analyzed in [29, Section 5.2].

However, in some well-known social choice problems the gap between deterministic and nondeterministic communication costs proves to be small. This is trivially true when even nondeterministic communication proves almost as hard as full revelation (e.g., in the surplus maximizing combinatorial allocation problem considered in Section 7.2). More interestingly, the gap is also small in some cases in which much less communication than full revelation suffices. For example, in a convex economy with the "gross substitute" property, Walrasian tatonnement converges quickly to a Walrasian equilibrium, which verifies Pareto efficiency [32, Section 17.H]. Similarly, in the many-to-one matching problem with strict substitutable preferences, a Gale–Shapley

[39] In fact, the proving method of [16] cannot be extended to general communication. The proofs in [16] use a "fooling set" in which all firms have the same and known preferences over workers. On this fooling set, we could use a simple protocol in which workers sequentially, in the reverse order of their desirability, chose firms from those that remain available. This protocol finds a stable matching with W steps and communicates at most $\log_2 F$ bits per step.

[40] Even if a firm's preference relation is known to be strict and substitutable, the communication burden of describing such a relation is still exponential in W, as shown by Echenique [10, Corollary 5].

deferred acceptance algorithm converges quickly to a "match-partitional" equilibrium, which verifies stability [44, Section 6.1]. In both these mechanisms, at each step, the designer offers budget sets for the agents, and the agents report their optimal choices from their respective budget sets. If the choices are inconsistent, the designer adjusts the budget sets to be "closer" to being an equilibrium. A "substitutability" condition on the agents' preferences allows to construct an adjustment process that is monotonic, and therefore converges quickly (enormously faster than full revelation). Some of the agents in such mechanisms even have the incentives to report truthfully (e.g., nonatomic agents in Walrasian tatonnement, the proposing agents in a deferred acceptance algorithm).

9. Relation to incentives

The concept of budget equilibrium has naturally arisen in mechanism design (implementation) with incentives. Indeed, any mechanism defines a "budget set" for each agent as the set of outcomes he can attain via possible strategies, and incentive compatibility requires that the agent not have a strictly preferred outcome within his budget set. With a single agent, this observation is known as the "Taxation Principle" [15]. In designing mechanisms for Nash implementation with many agents [33], budget sets can be constructed in a similar manner. Namely, take a mechanism that describes a strategy space S_i for each agent i and an outcome function $g : S_1 \times \cdots \times S_N \to X$. A strategy profile $s \in S_1 \times \cdots \times S_N$ is a Nash equilibrium of the mechanism if and only if $g(s)$ is each agent i's optimal alternative in his attainable set $B_i(s) = \left\{g\left(s_i', s_{-i}\right) : s_i' \in S_i\right\}$. Thus, the mechanism is equivalent to the budget protocol with the message space $M = \{(B_1(s), \ldots, B_N(s), g(s)) : s \in S_1 \times \cdots \times S_N\}$.[41] In particular, by this argument, Theorem 1 implies that any Nash implementable choice rule is monotonic.[42]

A key contribution of the present paper relative to the implementation literature is in showing that incentive considerations are not *necessary* for the revelation of supporting budget equilibria: By Theorem 2, *if agents have private knowledge of their preferences*, supporting budget sets must be revealed in *any* mechanism solving an intersection-monotonic social choice problem, even when agents can be relied upon to be truthful rather than selfish.

We can also make the point that incentives are not *sufficient* for the revelation of supporting budget sets when agents *have complete information* about one another's preferences. At first glance, this point appears to contradict the earlier observation that equilibrium strategies in Nash implementation must reveal supporting budget sets. However, there is no contradiction once we consider implementation using *multi-stage* mechanisms, in which equilibrium play does not reveal the agents' (contingent) strategies, and therefore need not reveal supporting budget sets. In particular, any Nash implementable choice rule can be implemented with a two-stage mechanism in whose equilibrium agents agree on a desirable alternative in the first stage without announcing any other information, and only in case of a disagreement, the mechanism would proceed to the

[41] This observation was made by Williams [53, Theorem 1] and by Dutta et al. [9]. The latter paper focused on Nash implementation of interior Pareto efficiency in smooth convex economies (considered in Section 7.1), and showed that for this problem, a supporting *Walrasian* equilibrium must be revealed, along the lines of the proof of Proposition 1.

[42] The converse is not true [33, Example 2], because not every budget protocol can be derived from a game form. For the same reason, even when a choice rule *can* be Nash implemented, this may require more communication than realizing it with a budget protocol. For example, Reichelstein and Reiter [43] examined the increase in communication required to Nash implement the Walrasian equilibrium choice rule.

second stage in which supporting budget sets are revealed.[43] Intuitively, when agents know one another's preferences, they already know supporting budget sets, and so the sets need not be communicated in equilibrium.

Finally, observe that just because a communication protocol must reveal supporting budget sets does not imply that the protocol must be incentive compatible. To be sure, in a budget protocol, no agent would have an incentive to deviate by proposing another alternative within his budget set. However, a budget protocol, being nondeterministic, does not specify what alternative an agent could get by "rejecting" the budget equilibrium announced by the oracle. Incentive compatibility must instead be examined in the context of *deterministic* communication. When a budget equilibrium correspondence is realized with a deterministic protocol, an agent may be able to manipulate his messages to influence his budget set to his advantage [32, Example 23.B.2].[44] Thus, in general, the restriction to incentive-compatible protocols increases the communication cost [42,12].

10. Conclusion

The "market design" literature has examined the attainment of socially desirable allocations using "price discovery" mechanisms, such as ascending auctions, tatonnement, and deferred acceptance algorithms. However, this literature has not answered two fundamental questions: (1) Why and when is the restriction to "price discovery" mechanisms justified? and (2) How should the "necessary," or "minimal," price space for a given problem be constructed? Instead, a few papers have proposed *ad hoc* price spaces for specific problems and established fundamental welfare theorems for them [35,31,2,3,28,18].

The present paper answers both questions by analyzing the minimal information that must be communicated in order to solve a given social choice problem when the preference information is distributed among the agents. The analysis answers (1) by characterizing the social choice problems for which any minimally informative verifying message is a price equilibrium (more generally "budget equilibrium"), and answers (2) by constructing the minimally informative verifying price equilibria for any given social choice problem. Thus, the paper provides a justification for and characterizes the scope of the "market design" approach (as opposed to more general mechanism design), and characterizes the form of "prices" that must be discovered to solve a given social choice problem. Contrary to widespread belief, prices are necessary not in order to incentivize the agents, but in order to aggregate distributed information about their preferences into a socially desirable decision. The necessity of revealing prices proves to be a useful step for identifying the communication costs of social choice rules.

To be sure, the paper does not fully solve the general "market design" problem of solving a given social choice problem with a practical mechanism that is *deterministic* and *incentive-compatible.*

[43] For example, take the one-stage mechanism proposed by Maskin [33, Theorem 3] to implement any monotonic choice rule F satisfying a "no veto power" condition with $N \geqslant 3$ agents. This mechanism can be converted into the following two-stage mechanism: In the first stage, agents simultaneously announce an alternative. If they agree on an alternative, it is implemented, otherwise we move to the second stage, in which each agent announces a state and an integer (without observing the others' first-stage messages). The outcome function is the same as in Maskin's mechanism. Applying Maskin's arguments, it is easy to check that the two-stage mechanism still Nash implements F, yet in any equilibrium the agents agree on an alternative in the first stage.

[44] An exception is given by "nonatomic" convex economies, in which individual agents have no influence on the Walrasian equilibrium prices. Another exception is when an agent's budget set depends only on other agents' types, as in the Vickrey–Groves–Clarke mechanism.

However, the paper has two important implications for this problem. The first implication is that in some social choice problems (such as the efficient combinatorial allocation problem), the space of prices that must be discovered proves to be prohibitively large, and the communication of such prices proves to be almost as hard as full revelation of preferences. In such cases, the designer of a practical mechanism must either moderate her goals or restrict attention to a smaller preference domain. The second implication is for the problems for which the required space of supporting prices proves to be manageable, and their communication proves much simpler than full revelation. For such problems, the characterization of the price space offers some clues for the design of practical mechanisms that must *find* an equilibrium from this space. In some important cases, mentioned in Section 8, a price (budget set) adjustment process can be constructed to converge quickly to a verifying budget equilibrium and to provide agents with the incentives for truthful reporting. Identifying more general approaches to constructing deterministic and incentive-compatible mechanisms solving a given social choice problem with minimal communication is an important question for further research.

Acknowledgements

I am grateful to Susan Athey, James Jordan, Jonathan Levin, Eric Maskin, Paul Milgrom, Andy Postlewaite, Ariel Rubinstein, Thomas Sjostrom, and the participants of innumerable seminars and workshops at which earlier drafts of this paper have been presented. Azeem Shaikh and Ronald Fadel provided excellent research assistance. I thank the Institute for Advanced Study at Princeton for its hospitality, and the Guggenheim Foundation and the National Science Foundation (Grants SES-0214500, SES-0427770) for financial support.

References

[1] D.D. Aliprantis, K.C. Border, Infinite Dimensional Analysis, Springer, New York, 1999.
[2] S. Bikhchandani, J. Mamer, Competitive equilibrium in an exchange economy with Indivisibilities, J. Econ. Theory 74 (1997) 385–413.
[3] S. Bikhchandani, J. Ostroy, The package assignment model, J. Econ. Theory 107 (2002) 377–406.
[4] X. Calsamiglia, Decentralized resource allocation and increasing returns, J. Econ. Theory 14 (1977) 262–283.
[5] P. Cramton, Y. Shoham, R. Steinberg (eds.), Combinatorial Auctions, MIT Press, Cambridge, 2006.
[6] G. Debreu, The coefficient of resource utilization, Econometrica 19 (1951) 273–292.
[7] G. Debreu, A social equilibrium existence theorem, Proc. Nat. Acad. Sci. 38 (1952) 886–893.
[8] G. Debreu, Neighboring economic agents, in: Mathematical Economics: Twenty Papers of Gerard Debreu, Cambridge University Press, New York, 1983, pp. 173–178.
[9] B. Dutta, A. Sen, R. Vohra, Nash implementation through elementary mechanisms in economic environments, Econ. Design 1 (1995) 173–203.
[10] F. Echenique, Counting combinatorial choice rules, Games Econ. Behav. 58 (2007) 231–245.
[11] G.E. Edgar, Measure, Topology, and Fractal Geometry, Springer, New York, 1990.
[12] R. Fadel, I. Segal, The communication cost of selfishness, Working Paper, Stanford University.
[13] J. Greenberg, The Theory of Social Situations: An Alternative Game-Theoretic Approach, Cambridge University Press, Cambridge, 1990.
[14] O.A. Gross, Preferential arrangements, Amer. Math. Mon. 69 (1962) 4–8.
[15] R. Guesnerie, A Contribution to the Pure Theory of Taxation, Cambridge University Press, Cambridge, 1995.
[16] D. Gusfield, R. Irving, The Stable Marriage Problem: Structure and Algorithms, MIT Press, Cambridge, 1989.
[17] P. Hammond, Game forms versus social choice rules as models of rights, in: K.J. Arrow, A.K. Sen, K. Suzumura (Eds.), Social Choice Re-examined, vol. II (IEA Conference, vol. 117), Macmillan, London, 1997, pp. 82–95, (Chapter 11).
[18] J.W. Hatfield, P.R. Milgrom, Matching with contracts, Amer. Econ. Rev. 95 (2005) 913–935.
[19] F.A. Hayek, The use of knowledge in society, Amer. Econ. Rev. 35 (1945) 519–530.

[20] L. Hurwicz, On the dimensional requirements of informationally decentralized Pareto-satisfactory processes, in: K.J. Arrow, L. Hurwicz (Eds.), Studies in Resource Allocation Processes, Cambridge University Press, New York, 1977, pp. 413–424.
[21] L. Hurwicz, E. Maskin, A. Postlewaite, Feasible Nash implementation of social choice correspondences when the designer does not know endowment or production sets, in: J. Ledyard (Ed.), The Economics of Informational Decentralization: Complexity, Efficiency, and Stability, Kluwer Academic Publishing, Dordrecht, 1995.
[22] L. Hurwicz, S. Reiter, Designing Economic Mechanisms, Cambridge University Press, Cambridge, 2006.
[23] T. Ishikida, T. Marschak, Mechanisms that efficiently verify the optimality of a proposed action, Econ. Design. 2 (1996) 33–68.
[24] J.S. Jordan, The competitive allocation process is informationally efficient uniquely, J. Econ. Theory 28 (1982) 1–18.
[25] B.-G. Ju, Nash implementation and opportunity equilibrium, Working Paper, University of Kansas, 2005.
[26] Y. Kannai, Continuity properties of the core of a market, Econometrica 38 (1970) 791–815.
[27] H. Karloff, Linear Programming, Birkhäuser Verlag, Basel, 1991.
[28] A.S. Kelso, V.P. Crawford, Job matching, coalition formation, and gross substitutes, Econometrica 50 (1982) 1483–1504.
[29] E. Kushilevitz, N. Nisan, Communication Complexity, Cambridge University Press, Cambridge, 1997.
[30] R.J. Lipton, E. Markakis, E. Mossel, A. Saberi, On approximately fair allocations of indivisible goods, in: Proceedings of the Fifth ACM Conference on Electronic Commerce, 2004, pp. 125–131.
[31] A. Mas-Colell, Efficiency and decentralization in the pure theory of public goods, Quart. J. Econ. 94 (1980) 625–641.
[32] A. Mas-Colell, M.D. Whinston, J. Green, Microeconomic Theory, Oxford University Press, New York, 1995.
[33] E. Maskin, Nash equilibrium and welfare optimality, Rev. Econ. Stud. 66 (1999) 23–38.
[34] R.P. McLean, A. Postlewaite, Excess functions and nucleolus allocations of pure exchange economies, Games Econ. Behav. 1 (1989) 131–143.
[35] J.-C. Milleron, Theory of value with public goods: a survey article, J. Econ. Theory 5 (1972) 419–477.
[36] E. Miyagawa, Reduced-form implementation, Columbia University Discussion Paper No. 0203-09, 2002.
[37] H. Moulin, B. Peleg, Cores of effectivity functions and implementation theory, J. Math. Econ. 10 (1982) 115–145.
[38] K. Mount, S. Reiter, The information size of message spaces, J. Econ. Theory 28 (1974) 1–18.
[39] N. Nisan, The communication complexity of approximate set packing and covering, in: 29th International Colloquium on Automata, Languages, and Programming, 2002.
[40] N. Nisan, I. Segal, The communication requirements of efficient allocations and supporting prices, J. Econ. Theory 129 (2006) 192–224.
[41] D.C. Parkes, Price-based information certificates for minimal-revelation combinatorial auctions, in: Padget et al. (Ed.), Agent-Mediated Electronic Commerce IV, Lecture Notes in Artificial Intelligence, vol. 2531, Springer, Berlin, 2002, pp. 103–122.
[42] S. Reichelstein, Incentive compatibility and informational requirements, J. Econ. Theory 34 (1984) 32–51.
[43] S. Reichelstein, S. Reiter, Game forms with minimal message spaces, Econometrica 56 (1988) 661–692.
[44] A.E. Roth, M.A.O. Sotomayor, Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis, Cambridge University Press, Cambridge, 1990.
[45] A. Rubinstein, The single profile analogous to multi profile theorems: mathematical logic's approach, Internat. Econ. Rev. 25 (1984) 719–730.
[46] F. Sato, On the informational size of message spaces for resource allocation processes in economies with public goods, J. Econ. Theory 24 (1981) 48–69.
[47] A.K. Sen, The impossibility of a Paretian Liberal, J. Polit. Economy 78 (1970) 152–157.
[48] R. Serrano, O. Volij, Walrasian allocations without price-taking behavior, J. Econ. Theory 95 (2000) 79–106.
[49] L. Shapley, M. Shubik, Quasi-cores in a monetary economy with nonconvex preferences, Econometrica 34 (1966) 805–827.
[50] T. Sjostrom, Implementation by demand mechanisms, Econ. Design 1 (1996) 343–354.
[51] M. Sotomayor, Three remarks on the many-to-many stable matching problem, Math. Soc. Sci. 38 (1999) 55–70.
[52] G. Tian, A unique informationally efficient allocation mechanism in economies with consumption externalities, Int. Econ. Rev. 45 (2004) 79–111.
[53] S.R. Williams, Realization and Nash implementation: two aspects of mechanism design, Econometrica 54 (1986) 139–152.