

Communication in Economic Mechanisms*

Ilya Segal[†]

This draft: January 19, 2006

Abstract

This chapter considers the problem of to find allocations that satisfy certain social goals when economic agents have private information about their preferences. While economists have traditionally considered the problem of providing incentives for agents to fully reveal their preferences, such full revelation is often impractical or undesirable, for several reasons: (1) it may require a prohibitive amount of communication as measured in bits or real numbers (2) it may be costly for agents to evaluate their complete preferences, and (3) the revealed information may be exploited by the designer or other agents. Thus, we consider the question: What is the minimal information that must be elicited from the agents in order to achieve the goals? Note that the question arises even if agents can be counted on to report truthfully.

Segal (2005) shows that for a large class of social problems, any minimally informative way to verify that a given alternative is desirable is equivalent to giving each agent

*This chapter was prepared for the 9th World Congress of the Econometric Society, London, August 19-24, 2005. I gratefully acknowledge financial support of the National Science Foundation (grants SES 0214500, 0427770). I also thank my discussant, Roger Myerson, and the participants of too many seminars and conference presentations where some of the research surveyed in the chapter was presented. In particular, I am grateful to Susan Athey, Jonathan Levin, Eric Maskin, Paul Milgrom, Andy Postlewaite, Thomas Sjostrom, and James Jordan for discussions related to this topic. Last but not least, I am indebted to my coauthors Noam Nisan and Ronald Fadel for helping clarify my understanding of the topics discussed in this chapter.

[†]Department of Economics, Stanford University, Stanford, CA 94305. Email: ilya.segal@stanford.edu

a “budget set” – a subset of the social alternatives (which could in general be described by personalized nonlinear prices), and asking each agent to verify that the proposed alternative is optimal to him within his budget set. Therefore, any communication mechanism that yields a solution to the social problem must also yield a supporting price equilibrium. This result formalizes Hayek’s insight about the role of prices as minimal communication needed to solve the social coordination problems. The class of problems for which price revelation is necessary proves quite large. For example, it includes such social goals as exact or approximate efficiency, voluntary participation, stability to group deviations, and some notions of fairness. For such goals, price revelation is necessary regardless of the preference domain, which allows for nonconvexities or discrete decisions (e.g., as in combinatorial auctions or matching problems). On the other, the particular form of prices to be used depends on the problem. Segal (2005) suggests an algorithm for deriving the form of price equilibria that verify the solution of a given problem with minimal information revelation. Applied to several well-known social problems, the algorithm generates the price equilibrium concepts that have been proposed for these problems. The necessity of revealing such prices bounds below the communication costs of the problem, measured in bits (“communication complexity”), real numbers (“dimension of message spaces”), evaluation costs, or in other ways. These results indicate which problems can be solved in a practical way and which problems cannot, and what role prices have in mechanisms that solve them.

This chapter outlines the results described above, a substantial body of related work in both economics and computer science, and potential extensions. In particular, it discusses how to provide incentives for truthfull communication, how to distribute communication to reduce individual agents’ communication costs, the role of prices in achieving probabilistic (average-case) social goals, and in problems with interdependent values.

1 Introduction

This chapter considers the problem of finding allocations that satisfy certain social goals when economic agents have private information regarding their preferences. This problem has been discussed since at least the early 20th century debate on alternative economic mechanisms, but it has received renewed attention recently in the literature on “market design,” which proposes mechanisms to solve various allocation problems of practical importance. For example, the “two-sided matching problem” arises to allocating workers across firms, students across schools, or medical interns across medical schools (Roth and Sotomayor (1990)). The “combinatorial auction problem” arises in allocating bundles of indivisible items among bidders (Cramton et al. (2006)). Both agents’ preferences and social goals in these problems differ substantially from those in the classical economies studied earlier. In particular, agents’ preferences often exhibit nonconvexities and indivisibilities, and the social goals may include exact or approximate efficiency, voluntary participation, stability to group deviations, and even some notions of fairness.

A major theme in the “market design” literature is that the choice of mechanism is not determined by incentives alone. Indeed, if incentive compatibility were the only concern, it could be verified with a direct revelation mechanism. However, full revelation of agents’ preferences is often impractical or undesirable, for several reasons: First, sometimes full revelation requires a prohibitive amount of communication—e.g., a bidder in a combinatorial auction would have to announce his valuations for all possible bundles of objects, whose number is exponential in the number of objects. Second, agents may have to incur “evaluation costs” to learn their own preferences. Finally, the more information is revealed, the more deviations exploiting the revealed information become available to agents or the designer, as noted in the literature on communication and mechanisms without perfect commitment (e.g., Myerson (1991, Section 6) and Salanie (1997, Section 6)). For all these reasons, the “market design” literature has examined a variety of mechanisms that aim to achieve the desired goals without fully revealing agents’ preferences. For example, in the many proposed “iterative” combinatorial auction designs, bidders submit and modify their bids for various bundles over time. This raises the question: What is the minimal information that must be

elicited from the agents in order to achieve the goals? Note that the question arises even agents are willing to communicate sincerely.

An early discussion of the communication problem can be found in Hayek's (1945) critique of socialist planning. Hayek 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." He 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 it 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 Hayek's ideas inspired economists to study the workings of price mechanisms, their place among all other conceivable allocation mechanisms and their domain of applicability have remained unclear. For example, consider the best-known results about price mechanisms — the Fundamental Welfare Theorems. The First Welfare Theorem says that announcing supporting prices is *sufficient* to verify the Pareto efficiency of an allocation, but not that it is *necessary*. The Second Welfare Theorem says only 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, an efficient allocation can be computed and imposed directly, without using prices.¹ The theorems have nothing to say about possible efficient non-price mechanisms in an economy with distributed knowledge of preferences.

A major advance in understanding the role of prices was made by the literature on the "informational efficiency" of Walrasian equilibria, spurred by Hurwicz (1977) and Mount and Reiter (1974). In contrast to the Fundamental Welfare Theorems, the literature fol-

¹There are many computational optimization techniques that do not compute supporting dual variables – e.g., the ellipsoid method or the simplex method for linear programming (Karloff 1991).

lowed Hayek in modeling allocation mechanisms in an economy with *decentralized knowledge of preferences*. (Similar techniques were independently developed in the computer science literature on “communication complexity,” which considered discrete communication problems—see Kushilevitz and Nisan (1997)). The literature considered the problem of verifying Pareto efficiency in economies with convex preferences, and showed that the Walrasian equilibrium verifies an efficient allocation using the minimal number of real variables among all continuous verification mechanisms. However, the recent “market design” problems have different preference domain (e.g., with nonconvexities and indivisibilities, Walrasian equilibria or continuous mechanisms may not exist), different social goals (e.g., coalitional stability or approximate efficiency), and different relevant communication costs (e.g., the number of bits, or the cost of evaluating preferences).

It turns out that the necessity of price revelation can be demonstrated in a general social choice setting that covers most recent “market design” problems. This is shown in Segal (2005), who characterizes the class of social choice problems (defined by preference domains and social goals) for which any communication mechanism must reveal supporting “prices” (which in general take the form of abstract subsets of alternatives offered to the agents). The class turns out to include a number of important economic problems. Segal (2005) also suggests an algorithm for deriving the form of budget sets that need to be used to verify the solution of a given problem with minimal information revelation. These results have implications for the communication costs of various social choice problems, measured in bits, real numbers, evaluation costs, or in other ways. In particular, the results are used to see which problems can be solved in a practical way and which problems cannot, and what role prices must have in mechanisms that solve them.

The objective of this chapter is to survey the results described above, a substantial body of related work, and some potential extensions. We begin in Section 2 with a very simple example in which the concepts of communication and minimally informative messages are defined, and the necessity of price revelation is demonstrated. Section 3 extends these ideas to a large class of social choice problems. In Section 4 we apply the general analysis to several social choice problems, including classical convex economies, combinatorial auctions and two-sided matching. In each of these applications, we derive the space of budget equilibria

corresponding to minimally informative messages, and use this space to identify the communication cost. Section 5 discusses and relates several alternative measures of communication cost, such as the number of real variables versus bits transmitted, communication cost of individual agents rather than in the aggregate, the number of preference evaluations performed by agents, and some notions of privacy preservation. Section 6 discusses some further issues, such as comparison between the costs of communication and verification, probabilistic (average-case) social goals, the additional communication cost of incentivizing agents, and the role of prices when agents' utilities are interdependent. Many of the questions raised in Section 6 are still open and require further investigation.

2 A Simple Example

We illustrate the main ideas with a very simple example: One object is to be allocated between two agents with valuations v_1, v_2 . Each agent's valuation is his privately observed *type*, and the valuation pair (v_1, v_2) is called the *state*. Suppose that we know a priori that the valuations lie in the set $\{0, 1, 2, 3\}$. The goal is to find an "optimal" allocation, which for now we define as efficiency—giving the object to the agent with the higher valuation (when the valuations coincide, both allocations are optimal). What communication is needed to find an optimal allocation?

To begin with, we measure the communication cost as the number of bits needed to encode the agents' messages, as in the "communication complexity" literature (Kushilevitz and Nisan 1997).² While in this simple example the communication cost proves to be trivial, the ideas developed in this section will prove useful in much more complex settings.

²Thus, agents are forced to communicate using only binary messages (bits). If instead they could communicate using a k -letter alphabet, a letter from the alphabet could be encoded using $\log_2 k$ bits, so the communication length would only be reduced by the constant factor $\log_2 k$. Thus, the choice of the alphabet is relatively unimportant in large problems.

2.1 Communication Protocols

An obvious way to find an optimal outcome is by asking agents to reveal their private information:

Protocol 1 (Full Revelation): The agents announce their valuations v_1, v_2 (encoded in bits).

Since each agent needs $\log_2 4 = 2$ bits to encode his valuation, in total 4 bits are sent.

The object is allocated to agent 1 if $v_1 > v_2$ and to agent 2 otherwise.

Can we find an optimal allocation with less communication? The answer is yes, by letting agents make announcements sequentially and condition their announcements on the past announcements. Thus, we define sequential communication as follows:

Definition 1 *A communication protocol is (i) an extensive-form game form in which all moves are binary, (ii) agents' strategies in this game (each agents' strategy contingent on his private type as well as history), and (iii) an assignment of allocations to the terminal nodes of the game.*

We assume that agents obey the prescribed strategies—e.g., agents could well be computers who follow their programs. (The problem of providing incentives not to deviate is discussed in Subsection 6.2 below.) We want the protocol to implement in every state an optimal allocation for this state. Consider the following example:

Protocol 2 (One-sided Revelation): Agent 1 announces his valuation v_1 (encoded in 2 bits),

then agent 2 announces an allocation of the object (1 bit). Thus, 3 bits are sent in total. Agent 2's strategy is to allocate the object to agent 1 if $v_1 > v_2$ and to himself otherwise.

In Protocols 1 and 2, the number of bits sent is the same in any state (v_1, v_2) . In other protocols, the amount of communication may differ across states:

Protocol 3 (English auction): The protocol starts with a price $p = 0$, and then agents send messages in sequence:

1. Agent 2 says “stop” or “raise”. If he says “stop,” allocate the object to agent 1, otherwise set $p = 1$ and continue.
2. Agent 1 says “stop” or “raise.” If he says “stop,” allocate the object to agent 2, otherwise set $p = 2$ and continue.
3. Agent 2 says “stop” or “raise.” If he says “stop,” allocate the object to agent 1, otherwise allocate the object to agent 2.

Each agent’s strategy is to say “raise” when his valuation exceeds the current price p and say “stop” otherwise. Given these strategies, the protocol always implements an optimal allocation. Depending on the agents’ valuations, the protocol may stop after the agents send 1, 2, or 3 bits.

We now focus on the simplest measure of the communication cost, known as “worst-case” communication complexity – the largest number of bits sent across all states.³ Can we find a protocol with a lower communication cost than the protocols above?

Protocol 4 (Bisection): Agent 1 says “low” if $v_1 \in \{0, 1\}$ or “high” if $v_1 \in \{2, 3\}$ (1 bit).

Then agent 2 announces an allocation (1 bit). Agent 2’s strategy is as follows: If agent 1 said “low,” agent 2 allocates the object to agent 1 if $v_2 = 0$ and to himself otherwise. If agent 1 said “high,” agent 2 allocates the object to himself if $v_2 = 3$ and to agent 1 otherwise. This protocols finds an optimal allocation using 2 bits.

Can we find an optimal allocation using fewer than 2 bits in the worst case? In general, how can we find the communication complexity of a given *problem*, defined as the *minimal* communication complexity of a protocol solving this problem? To tackle this question, it is convenient to represent communication geometrically in the state space. In our example,

³Alternatively, given a probability distribution over valuation pairs (v_1, v_2) , we could consider “average-case” communication complexity as the *expected* number of bits sent in the protocol (this is also known as “distributional” complexity). In Protocol 3, this expected number could be close to 1 if the valuations are very likely to be low. This is related to Shannon’s (1948) information measure, which allows coding more frequent messages with shorter strings of bits. We consider average-case communication complexity in Subsection 6.3 below.

		v_2			
		0	1	2	3
v_1	0	1, 2	2	2	2
	1	1	1, 2	2	2
	2	1	1	1, 2	2
	3	1	1	1	1, 2

Table 1: State space

the state space is described by a matrix, where in each state (cell) we put the set of optimal allocations (Table 1).

Each node of a communication game tree corresponds to an “event” – a subset of the state space in which the node is reached. Note that since agent 1’s message at any of his decision nodes depends only on his own type, it slices the corresponding event into sub-events horizontally; similarly, agent 2’s messages slice events vertically. Thus, by induction on the depth of the node we can see that the event corresponding to any node must be a product set. In computer science, such events are called “rectangles,” although they need not be geometric (i.e., contiguous) rectangles.

Now consider the rectangles corresponding to the terminal nodes of a protocol. Note that such rectangles must partition the state space (since in each state, exactly one terminal node is reached). Also, if the protocol finds an optimal allocation, then for each rectangle corresponding to a terminal node there must exist a single allocation that is optimal on the whole rectangle, and which could be assigned to the node. In computer science, rectangles with this property are called “monochromatic”. Thus, the terminal nodes of the protocol must partition the state space into monochromatic rectangles. The partitions generated by Protocols 1-4 are shown in Table 2.

The worst-case communication complexity W of a protocol is the maximal depth of the corresponding binary tree. Since the number T of terminal nodes in such a tree is at most 2^W , we must have $W \geq \log_2 T$. Thus, the worst-case communication complexity of finding an optimal allocation can be bounded below by bounding below the size of any partition of

		v_2			
		0	1	2	3
v_1	0	2	2	2	2
	1	1	2	2	2
	2	1	1	2	2
	3	1	1	1	2

Protocol 1

		v_2			
		0	1	2	3
v_1	0	2			
	1	1	2		
	2	1		2	
	3	1			2

Protocol 2

		v_2			
		0	1	2	3
v_1	0				
	1	1	2		
	2			1	2
	3				

Protocol 3

		v_2			
		0	1	2	3
v_1	0	1	2		
	1				
	2	1			2
	3				

Protocol 4

Table 2: Communication Partitions

the state space into monochromatic rectangles.⁴

2.2 Verification Protocols

Since characterizing all communication protocols has proven to be very hard, a lot of attention has been put into providing lower bounds on communication complexity. As discussed before, such a bound can be obtained by finding the minimal size of a partition of the state space into monochromatic rectangles. We can further simplify the problem by allowing the rectangles to overlap, i.e., allow *coverings* rather than partitions of the state space.

A covering of the state space into monochromatic rectangles can be interpreted as a *verification protocol* (also called “nondeterministic communication” in computer science – see Kushilevitz and Nisan (1997, Chapter 2)). To understand verification, imagine an omniscient oracle who knows the agents’ valuations and consequently the optimal allocation(s), but needs to prove to an ignorant outsider that an allocation x is indeed optimal. The oracle 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. (Thus, the set of states on which the message is accepted is a rectangle.) The acceptance of message m by all agents must *verify* to the outsider that allocation x is optimal. (Thus, the rectangle is monochromatic.) The (worst-case) complexity of a verification protocol with message space M is the minimum number of bits needed to encode a message, which is $\log_2 M$.⁵

While verification protocols are patently unrealistic, their examination proves useful for the following reasons:

1. Any communication protocol can be verified by the oracle sending all the messages instead of the agents, and having each agent accept the message sequence if and only

⁴The bound in general will not be tight, for two reasons: First, some partitions of the state space cannot arise in any communication protocol (Kushilevitz and Nisan 1997, Figure 2.1). Second, the inequality $W \geq \log_2 T$ is tight only in balanced trees (such as Protocols 1,2, and 4) and strict in unbalanced trees (such as Protocol 3).

⁵Such communication is called “nondeterministic” in computer science because the oracle “guesses” an acceptable message (and there may be more than one such message in a given state). In contrast, the communication protocols defined in the previous subsection are called “deterministic”.

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 terminal nodes (message sequences). Therefore, verification is a generalization of communication, and so communication cost is bounded below by verification cost.

2. A famous economic example of verification is Walrasian equilibrium. The role of the oracle is played by the “Walrasian auctioneer,” who announces the equilibrium prices and allocation. Each agent accepts the announcement if and only if his announced allocation constitutes his optimal choice from the budget set delineated by the announced prices. We will describe a natural extension of such price-based verification mechanisms to general social choice problems in Section 3.
3. A verification protocol may be viewed as the steady state of an iterative communication protocol. 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 iterative processes include “tatonnement” processes for finding Walrasian equilibria, “deferred acceptance algorithms” for finding stable matchings, and ascending-bid auctions for finding efficient combinatorial allocations. In some settings, the iterative processes converge very quickly, though in general this cannot be guaranteed (see Subsection 6.1 below.)

2.3 Minimally Informative Messages and Prices

In order to verify optimality using the smallest number of bits, we need to find a minimal covering of the state space with monochromatic rectangles. For this purpose, we want to use larger rectangles, corresponding to messages that reveal less information about the agents’ types. Formally, we define the following partial “informativeness” order on messages:

Definition 2 *Message m is less informative than (or verified by) message \tilde{m} if m is accepted on a larger set of states (rectangle) than \tilde{m} . Also, m is a minimally informative message verifying (the optimality of) allocation x if any less informative message verifying x is as informative as m .*

		v_2			
		0	1	2	3
v_1	0	2	2	2	2
	1		2	2	2
	2			2	2
	3				2

Table 3: Minimally Informative Message

Graphically, a minimally informative message m verifying x corresponds to a maximal rectangle contained in the set of states in which x is optimal. Typically, a given allocation may be verified by many minimally informative messages, which are not comparable in the informativeness order. For example, with two agents, one minimally informative rectangle could be tall and narrow (revealing little information about agent 1 and much about agent 2), while another short and wide (revealing much about agent 1 and little about agent 2).

It can be seen that for any message m verifying x there exists a less informative message m' that is a minimally informative message verifying x .⁶ Thus, starting with any verification protocol, we can replace every message with a minimally informative message verifying the same allocation, and obtain a verification protocol with the same number of messages that uses only minimally informative messages. (Furthermore, this replacement may allow us to discard some of the messages while still covering the state space with the remaining rectangles.)

We proceed to characterize the minimally informative messages in our simple example. The states in which allocating the object to agent 2 is optimal are marked with “2” in Table 3. The minimally informative messages verifying the optimality of allocating the object to agent 2 correspond to the largest rectangles that fit into this set, i.e., that do not extend below the diagonal. These are exactly the geometric rectangles with one corner on the diagonal and another in the top-right state $(v_1, v_2) = (0, 3)$.

⁶This observation is trivial when the state space is finite. For general state spaces, this is shown in Segal (2005, Lemma 2).

Note that any minimally informative message verifying allocation 2 can be described as a “price equilibrium:” The oracle names a price $p \in \{0, 1, 2, 3\}$ and the allocation of the object to agent 2, and each agent accepts if and only if the allocation is optimal to him given the price. That is, agent 2 accepts if and only if he is willing to buy at price p (i.e., $v_2 \geq p$), and agent 1 accepts if and only if he is willing *not* to buy at price p (i.e., $v_1 \leq p$). (The rectangle depicted in Table 3 corresponds to $p = 1$.) Thus, *the minimally informative messages verifying allocation to agent 2 are characterized as price equilibrium messages for prices $p \in \{0, 1, 2, 3\}$* . Symmetrically, the same is true for minimally informative messages verifying allocation to agent 1. This implies that any communication protocol that finds an optimal allocation must reveal enough information to construct a supporting price equilibrium.

This observation implies a lower bound on the communication cost: Since each price $p \in \{0, 1, 2, 3\}$ has to be used in the diagonal state $(v_1, v_2) = (p, p)$ (regardless of which of the two optimal allocations is verified in this state), we need to use at least 4 messages. Thus, the worst-case communication cost is at least $\log_2 4 = 2$ bits. This lower bound is achieved by Protocol 4.⁷

Suppose now that the agents’ valuations instead lie in the $[0,1]$ interval. The minimally informative messages verifying an allocation again correspond to price equilibria (see Figure 1), but now any price $p \in [0, 1]$ is a unique equilibrium price in the diagonal state $(v_1, v_2) = p$, and so any verification protocol must use an infinite number of messages. Formally, we will allow infinite protocols with infinite message spaces, and measure their “dimensionality” — i.e., how many real numbers are announced by the agents or the oracle (see Subsection 5.1 below for technical details). Intuitively, the message space in the example must have at least the same dimensionality as the diagonal — i.e., be at least one-dimensional. This lower bound is tight: just like in Protocol 2, we can find an optimal allocation with agent 1 revealing his

⁷A set of states with the property that no two elements of the set can share a message is called a “fooling set” in computer science, and “a set with the uniqueness property” in the economic literature on communication. The size of such a set bounds below the size of the message space. The novelty here is that the fooling set (in our example, the diagonal) is not chosen *ad hoc* but characterized as the set of states with a unique supporting price. This characterization can be extended to a large class of social choice problems.

valuation with one real number, and then agent 2 reporting an optimal allocation with 1 bit.

2.4 Other social goals

The result on the necessity price revelation can be extended to social goals other than efficiency:

Example 1 (Approximate efficiency): Take $\varepsilon > 0$, and say that allocating the object to agent i is “optimal” if and only if $v_i \geq v_{-i} - \varepsilon$. Minimally informative messages verifying that allocation to agent 2 is optimal are described by the geometric rectangles with one corner on the line $v_2 = v_1 - \varepsilon$ and another in the top-right corner of the state space (see Figure 2). Such messages can be interpreted as price equilibria in which the agents face different prices p_1, p_2 for the object such that $p_2 = p_1 - \varepsilon$. This observation can again be used to bound below the communication cost. Note that two diagonal states with coordinates further apart than ε cannot share a price equilibrium (regardless of which allocation it supports). With continuous valuations in $[0,1]$, we can find $1/\varepsilon$ such distinct diagonal points, hence we need to use at least $1/\varepsilon$ distinct messages, and the communication cost it at least $\log_2(1/\varepsilon)$ bits. This lower bound is almost achieved by letting agent 1 announce his valuation rounded off to a multiple of ε , and agent 2 then report an optimal allocation.⁸

Note also that there exist social goals that *cannot* be verified with a price equilibrium:

Example 2 (Minimize efficiency): An allocation is “optimal” if it allocates the object to the agent with the *lower* valuation. No price equilibrium with prices p_1, p_2 supporting allocation of the object to agent 2 can verify that the allocation is optimal: if it is ever an equilibrium, it will remain an equilibrium in state $(v_1, v_2) = (0, 3)$, in which the object must go to agent 1.

⁸Compare this to the earlier finding that *exact* efficiency with continuous valuations would require one-dimensional continuous communication. The relationship between continuous communication and discrete approximation is discussed in more detail in Subsection 5.1 below.

Example 3 (Egalitarian efficiency): In addition to allocating the object efficiently, we must also determine a payment between the agents to equalize their utilities — i.e., if agent i is “wins” the object, he must pay $v_i/2$ to the “loser.” This payment cannot be verified with a price equilibrium: Any price equilibrium would remain an equilibrium when the winner’s value goes up, but egalitarian efficiency requires that the winner’s payment to the loser must increase.

The examples suggests that price equilibria can only be used to verify social goals that are somehow “congruent” with private preferences (such as efficiency or approximate efficiency), but not those opposing or orthogonal to private preferences (such as minimization of efficiency or equalization of utilities).

Finally, note the difference between whether (a) price equilibria *can* be used to verify a social goal, and (b) the *minimally informative messages* verifying the social goal are price equilibria. (a) means that in any state, for any optimal allocation x in the state there exists a price equilibrium that verifies the optimality of x . E.g. the efficiency of an allocation in state (v_1, v_2) can always be verified with a price equilibrium, say, setting price $p = (v_1 + v_2)/2$. This is similar to the traditional Fundamental Welfare Theorems (although the example fails the usual convexity assumptions of the theorem). In the previous subsection, we have also shown that (b) holds for the goal of efficiency. Yet, for social goals other than efficiency, (a) does not imply (b):

Example 4: Suppose there are three possible allocations and a single agent. (We could add a second agent with a constant utility over the allocations.) A state is described by the agent’s valuations (v_1, v_2, v_3) for the three allocations. An allocation is defined as “optimal” if the agent’s utility from it is at least as high as from *at least one* of the other two allocations. Any optimal allocation in any state can be verified with a price equilibrium, e.g., with prices $(p_1, p_2, p_3) = (v_1, v_2, v_3)$ for the three allocations. However, consider a message in which the agent verifies that allocation 1 is optimal. This is a minimally informative message verifying allocation 1, but it is not equivalent to a price equilibrium: it does not reveal any prices at which the agent prefers allocation 1 to allocation 2 or to allocation 3, since it does not bound above either $v_2 - v_1$ or

$v_3 - v_1$ (it only reveals that *one* of the differences is nonpositive, but does not reveal which one).

3 General Social Choice Problems

3.1 Setup

We now extend the observations made in Section 2 to general social choice problems. 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 that are rational (i.e., complete and transitive). Each agent i 's preference relation is assumed to be his privately observed *type*, and the set of his possible types is denoted by $\mathcal{R}_i \subset \mathcal{P}$. A *state* is a preference profile $R = (R_1, \dots, R_N) \in \mathcal{R}_1 \times \dots \times \mathcal{R}_N \equiv \mathcal{R}$, where \mathcal{R} is the *state space*, also known as *preference domain*. The goal of communication is to implement a *choice rule*, which is a correspondence $F : \mathcal{R} \rightarrow X$. For every state $R \in \mathcal{R}$, the set $F(R) \subset X$ describes the *optimal* alternatives in this state.

We focus on the verification problem described in Section 2: An omniscient oracle knows the agents' valuations and consequently the optimal allocation(s), but he needs to prove to an ignorant outsider that an allocation x is indeed optimal. 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. The acceptance of message m by all agents must verify to the outsider that allocation x is optimal.

We can define two notions of verification:

Definition 3 *A verification protocol verifies choice rule F if $\forall R \in \mathcal{R} \exists x \in F(R) \exists m \in M$ that is acceptable in state R and verifies x . The protocol fully verifies choice rule F if $\forall R \forall x \in F(R) \exists m \in M$ that is acceptable in state R and verifies x .*

Thus, simple verification requires only *one* optimal alternative to be verifiable in each state, while full verification requires *all* optimal alternatives to be verifiable. We are ul-

mately interested in simple verification (communication is not required to find more than one optimal alternative), but full verification will prove a useful intermediate concept.

3.2 Verification with Budget Equilibria

We extend the notion of a “price equilibrium” to this general social choice setting, in which we may not even have any divisible goods in which prices could be measured. Thus, we consider abstract budget sets, which are general subsets of the space of alternatives (and which may or may not be delineated by prices). A *budget equilibrium message* consists of a proposed alternative $x \in X$ and a *budget set* $B_i \subset X$ for each agent i . Each agent $i \in N$ accepts message (B_1, \dots, B_N, x) if and only if there is no alternative in his budget set B_i that he strictly prefers to the proposed alternative x . (B_1, \dots, B_N, x) is a *budget equilibrium in state* $R \in \mathcal{R}$ if it is accepted by all agents in this state.⁹ It is convenient to define $L(x, R) = \{y \in X : xRy\}$ — the *lower contour set* of preference relation R_i at alternative x . Then the budget equilibrium condition can be written as $B_i \subset L(x, R_i)$ for all agents i .

To represent a budget equilibrium message graphically, it is convenient to “order” the agents’ preferences by the ranking of alternative x , i.e., by the set inclusion order on $L(x, R_i)$ (see Figure 3). Since in general this is not a complete order, a one-dimensional axis can only represent a “slice” of an agent’s type space. (The setting studied in Section 2 was a special case in which each agent’s type in fact *was* one-dimensional—ordered by his willingness to pay for the object.) Yet, however imprecise, Figure 3 allows us to develop some useful intuitions. A budget equilibrium message (B_1, B_2, x) is the set of states in which $B_i \subset L(x, R_i)$ for $i = 1, 2$, and in the figure it is represented with a geometric rectangle with one corner at (B_1, B_2) and another in the right-hand corner of the state space (where $L(x, R_1) = L(x, R_2) = X$).

Figure 3 also makes it clear that increasing budget sets makes a budget equilibrium more informative: budget equilibrium (B', x) is more informative than budget equilibrium (B, x)

⁹A number of related concepts have been suggested, including “social equilibrium” (Debreu 1952), “social situations” (Greenberg 1990), “effectivity functions” (Moulin and Peleg 1982), “effectivity forms” (Miyagawa 2002), “opportunity equilibrium” (Ju 2001), and “interactive choice sets” (Serrano and Volij 2000). However, all these papers have motivated the concept by incentives, rather than deriving it from communication among sincere agents.

whenever $B_i \subset B'_i$ for all agents i , Graphically, the rectangle corresponding to (B', x) is then included in the rectangle corresponding to (B, x) .

We can now define a *budget protocol* as a verification protocol in which the oracle’s message space M is a collection of budget equilibria, such that each equilibrium (B_1, \dots, B_N, x) from M verifies the equilibrium alternative x . Which choice rules can be verified with a budget protocol? Traditional Fundamental Welfare Theorems say that in a convex exchange economy, an allocation is Pareto efficient if and only if it can be verified with a Walrasian equilibrium (which is a kind of budget equilibrium). The theorems have been extended to some “non-classical” social choice problems, for which different kinds of budget equilibria have been proposed.¹⁰ We extend these results to general social choice rules, by characterizing choice rules F that are fully verified with a budget protocol.

According to the definition of full verification, we want to check that for any alternative $x \in X$, in each state $R \in \mathcal{R}$ in which x is optimal there exists a budget equilibrium (B, x) verifying x . To check this, it suffices to check the largest budget sets supporting x in state R , i.e., $B'_i = L(x, R_i)$ for each i (see Figure 3). That this budget equilibrium (B', x) verifies x means that x must remain optimal in any state R' “above” R , i.e., in which $L(x, R_i) = B'_i \subset L(x, R'_i)$ for each i . This property of choice rules is formally known as follows:

Definition 4 (Maskin (1999)) *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 verified by a budget protocol if and only if it is monotonic.*¹¹

Results equivalent to Theorem 1 are stated in Williams (1986, Theorem 2), Miyagawa’s (2002, Theorem 1), Ju (2001), and Greenberg (1990, Theorem 10.1.2). The present formulation and the idea of the proof are from Segal (2005).

¹⁰Including the Pareto rule in public-good economies (Milleron 1972) and general economies with numeraire (Mas-Colell 1980; Bikhchandani and Mamer 1997; Bikhchandani and Ostroy 2002), and stable many-to-one matching problems with and without transfers (Kelso and Crawford 1982; Hatfield and Milgrom 2005).

¹¹This implies that F is *verified* by a budget protocol if and only if it has a nonempty-valued monotonic subcorrespondence.

The deficiency of Theorem 1 is that, just like the traditional Fundamental Welfare Theorems, it does not rule out that choice rule F could be verified with a non-budget protocol that might reveal less information and have lower communication costs than any budget protocol verifying F . To rule this out, we would like to require the following stronger property:

Definition 5 *Choice rule F satisfies the Budget Equilibrium Revelation Property (BERP) if for any message verifying the optimality of an alternative $x \in X$ there exists a less informative budget equilibrium (B, x) that verifies the optimality of x .*

BERP is illustrated in Figure 4. When applied to a message m that fully reveals a state R (i.e., corresponds to a single point $\{R\}$ in Figure 4), 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 that F is fully verified with a budget protocol, and so by Theorem 1 that F is monotonic. However, BERP is stronger, since it requires a budget equilibrium verifying x to be constructed without knowing the exact state, upon observing *any message* verifying x . Note that BERP ensures that any minimally informative message verifying an alternative in F must be equivalent to a budget equilibrium message.

Contrary to the impression created by Figure 4, not all monotonic choice rules satisfy BERP. Figure 4 is misleading when feasible contour sets $L(x, R_i)$ cannot be ordered, in which case there do exist monotonic choice rules that do not satisfy BERP (see Example 4 in Section 2). Yet, the figure can be still used to develop intuition for which rules do satisfy BERP. To check whether a message $m = m_1 \times m_2$ verifies some budget equilibrium (B_1, B_2, x) that verifies x , it again suffices to check the largest budget sets that support x in all states from m , which are $B'_i = \cap_{R_i \in m_i} L(x, R_i)$ for each i (see Figure 4). Thus, it suffices to check that this equilibrium verifies x , i.e., that x is optimal in any state R' in which $B'_i \subset L(x, R'_i) \forall i \in N$. Formally, this property can be defined as follows

Definition 6 *Choice rule F is Intersection-Monotonic (IM) if $\forall m = m_1 \times \dots \times m_N \subset \mathcal{R}$, $\forall x \in \cap_{R \in m} F(R)$, and $\forall R' \in \mathcal{R}$ such that $\cap_{R \in m} L(x, R_i) \subset L(x, R'_i) \forall i \in N$, we have $x \in F(R')$.*

Theorem 2 *Choice rule F satisfies the Budget Equilibrium Revelation Property if and only if it is Intersection-Monotonic.*

Intersection monotonicity 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 $m_j = \{R'_j\}$ for $j \neq i$) –the full property 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 its verification requires revelation of supporting budget sets.

3.3 Examples of Intersection-Monotonic Rules

Segal (2005) shows that a number of important choice rules are intersection-monotonic on the universal preference domain \mathcal{P}^N (and therefore on any smaller domain), including

- Weak Pareto efficiency.¹²
- A notion of approximate Pareto efficiency (e.g., with quasilinear utilities, approximating the maximal achievable total surplus within ε).
- The weak core.
- Stable matching.
- The envy-free rule (requiring that no agent envies another agent's allocation).

More generally, the class of IM rules includes any rule from the following class:

Definition 7 *Choice rule F is a Coalitionally Unblocked (CU) choice rule if for some blocking correspondence $\beta : X \times 2^N \rightarrow X$,*

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

In words, for each coalition $S \subset N$ and each candidate alternative $x \in X$, the blocking correspondence defines a “blocking set” $\beta(x, S) \subset X$. An alternative $x \in X$ is optimal in state R if and only if no coalition $S \subset N$ can find a strict Pareto improvement over x in

¹²The 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.

its blocking set $\beta(x, S)$.¹³ It is easy to see that all the above examples of choice rules are CU rules, for different specifications of the blocking correspondence. Segal (2005) shows that any CU choice rule is IM. There do exist IM rules that are not CU, but their economic significance is unclear. A Venn diagram for choice rules summarizing the above results is drawn in Figure 5.

3.4 The Budget-Shrinking Algorithm

Now we look for minimally informative messages verifying a given choice rule, which under BERP must be equivalent to budget equilibria. We propose an algorithm to construct such budget equilibria for any given IM choice rule. Thus, for any given social choice problem, the algorithm constructs and characterizes the budget equilibria that verify the problem with minimal revelation of information. For simplicity, we restrict attention to IM choice rules that are extendable to the universal preference domain $\mathcal{R} = \mathcal{P}^N$. (In particular, note that any CU choice rule is extendable to \mathcal{P}^N using the same blocking correspondence.)

The proposed algorithm obtains a minimally informative message verifying a given alternative x by starting with any message verifying x and stretching the corresponding rectangle sequentially agent-by-agent.¹⁴ For an IM choice rule, we can focus on budget equilibrium messages, and their stretching corresponds to shrinking the agents' budget sets. As illustrated in Figure 6, we can start with a budget equilibrium (B_1, B_2, x) verifying alternative x , and “stretch” the rectangle in the direction of agent 1 as much as possible, while still verifying x . This stretching, illustrated with the horizontal arrow, corresponds to “shrink-

¹³CU choice rules have been also known as “respecting group rights,” with $y \in \beta(x, S)$ interpreted as the “one-way right” of coalition S to block alternative x with alternative y (Hammond 1997, Section 5). The “rights” literature, initiated by Sen (1970), 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 4 below, the preference domains and coalitional rights will be defined to ensure nonempty-valuedness.

¹⁴The algorithm is independently proposed by Segal (2005) and Hurwicz and Reiter (2006, who call it the “rectangle method”). However, Segal’s (2005) application of the algorithm to the special case of intersection-monotonic choice rules allows to focus on budget equilibrium messages, and stretch them by shrinking the agents’ budget sets.

ing” agent 1’s budget set from B_1 to B'_1 . Next, “stretch” the rectangle described by budget equilibrium (B'_1, B_2, x) in the direction of agent 2. This stretching, represented with the vertical arrow, corresponds to “shrinking” agent 2’s budget set from B_2 to B'_2 . This yields a budget equilibrium message (B'_1, B'_2, x) that can no longer be stretched, i.e., corresponds to a minimally informative verifying message. (The same procedure works with any number of agents: sequential agent-by-agent stretching yields a minimally informative verifying message.)

Note that the resulting equilibrium (B'_1, B'_2, x) can be described by the “boundary” state $R \in \mathcal{P}^N$ in which the agents’ lower contour sets at x coincide with B'_1, B'_2 , and x is on the verge of becoming non-optimal. Formally, the boundary states R for alternative x and the corresponding minimally informative budget equilibria are characterized by the condition

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

In words, each agent i ’s budget set is his smallest lower contour set for which x is still optimal, holding other agents’ preferences fixed.

When the preference domain \mathcal{R} is a strict subset of \mathcal{P}^N , we face the following complications:

- There typically exist many budget equilibria that are equally informative to (*) but have even smaller budget sets. For example, in exchange economies in which preferences are known to be monotone in consumption, a Walrasian budget equilibrium, in which the budget sets are half-spaces, is equivalent to the budget equilibrium in which the half-spaces are replaced with their boundary hyperplanes (i.e., waste is not allowed). The budget equilibria characterized by (*) have the largest budget sets among those that are equally informative, and it proves convenient to focus on them (if only because they are guaranteed to exist). Thus, in shrinking agent i ’s budget set, we only shrink it to the intersection of the feasible lower contour sets in \mathcal{R}_i for which x is still optimal, and not any further, even when such shrinking might yield an equally informative message.
- Since not all subsets of X may serve as lower contour sets, the “boundary states”

characterized by (*) are not guaranteed to be in \mathcal{R} . However, it is still true that (*) with $R \in \mathcal{P}^N$ characterizes (up to equivalence) the minimally informative verifying budget equilibria .

If a boundary state R satisfying (*) *does* fall in the preference domain \mathcal{R} , then we can see that $(L(x, R_1), \dots, L(x, R_N), x)$ is a unique (up to equivalence) budget equilibrium verifying x in state R . Such an equilibrium cannot be discarded if we want to verify alternative x in state R with a budget protocol. This observation will prove useful for bounding below the size of the message space, and thus the communication cost. (A simple example of this occurred in Section 2, in which the boundary states were those on the diagonal). A complication arises when there are many optimal alternatives in state R : since we do not require *full* verification, we do not have to verify any given $x \in F(R)$. In such situations, we resort to additional application-specific tricks to bound below the number of budget equilibria needed for verification.

4 Some Applications

4.1 Pareto Efficiency in Convex Economies

In a *smooth convex exchange economy*, 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 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 positive aggregate endowment $\bar{x} \in \mathbb{R}_{++}^L$: $\bar{X} = \{x \in X : \sum_i x_i = \bar{x}\}$.¹⁵ The goal is to verify an allocation that is Pareto efficient within \bar{X} .

We use the budget-shrinking algorithm described in Subsection 3.4 to derive minimally informative messages verifying the Pareto efficiency of an allocation $x \in \bar{X}$ with $x \gg 0$.¹⁶

¹⁵We consider a space X of alternatives that is larger than the feasible set \bar{X} , to allow budget sets to include infeasible allocations, as the Walrasian budget sets do.

¹⁶We restrict attention to $x \gg 0$ to avoid the problem of non-existence of supporting Walrasian prices (see, e.g., Mas-Colell et al. (1995, Figure 16.D.2)).

The derivation can be illustrated in the standard Edgeworth box depicted in Figure 7. 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 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, 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 budget-shrinking algorithm yields a Walrasian equilibrium. Furthermore, any Walrasian equilibrium is invariant to budget shrinking - i.e., satisfies (*). A formalization of this argument yields

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 economy¹⁷ if and only if it is equivalent to a Walrasian equilibrium supporting x , i.e., a budget equilibrium (B, x) with*

$$B_i = \{y \in X : p \cdot y_i \leq p \cdot x_i\} \quad \forall i \in N \quad (1)$$

for some commodity 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.

The proposition implies that the minimal message space required for verifying any interior Pareto efficient allocation in any convex economy is the space of Walrasian equilibria. We now discuss the implications of this finding for the verification cost measured as the dimension of the message space. (We keep the arguments informal; see Subsection 5.1 for how the dimension could be formally defined.) Informally, 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 $L - 1$

¹⁷If non-smooth preferences are allowed, Walrasian equilibria remain minimally informative messages verifying Pareto efficiency, but other such messages emerge – see Segal (2005) for details.

real variables, the space of Walrasian equilibria has dimension $(L - 1) + (N - 1)L = NL - 1$. This compares favorably to full revelation of agents' utility functions, which would require an infinite-dimensional message space.

If we don't want *full* verification, and only need to verify *one* efficient allocation in each state, we can further reduce the dimension of the state space. In fact, it is possible to verify Pareto efficiency without any communication—e.g., by always giving all the endowment to agent 1. We rule out such corner allocations, focusing on “non-dictatorial” Pareto efficiency. Note that the nondictatorial Pareto rule can be verified by fixing an “endowment allocation” $\omega \in \bar{X}$ with $\omega \gg 0$ and announcing a Walrasian equilibrium (B, x) such that $\omega \in B_i$ for all i , which exists in any convex economy (Mas-Colell et al. 1995, Section 17.BB). Since such equilibria satisfy the additional “budget constraints” $\sum_l p_l \omega_{il} = \sum_l p_l x_{il}$ for all i , they can be communicated using $(L - 1) + (N - 1)(L - 1) = N(L - 1)$ real numbers.

In fact, it is impossible to verify nondictatorial Pareto efficiency using fewer than $N(L - 1)$ real numbers. This can be shown using a “fooling set” consisting of the *Cobb-Douglas economies*, in which each agent i 's utility function takes the form $u_i(x_i) = \prod_l x_{il}^{\alpha_{il}}$ with a positive parameter vector $\alpha \in \mathbb{R}_{++}^L$, with the normalization $\sum_l \alpha_{il} = 1$. Note that all nondictatorial Pareto efficient allocations in a Cobb-Douglas economy are interior, and the first-order equilibrium conditions imply that no two distinct Cobb-Douglas economies share an interior Walrasian equilibrium. Therefore, verification requires using a subspace of Walrasian equilibria whose dimension is at least that of Cobb-Douglas economies, which is $N(L - 1)$:

Corollary 1 *The verification cost of nondictatorial 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.*

Corollary 1 was first established in the “informational efficiency” literature (Hurwicz 1977; Mount and Reiter 1974) for verification protocols satisfying a continuity property. Here it has been derived in a different way—from the purely set-theoretic characterization of minimally informative messages as Walrasian equilibria (Proposition 1). Unlike the old approach, the set-theoretic approach does not require any topological restrictions on communication or any scalar measure of the communication cost, and easily extends to other

social choice problems, including those considered in the “market design” literature.¹⁸

4.2 Efficiency in Quasilinear Economies

In *economies with numeraire*, the space of alternatives take the form $X = K \times \mathbb{R}^N$, where K is a finite set of (*non-monetary*) allocations, and \mathbb{R}^N describes the *transfers* of numeraire (money) to the agents. The feasible set takes the form $\bar{X} = \{(k, t) \in X : \sum_i t_i = 0\}$, i.e., requires a balanced budget.

For simplicity, we let each agent i 's preference domain \mathcal{R}_i consist of preferences R_i over $(k, t) \in X$ that are quasilinear in his consumption of numeraire, i.e., described by a utility function of the form $u_i(k) + t_i$.¹⁹ Pareto efficiency is then equivalent to requiring that the non-monetary allocation $k \in K$ maximize the total surplus $\sum_i u_i(k)$, regardless of the allocation of numeraire. (The example in Section 2 was a special case with two non-monetary allocations, in which surplus-maximization required giving the object to the agent with the higher valuation.)

We use the budget-shrinking algorithm of Subsection 3.4 to derive minimally informative messages verifying Pareto efficiency. We illustrate this algorithm in an Edgeworth box depicted in Figure 8, in which the vertical dimension represents allocations of numeraire between the agents, and the horizontal dimension represents the non-monetary 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 . Shrink the lower contour set of agent 1 as much as

¹⁸The analysis also extends to convex economies with public goods. For such economies, the budget-shrinking algorithm yields *Lindahl equilibria*, i.e., budget equilibria described by linear anonymous prices for the private goods and linear personalized “Lindahl” prices for the public goods. This can in turn be used to derive the dimensionality of the message space needed to verify Pareto efficiency (which was first obtained by Sato (1981)).

¹⁹In fact, the analysis of this subsection holds on the larger preference domain where each agent i 's preferences are (i) independent of other agents' transfers t_{-i} , (ii) continuous and nondecreasing in his own transfer t_i , and (iii) allow compensation (i.e., for any $x \in X$ and any $k \in K$ there exists $t \in \mathbb{R}$ such that $(k, t) R_i x$). This follows from the observation that any lower contour set of a preference relation satisfying (i)-(iii) is also a lower contour set of some quasilinear preference relation.

possible while preserving the Pareto efficiency of x (as illustrated with the downward arrows in the figure). The furthest we 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'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 (*). 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 $(k, t) \in \bar{X}$ in a quasilinear economy if and only if it is equivalent to a valuation equilibrium supporting (k, t) , i.e., a budget equilibrium $(B, (k, t))$ in which*

$$B_i = \{(k', t') \in X : p_i(k') + t'_i \leq p_i(k) + t_i\} \quad \forall i \in N \quad (2)$$

for some price vector $p \in \mathbb{R}^{NK}$ satisfying

$$\sum_i p_i(k') = \sum_i p_i(k) \quad \text{for all } k' \in K. \quad (3)$$

Any such equilibrium is a unique valuation equilibrium supporting allocation (k, t) in the state given by the agents' utility functions $u_i = p_i$ for each i .

Valuation equilibria were introduced by Mas-Colell (1980) and studied by Bikhchandani and Mamer (1997) and Bikhchandani and Ostroy (2002). These papers have extended classical welfare theorems to such equilibria: An allocation is Pareto efficient if and only if it is supported by a valuation equilibrium. The contribution of Proposition 2 lies in showing that valuation equilibria constitute *minimally informative* verification of Pareto efficiency in an economy with numeraire.

Proposition 2 implies that the minimal message space required for verifying any efficient allocation in an 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 describe a price vector satisfying (3) using $(N - 1)(K - 1)$ real numbers.

If we don't require *full* verification, 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 $p \in \mathbb{R}^{NK}$ satisfying (3) still must be used. Indeed, while in the “boundary” state given by utility functions $(u_1, \dots, u_N) = (p_1, \dots, p_N)$ all allocations are efficient by (3), by the second part of Proposition 2, the agents' budget sets must be described by the same prices p no matter which allocation the equilibrium supports. Therefore, verifying Pareto efficiency with quasilinear preferences requires the announcement of an $(N - 1)(K - 1)$ -dimensional price vector.

This lower bound on the communication cost is in fact achieved by the communication protocol in which the first $N - 1$ agents announce their normalized utility functions, and then the last agent chooses a surplus-maximizing allocation. To summarize:

Corollary 2 *The continuous verification cost Pareto efficiency in a quasilinear economy is $(N - 1)(K - 1)$ real numbers, and it is achieved with a communication protocol.*

A large number of problems with more restricted quasilinear preferences have been considered, and we describe two of them below.

4.2.1 Combinatorial Allocation

In this problem, a set L of objects to be allocated among the agents, and so the allocation set can be written as $K = N^L$. The preference domain consists of those quasilinear preferences in which each agent i 's utility depends only on his own consumption bundle $k^{-1}(i)$ and is nondecreasing in this bundle (in the set inclusion order). For the particular case of $N = 2$, the budget-shrinking algorithm yields valuation equilibria in which each agent i 's price $p_i(k)$ is nondecreasing in his bundle $k^{-1}(i)$, and we can normalize prices so that $p_i(k) = 0$ when $k^{-1}(i) = \emptyset$.²⁰ In the state $(u_1, u_2) = (p_1, p_2)$, all allocations are efficient by (3), but the normalized price vector in any valuation equilibrium must coincide with p . Thus,

²⁰Application of the budget-shrinking algorithm to the case of $N > 2$ agents appears more complicated, and we have not attempted it.

the communication cost is bounded below by the dimensionality of this price space, which is $2^L - 1$. This lower bound is achieved with a communication protocol in which agent 1 announces his utility function and agent 2 chooses an efficient allocation. To summarize:

Corollary 3 *The continuous verification cost of efficient combinatorial allocation of L objects between two agents is $2^L - 1$, and it is achieved with a communication protocol*

Corollary 3 was obtained by Nisan and Segal (2004). A number of other results have been obtained on the potential communication savings in combinatorial allocation problems when agents utility functions are a priori restricted to lie in certain classes, such as those complement-free utilities, submodular utilities, utilities with substitute objects, utilities with homogeneous objects, etc. For some of these results, see Nisan and Segal (2004), Dobzinski et al. (2005), and Babaioff and Blumrosen (2005).

4.2.2 Binary Utilities

Suppose that agents' utilities are known to be $u_i(k) \in \{0, 1\}$ for all $k \in K$. Then the budget-shrinking algorithm yields valuation equilibria described by prices $p_i(k) \in \{0, 1\}$ for all i, k , and we can normalize prices so that $p_i \neq (1, \dots, 1)$ for each agent i (since this price would be equivalent to $p_i = (0, \dots, 0)$). In the state $(u_1, \dots, u_N) = (p_1, \dots, p_N)$, all allocations are efficient by (3), but the normalized price vector in any valuation equilibrium must coincide with p . Thus, the communication cost measured in bits is bounded below by the binary logarithm of size of this price space. The number of possible price vectors in $\{0, 1\}$ satisfying $\sum_i p_i(k) = r$ for a given integer r is $\binom{N}{r}^K$, since for each allocation we allocate r "1's" among N agents' utilities. For simplicity taking $r = N/2$ (with N even) and using Stirling's formula yields the lower bound

Corollary 4 *The communication cost of efficiency with binary utilities is asymptotically at least NK bits as $N \rightarrow \infty$.*

Thus, as the number of agents grows, the cost is asymptotically the same as that for full revelation of utilities.

This setting can be interpreted as “approval voting,” interpreting $u_i(k) = 1$ as agent i ’s “approval” of allocation k , with the goal being to find an allocation approved by most agents. Conitzer and Sandholm (2005) derive the above result with a different proof.²¹

4.3 Approximate Efficiency in Quasilinear Economies

When finding an exactly efficient allocation may be prohibitively costly, we may want to allow approximate efficiency. Consider again quasilinear economies, fix $\varepsilon > 0$, and say that an allocation is ε -efficient if it maximizes the total surplus within ε (regardless of the transfers).²² Note that the goal of ε -efficiency can be described as a CU choice rule (see subsection 3.3), by requiring the grand coalition to pay a tax ε in numeraire for blocking a candidate allocation (and not allowing any smaller coalition to block). Then an allocation is unblocked if and only if it is ε -efficient. Thus, the defined choice rule is IM, and therefore satisfies the Budget Equilibrium Revelation Property.

To characterize the minimally informative budget equilibria verifying ε -efficiency, we again use the budget-shrinking algorithm. Note that in the Edgeworth box depicted in Figure 8, an allocation x is ε -efficient if and only if agent 1’s indifference curve passing through x does not fall below agent 2’s indifference curve passing through x by more than ε . Shrinking agent 1’s lower contour set yields a “budget line” that is ε below agent 2’s indifference curve at all off-equilibrium allocations. After that, agent 2’s lower contour set cannot be shrunk. Thus, the sum of the prices delineating the agent’s budget sets must be higher by ε for any off-equilibrium allocation than for the equilibrium allocation. (Intuitively, agents should be “penalized” for deviations to off-equilibrium allocations.) Formally, we have

Proposition 3 *A message is a minimally informative message verifying ε -efficiency of allocation $x = (k, t) \in \bar{X}$ in a quasilinear economy if and only if it is equivalent to an ε -valuation*

²¹Conitzer and Sandholm (2005) also characterize the communication costs of several other common voting rules. Some of these rules, such as approval voting and the majority rule, are interesection-monotonic, and so their results can be alternatively derived by characterizing supporting budget sets. Others are not even monotonic, and the results are proven using different “fooling sets.”

²²This is a “worst-case” notion of approximation. Average-case approximation is discussed in Subsection 6.3 below.

equilibrium supporting x , i.e., a budget equilibrium (B, x) with budget sets described by (2) for some price vector $p \in \mathbb{R}^{NK}$ satisfying

$$\sum_i p_i(k') = \sum_i p_i(k) + \varepsilon \text{ for all } k' \in K \setminus \{k\}. \quad (4)$$

Any such equilibrium is a unique ε -valuation equilibrium in the state given by the agents' utility functions $u_i = p_i$ for all i .

Observe that if agents' utility functions are bounded, then any approximation $\varepsilon > 0$ can be achieved with finite communication in which agents announce their utilities rounded off to multiples of ε/N . Thus, arbitrarily close approximation can be achieved with discrete communication, and so the communication cost of approximation should be measured in bits. In Subsection 5.1 below we discuss how this cost relates to the cost of exact efficiency measured in real numbers.

Now we focus on the setting of “binary utilities” described in the previous subsection. In this setting, the agents' utilities are in $\{0, 1\}$, and the budget-shrinking algorithm yields prices in $\{0, 1\}$. Note that approximation within $\varepsilon = N - 1$ can be achieved with a “dictatorial” protocol in which one agent announces an allocation that maximizes his utility. Approximation within $\varepsilon < N - 1$ requires finding an allocation that gives utility 1 to at least two agents. The communication complexity of this can be bounded below by counting how many “diagonal” states, i.e., states with total surplus 1 for all allocations, can be “covered” with a given ε -valuation equilibrium, and dividing by the total number of diagonal states. This gives a lower bound on the number of price equilibria that need to be used, yielding (see Segal 2005):

Corollary 5 *With binary utilities, the communication cost of achieving a better approximation of efficiency than letting one agent choose an allocation is at least $(K - 1) \log_2(1 + 1/(N - 1))$ bits.*

Interpreting the problem as “approval voting”, this means that the cost of finding even an alternative that is approved by more than one voter is proportional to the number of alternatives. The result can also be applied to the combinatorial allocation problem, by

constructing a “large” subset K of allocations such that the agents can have arbitrary utilities for allocations from K , and that all allocations that are better than dictatorial allocations lie in K . Nisan and Segal (2004) construct such a set K whose size is exponential in the number of objects. Corollary 5 then implies that any improvement upon giving all objects to one agent requires exponential communication.

4.4 Stable Many-to-One Matching

Now we consider the problem of stable many-to-one matching, which is studied in Roth and Sotomayor (1990), henceforth RS. In the problem, the set N of agents is partitioned into the set F of firms and the set W of workers. A *matching* between firms and workers is a binary relation $x \subset F \times W$. With a slight abuse of notation, we let $x(i)$ denote the set of agent i 's matching partners in matching x . We restrict the space of alternatives to include only *many-to-one* matchings, in which a worker cannot match with more than one firm: $X = \{x \subset F \times W : |x(w)| \leq 1 \forall w \in W\}$. We examine 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.

A coalition S can deviate from a candidate match $x \in X$ by (i) breaking any matches and (ii) creating new matches between its members; formally, it can deviate to any match $y \in X$ such that $y \setminus (S \times S) \subset x \setminus (S \times S)$.²³ This describes a CU choice rule as defined in Subsection 3.3 above, which is therefore intersection-monotonic, hence satisfies the Budget Equilibrium Revelation Property. We proceed to characterize the minimally informative budget equilibria verifying stability.

Intuitively, since a worker's preferences depend only on his employer, his budget set can be described in terms of the available employers. On the other hand, a firm has preferences over *groups* of workers, and so its budget sets can be described in terms of such available groups. Describing such a combinatorial budget set for a firm would require exponential communication (2^W bits).

²³We might also ban a coalition from breaking matches between outsiders, but this is irrelevant when externalities in preferences are ruled out.

Fortunately, it turns out that *minimally informative* budget equilibria verifying stability don't use combinatorial budget sets for firms. To see this, note that a budget equilibrium verifies stability if and only if each firm f 's budget set includes any group consisting of some 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 hires strictly better off. In the minimally informative budget equilibria, characterized by (*), firms must have minimal budget sets necessary for verification, which means that each firm f 's budget set must include *exactly* the groups consisting of some of f 's current employees and some of those workers who do not have f in their budget set. Thus, 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. In such an equilibrium, each potential off-equilibrium match is allocated to either the firm's or the worker's budget set but not both. (Such an equilibrium is illustrated in Figure 9, in which the equilibrium matching is described with dashed vertical lines, firm's budget sets are described with downward arrows and workers' budget sets are described with upward arrows.) Formally, the argument yields

Proposition 4 *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 , i.e., a budget equilibrium (B, x) satisfying*

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

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 $i \in N$.

The finding 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 ω are implied by the conditions $\phi \cap \omega = x$ and $\phi \cup \omega = F \times W$. Thus, the cost

of verifying a stable matching is $O(FW)$ as $F, W \rightarrow \infty$. This is exponentially smaller than that of full revelation of a firm’s preference rankings over subsets of workers, which asymptotically takes $\log_2(2^W!) \sim W \cdot 2^W$ bits as $W \rightarrow \infty$ (using Stirling’s formula).

If we are not required to *fully* verify stability, we only need to verify *one* stable matching in each state, and so need not use all match-partitional equilibria. However, we can show that “almost” all such equilibria need to be used. 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, and so we can restrict attention to matchings x in which $|x(i)| \leq 1$ for all $i \in N$. With such preferences, Segal (2005, Lemma 5) shows that the uniqueness of a stable matching in state R can be ensured by adding one matched firm-worker pair, and completing other agents’ preferences in a way consistent with R . Therefore, using the second part of Proposition 4, 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. Thus, we can bound below the communication cost of stability by that of describing a budget equilibrium with $F - 1$ firms and $W - 1$ workers. Since any worker’s budget set may include any of the firms in addition to its current employer (if in fact he is employed), we have the following lower bound

Corollary 6 *The verification cost of stable matching between W workers and F firms with strict one-to-one preferences is at least $(F - 2)(W - 1)$ bits. The communication cost of finding a 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 at least FW as $F, W \rightarrow \infty$.*

Corollary 6 generalizes quadratic lower bounds obtained by Gusfield and Irving (1989) 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” (their Theorem 1.5.1) or “what rank partner i has in your preference ranking” (their Theorem 1.5.2). The corollary establishes that allowing general communication does

not reduce the communication cost.

The communication cost of actually of *finding* a stable matching may in principle be substantially higher than that of verification. However, when firms’ preferences are restricted to be strict and substitutable (RS Definition 6.2), a stable matching exists and can be found using only slightly more communication. This can be done with a Gale-Shapley “deferred acceptance algorithm” (RS 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 still exponentially less than full revelation of firms’ preferences over combinations of workers.²⁴

5 Different Measures of Communication Cost

5.1 Continuous vs. Discrete Communication

Here we discuss in greater detail the definition of continuous communication cost, and its relation to the discrete communication cost measured in bits. In a continuous communication protocol, agents should be able to send real-valued elementary messages, but we also want to allow finite-valued messages (say, to communicate discrete allocations), without counting the latter toward the communication cost. Thus, the worst-case cost of continuous communication is defined as the maximum number of real-valued elementary messages sent in the course of the protocol. In a verification problem, we can identify the communication cost with the dimension of the oracle’s message space M , i.e., the number of real numbers needed to encode the oracle’s message. For this purpose, we must have a topology on M .

A well-known problem in continuous communication is the possibility of “smuggling” multidimensional information in a one-dimensional message space with a one-to-one encoding. Traditionally, dimension smuggling has been ruled out by imposing a continuity restriction

²⁴Indeed, it would take $\log_2(2^W)! \sim 2^W \cdot W$ bits to describe a strict preference rankings groups of workers when W is large (using Stirling’s formula). Even if a firm’s preference relation is known to be strict and substitutable, the number of bits needed to describe such a relation is still exponential in W , as shown by Echenique (2005, Corollary 5).

on the communication protocol (Abelson 1980; Luo and Tsitsiklis 1991; Mount and Reiter 1974; Walker 1977). For example, Mount and Reiter (1974) and Walker (1977) require the “message correspondence” from states into messages to have a continuous selection in any neighborhood. This requirement rules out *a priori* some important communication protocols, e.g., those in which agents announce discrete allocations.²⁵

A different way to rule out “smuggling” is proposed by Nisan and Segal (2004). They note that when many dimensions are “smuggled” into a one-dimensional message, a small error in the message would yield a huge error in its “meaning,” i.e., the set of states it represents. Thus, smuggling can be avoided by using a metric on messages that is not arbitrary but based on their meaning. Specifically, the distance between messages m and m' can be defined as the Hausdorff distance between the corresponding rectangles in the state space \mathcal{R} .²⁶ The communication cost is then defined as a metric dimension of the message space M .²⁷ In contrast to the traditional approach, this approach does not rule out any protocols, and in particular allows protocols that mix continuous and discrete messages.

Another advantage of the Nisan-Segal definition is that it implies a relation between continuous communication and discrete approximation:

Proposition 5 (*Nisan-Segal 2004*). *A protocol verifying a certain social goal with a message space whose box-counting dimension is d can be discretized into a protocol verifying approximation of the goal within ε using asymptotically $d \log \varepsilon^{-1}$ bits as $\varepsilon \rightarrow 0$.*

Intuitively, the oracle can communicate a message rounded-off within ε using roughly $d \log \varepsilon^{-1}$ bits, and the round-off yields a small distortion in the meaning of the message.

²⁵For example, consider the setting of Section 2 in which an object is allocated between two agents with valuations in $[0,1]$. The protocol in which agent 1 announces his valuation with 1 real number and then agent 2 reports an optimal allocation with 1 bit is discontinuous on the diagonal, where the optimal allocation switches. Insisting on continuity would require a two-dimensional message space (as in full revelation), which we believe overstates the communication cost in that example.

²⁶This distance is based on an underlying metric on the state space \mathcal{R} of preference relation profiles. In turn, the latter can be derived from a given metric on X along the lines suggested by Debreu (1983).

²⁷There are different notions of metric dimension—e.g., the Hausdorff dimension, the box-counting dimension, and the packing index (Edgar 1990), but in all economic examples considered they yield the same answers.

This means that the discretized protocol yields an allocation that is optimal for some state that is not too far from the true state, and therefore approximates an optimal allocation. Thus, metric dimension d of the message space is indicative of the communication complexity of achieving a “fast” approximation of efficiency, in which each additional bit reduces the error by the same factor $(e^{1/d})$.²⁸

On the other hand, it turns out a somewhat slower but still practical approximation is sometimes achieved with much less communication than that implied by the continuous cost of exact optimality. A dramatic example of this obtains in Calsamiglia’s (1997) model of allocating a homogeneous divisible good between two agents in a quasilinear economy. In this model, exact surplus-maximization requires infinite-dimensional communication (which can be shown by adapting Corollary 2 to an infinite set of allocations K), but Nisan and Segal (2004) demonstrate a protocol that approximates the maximal surplus within ε using $O(\varepsilon^{-1})$ bits. This approximation is still considered “fast” (polynomial) in computer science. In cases like this, the continuous measure of communication cost used in the economic literature seriously overstates the “hardness” of the problem.

5.2 Individual Communication Cost and Distributed Communication

We can reduce the communication costs of individual agents by not having them observe all the communication, i.e., by creating non-trivial information sets in the communication protocol. Also, the allocation need not be broadcast to all agents: instead, we could require that each agent i at the end of communication announce the component x_i of the alternative that he is concerned about. (Formally, we write the space of alternatives as $X = X_1 \times \dots \times X_N$, so that each agent i ’s preferences depend only on component x_i of $x = (x_1, \dots, x_N) \in X$.)²⁹ The individual communication cost of an agent can be defined as the number of elementary messages (bits or real numbers) that he must observe and send. This model of

²⁸Related observations are made by Hurwicz and Marschak (2003a,b).

²⁹Instead of requiring that agent i announce x_i we could require that he only *learn* x_i : If describing x_i is relatively “cheap”, as it is in most applications, then requiring that agent i announce x_i would not increase his communication burden substantially.

“distributed communication” better captures Hayek’s idea of decentralization.³⁰ When the number of agents is large, distributed communication could allow a substantial savings in agents’ individual communication costs.

Similarly to the aggregate communication cost, individual communication costs can be bounded below by considering a distributed version of the verification problem: The oracle has a message space M_i for each agent i , and he announces a “distributed message” $(m_1, \dots, m_N) \in M \subset M_1 \times \dots \times M_N$, where M is interpreted as the set of “legal” messages. Each agent i observes only his own message m_i , and accepts or rejects it based on his own type. Each agent i also has a function $h : M_i \rightarrow X_i$ that gives his allocation as a function of his message. Message $(m_1, \dots, m_N) \in M$ *verifies* the choice rule if whenever each agent i accepts his message m_i , the resulting alternative $(h_1(m_1), \dots, h_N(m_N))$ is optimal. The oracle should be able to verify an optimal alternative in each state. The communication cost of agent i is identified with the size of his message space M_i . Note that any distributed communication protocol can be converted into this distributed verification by letting M_i consist of agent i ’s information sets over the terminal nodes of the communication protocol. Thus, distributed verification offers a lower bound on distributed communication.

We say that choice rule F satisfies the *Distributed Budget Equilibrium Revelation Property (DBERP)* if for any distributed protocol verifying the choice rule there exists a function $b_i : M_i \rightarrow 2^{X_i}$ such that for any distributed message $(m_1, \dots, m_N) \in M$, each agent i can construct his budget set $b_i(m_i) \subset X_i$ on the basis of his own message m_i so that budget equilibrium $(b_1(m_1), \dots, b_N(m_N), h_1(m_1), \dots, h_N(m_N))$ verifies allocation $(h_1(m_1), \dots, h_N(m_N))$. The difference from BERP is that each agent should be able to construct his budget set on the basis of the communication *he observes*. Still, the same argument as that behind Theorem 2 shows that DBERP holds for any intersection-monotonic choice rule: Letting m_i represent the set of agent i ’s types for which he accepts the message, he can construct his budget set $b_i(m_i) = \bigcap_{R_i \in m_i} L(h_i(m_i), R_i)$, and by intersection monotonicity the resulting budget equilibrium $(b_1(m_1), \dots, b_N(m_N), h_1(m_1), \dots, h_N(m_N))$ verifies allocation

³⁰An intermediate model, in which publicly broadcast messages are followed by agents privately choosing their allocations, has been considered in economics under the name “parametric communication” (Calsamiglia 1987).

$(h_1(m_1), \dots, h_N(m_N))$. Thus, for IM choice rules, any distributed communication must reveal to each agent his own budget set, in addition to his own allocation. The necessity of observing one's own budget set can be used to bound below the size of the agent's message space, and therefore his individual communication cost. The "hard" cases for distributed communication are the ones in which individual communication cost grows with the number of agents, e.g., as it does in the matching problem.

The distributed communication model outlined above still requires a "center" to maintain the consistency of communication observed by different agents. In a verification protocol, the "center" must verify that $(m_1, \dots, m_M) \in M$. In a communication protocol, the "center" could be interpreted as a "communication device," which receives private inputs messages from many agents and sends private output messages to many agents. We could rule out such "communication devices" and consider a more restricted model of *pairwise communication*, in which only private messages between two agents are allowed. A verification version of such pairwise communication is considered by Marschak and Reichelstein (1998), who find that a certain number of agents must then become "coordinators:" in addition to observing their own prices (as they must under DBERP), they also get involved in relaying prices between other agents. Thus, the restriction to pairwise communication creates some "communication overhead."

Note that with pairwise communication, it may sometimes make sense to employ agents who possess no private information themselves, but can serve as "communication devices" by aggregating and/or disaggregating messages. While this could only increase the aggregate communication cost, it would now be spread among more agents, possibly reducing individual communication costs. For a survey of the literature on organizations with an endogenous number of agents, see van Zandt (1998).

An even more restricted model is that of *network* communication, which allows only pairwise communication between agents who share an edge in a fixed network. For example, the network could be given by existing Internet links or organizational structure. Marschak and Reichelstein (1998, Section 4) and Feigenbaum et al. (2003) consider a special case in which the communication network is a tree.³¹ A simple lower bound on communication

³¹The restriction to communication on trees may be justified by a large "fixed cost" of communication

along a given edge in a tree can be obtained by letting each agent sharing the edge have all the information on his “side” of the tree (i.e., the subtree obtained by cutting the edge). Feigenbaum et al. (2003) use this approach to show that implementation of some budget-balanced incentive-compatible allocation rules in trees requires the communication cost of a large number of agents to grow proportionately to the total number of agents. For modern internet multicast transmissions involving millions of users, such communication would certainly be impractical.

5.3 Evaluation Costs

Even when the communication cost measured in bits or real numbers is low, it may be costly for agents to evaluate their preferences to send the required messages. The costs of preference evaluation was noted in the computer science literature (Parkes 2000), and was modeled in economics as a cost of “information acquisition.”³² While these costs have recently gained attention in the mechanism design literature (see Bergemann and Valimaki 2005), here we focus on identifying the evaluation costs of a given choice rule under the maintained assumption that agents are sincere.

Just as with the communication cost, we can bound below the evaluation cost of finding an optimal alternative by that of verifying that a given alternative is optimal, and obtain the latter by using minimally informative verifying messages. Intuitively, the less informative a message is, the lower is each agent’s cost of confirming that his preferences are consistent with the message. Thus, BERP and our characterization of minimally informative budget equilibria again prove useful.

For an illustration, consider the many-to-one matching problem described in Subsection 4.4. Suppose that each agent has a cost of “evaluating” a potential matching partner, without incurring which he does not know his preferences regarding matchings with this

links. This argument was used by Arrow (1974) to explain the prevalence of hierarchies in firms.

³²Note that such “evaluation costs” depend not just on how many bits are sent, but on which information agents are asked to report with these bits. E.g., in the example in Section 2, it may be easier for an agent to answer the question “is your valuation above or below 1.5” than “is your valuation an even or odd number,” even though each answer would require 1 bit.

partner. According to Proposition 4, minimally informative messages verifying a stable matching are equivalent to match-partitional equilibria. To verify such an equilibrium, each potential match has to be evaluated by at least one of the partners, hence verification of stability requires at least FW evaluations. (Note that this cost must be expended in any state, and not just in the worst case.)

We could also allow different agents to have different evaluation costs. To take a simple example, suppose now it is costly for firms to evaluate workers, while workers do not have any evaluation costs. To verify stability with minimal evaluation costs for firms, we need to use a match-partitional equilibrium in which firms' budget sets are minimal. In fact, when firms have substitutable preferences, all the firms' budget sets can be minimized at once, by choosing the stable match that is Pareto worst for the firms (which exists by Roth and Sotomayor (1990, Theorem 6.8)), and letting each firm's budget set include only the workers who strictly prefer it to their current employer, along with the firm's current workers.

This verification procedure gives a lower bound on the firms' evaluation costs, but this lower bound is in fact achieved by the Gale-Shapley deferred acceptance algorithm in which workers propose (Roth and Sotomayor 1990, Theorem 6.8).³³ In this algorithm, each firm evaluates the minimal number of workers needed to find a stable match. This achieves a tangible evaluation savings over full revelation: e.g., in the one-to-one matching problem in which firms' and workers' preferences are uniformly and independently drawn, it can be calculated that in the worker-proposing deferred acceptance algorithm, a firm in expectation evaluates $1/3$ of all workers. It follows from our verification-based lower bound that this is the minimal expected number of evaluations by each firm that is needed to find a stable match.

5.4 Privacy

One reason to avoid full revelation is to prevent agents from learning about each other's private information — a goal known as *privacy*. In the economic literature, privacy is often

³³If we wanted to minimize *workers'* evaluation costs, we would achieve this with the deferred acceptance algorithm in which firms propose.

needed to prevent self-interested agents from deviating in ways that exploit the revealed information (see, e.g., Myerson (1991, Section 6)). The computer science literature studies privacy as a goal in itself.

Observe that privacy would be maximized if agents could reveal their information privately to a trusted “mediator” (“communication device”) who would then announce an optimal outcome. Then agents would then learn nothing about other agents’ private information beyond the implemented outcome— a situation known as *full privacy*. In reality, however, trusted mediators may not be available, and the question is how much privacy could still be maintained.

Without a trusted mediator, privacy could be enhanced using private pairwise communication between agents. In fact, with sufficiently many agents, private pairwise communication can usually achieve full privacy. This fact has been exploited in a number of papers that implement correlated equilibrium and communication equilibrium without a trusted mediator, by designing a communication protocols that reveals to each agent only his own prescribed action but nothing else, to prevent him from deviating (e.g., Forges 1990).

Suppose now that all communication is public. (Equivalently, we may assume that each agent is concerned that the other $N - 1$ agents’ would collude to share all their observed private messages to infer information about the agent’s type.) Privacy that can be achieved in in such setting is known as “unconditional privacy.”³⁴

To bound below unconditional privacy, we can again consider verification with minimally informative messages. To have an example, consider the many-to-one matching model described in Subsection 4.4, and suppose that we want to minimize revelation of information about firms’ preferences over workers. This is done using a match-partitional budget equilibrium in which the firms’ budget sets are minimal. Recall from Subsection 5.3 that when firms have substitutable preferences, all of their budget sets can be minimized at once

³⁴We assume that agents are not computationally constrained. If they are, then privacy can be achieved even with public communication using “public-key cryptography.” The idea is that a publicly communicated key from agent 1 to agent 2 can be used by agent 2 to encrypt information with a one-to-one function that is simple to compute but very hard to invert without a matching key for inversion, which only agent 1 has. This is the method currently used to implement secure Internet transactions.

using the Gale-Shapley deferred acceptance algorithm in which workers propose. The algorithm reveals nothing about firms’ preferences except their minimal budget sets, and so it maximizes the firms’ privacy. More generally, our results imply that public communication usually cannot achieve “full privacy,” since it must reveal supporting budget sets in addition to the outcome to be implemented.³⁵

6 Further Issues

6.1 Communication versus Verification

We have used the verification cost as a lower bound on the communication cost. This raises the following questions:

- *How tight is the verification bound?*

There are cases in which the gap between the communication cost and the verification cost of an allocation problem measured in bits can be exponential—an example is given in Segal (2005, Example 3). (The gap is never more than exponential, because starting with a b -bit verification protocol, which has at most 2^b messages, we can check all the messages sequentially until one is found that is accepted by all agents, which would take at most $N \cdot 2^b$ bits.)

- *In which cases is the verification bound fairly tight?*

³⁵Brandt and Sandholm (2005) show that with public communication and unrestricted preference domains, full privacy is not achievable for a large class of choice functions. However, with a restricted preference domain, full privacy may be achievable. For example, a Pareto efficient allocation x in a smooth convex economy can be verified by announcing a supporting Walrasian equilibrium. The supporting prices only reveal other agents’ marginal rates of substitution at x , which each agent would have learned from the allocation x itself by calculating his own marginal rates of substitution at x . Another question is whether “full privacy” is a relevant goal when we are implementing a choice rule that is a correspondence rather than a function. In this setting, so the revelation of information depends on *which* alternative x is implemented as a function of the state, and not just on what is revealed in addition to x .

The bound is trivially tight when even verification proves almost as hard as full revelation (e.g., in the combinatorial allocation problem considered in Subsection 4.2). More interestingly, there are some well-known social choice problems in which the gap between verification and communication proves to be small and both are much easier than that full revelation. For example, in the many-to-one matching problem considered in Subsection 4.4 in which firms have strict substitutable preferences, the Gale-Shapley deferred acceptance algorithm converges quickly to a “match-partitional” equilibrium, which verifies stability using only slightly more bits than that needed for verification, and exponentially less than that needed for full revelation of preferences. Similarly, in a convex economy with the “gross substitute” property, Walrasian tatonnement converges quickly to a Walrasian equilibrium, which verifies Pareto efficiency (Mas-Colell et al. (1995, Section 17.H)). Similar “tatonnement” mechanisms have been proposed for combinatorial auction problems with indivisible goods in which the objects are “substitutes” (e.g., Gul and Stachetti 2000, Nisan and Segal 2004). In all 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).

- *What is the role of price queries in communication?*

Many practical mechanisms, such as the ones mentioned in the above paragraph, are “demand-query protocols”: they quote to the agents a price list for the allocations (with prices sometimes allowed to be nonlinear and personalized) and ask them to submit demands given the prices, adjusting the prices according to some prespecified rules. Can we always restrict attention to such demand-query protocols without increasing the communication cost substantially? Nisan and Segal (2005) show that the answer is “no,” by constructing an allocation problem for which the restriction to demand-query protocols brings about an exponential blowup in the communication cost of finding an efficient allocation. Namely, for this class, an efficient mechanism exists that uses a number of bits that is linear the

number of items, but any demand query mechanism that achieves efficiency (or even any improvement upon the “dictatorial” allocation of all the items to one agent) must use an exponential number of demand queries. Contrast this to the verification problem, in which, according to Proposition 2, we can restrict attention to a demand-query mechanism (valuation equilibrium) without any increase in the communication cost.

To summarize, in several well-known cases the verification lower bound on communication is fairly tight, and efficient communication can be achieved with a demand-query mechanism. However, there are some problems in which these properties fail. It would be interesting to characterize social choice problems that satisfy both properties.

6.2 Incentives

So far we have ignored agents’ incentives to follow the strategies prescribed by the protocol. If the agents behave in their self-interest, the designer faces additional “incentive-compatibility” constraints requiring that no agent has an incentive to deviate from his prescribed strategy—i.e., the strategies constitute an equilibrium of the communication game. The number of bits by which these constraints increase the communication cost may be called the “communication cost of selfishness,” and it is examined in Fadel and Segal (2005, henceforth FS).

Note that the fact that the protocol must reveal supporting prices (by the Budget Equilibrium Revelation Property) does not ensure that it is incentive-compatible: agents may have the ability to manipulate the prices they face to their advantage. For example, take the setting of Section 2, in which one object is to be allocated between two agents, and consider Protocol 2, in which agent 1 announces his valuation v_1 , and agent 2 then announces an efficient allocation x . The protocol reveals a supporting price $p = v_1$. However, if agent 1 is charged this price for winning the object, then he will have an incentive to understate his valuation. In fact, as shown in FS, the protocol does not reveal enough information to compute a price that would motivate agent 1 to be truthful regardless of his beliefs about agent 2’s valuation v_2 . (Intuitively, when both agents’ valuations are in $[0,1]$, agent 1 can only be motivated to be truthful if he is charged price v_2 , as in the Vickrey auction, but this price is not revealed by the protocol.)

An agent’s incentive to deviate in a protocol depends on his information about the other

agents. FS consider two implementation concepts: Bayesian-Nash Incentive Compatibility (BIC), which requires that each agent has no incentive to deviate given his beliefs about other agents' types, and Ex Post Incentive Compatibility (EPIC), which requires that each agent has no incentive to deviate *regardless* of his beliefs about others' types. Both implementation concepts satisfy the Revelation Principle: if an allocation rule is implementable in *some* protocol, it is implementable in a *direct revelation* protocol, in which agents simultaneously announce their private information (but which may have a high communication cost). Thus, FS consider the communication cost of selfishness for those allocation rules that are implementable in a direct revelation protocol.

In general, agents' incentives in a protocol can be manipulated using two instruments: (1) monetary transfers (the agents' utilities are assumed to be quasilinear in such transfers, as in Subsection 4.2), and (2) information sets that hide information from the agents. For EPIC implementation, the protocol need not hide any information from the agents, and the communication cost of selfishness is entirely due to the need of computing motivating transfers in addition to the nonmonetary allocation. In contrast, for BIC implementation, the cost of selfishness is due to the need to hide information from the agents to restrict their contingent deviations (while computation of transfers does not entail any additional cost).

For both the EPIC and BIC case, FS provide an upper bound on the communication cost of selfishness:

$$\text{Incentive-Compatible Communication Complexity} \leq 2^{\text{Communication Complexity}}.$$

Since this bound is very weak, FS proceed to ask whether it is ever achieved or approached.³⁶

For BIC implementation, FS do show that the bound is tight, by providing an example in which the communication cost of selfishness is exponential. The example has two agents: An "expert" with private knowledge and a private utility function, and a "manager" with a privately known goal that determines how the expert's knowledge should be used. The expert will reveal his knowledge truthfully if he does not know how the manager's goal, but this revelation will take exponential communication in the number of outcomes. Communication

³⁶If the communication cost is measured as the *average-case* number of bits sent, as defined in footnote 3, FS show that the communication cost of selfishness can be unbounded, both for EPIC and for BIC.

could be reduced exponentially by having the manager first announce his goal and then letting the expert say how to achieve it, but this communication is not incentive-compatible — knowing the manager’s goal, the expert can manipulate her report to achieve her preferred outcome. FS show that any communication that satisfies the expert’s BIC constraints must be almost as long as full revelation of the expert’s knowledge.

For the EPIC case, it is not known whether the exponential upper bound is ever achieved or approached. In many studied cases, the communication cost of selfishness for EPIC proves to be low. For example, this is the case if we want to implement an efficient (surplus-maximizing) allocation.³⁷ Indeed, suppose that we have a communication protocol that finds an efficient allocation. After running the protocol, ask each agent to report his payoff $\pi_i = u_i(k)$ at the resulting allocation k , and pay each agent i a transfer $t_i = \sum_{j \neq i} \pi_j$.³⁸ Under this transfer scheme (first proposed by Reichelstein (1984)), each agent’s total payoff equals to the total surplus, and so the communication game becomes one of common interest (in the terminology of Marschak and Radner (1972), the agents become a “team”). Since the protocol is efficient, the resulting mechanism is EPIC: no deviation by an agent can increase the total surplus.³⁹

³⁷This argument extends to allocation rules that maximize nonnegative affine combinations of agents’ utilities, since they can be interpreted as efficient rules upon rescaling the agents’ utilities and adding a fictitious agent. Lavi et al. (2003) show that in some important settings, any dominant-strategy implementable allocation rule must take this form.

³⁸Technically, this requires agents to communicate real numbers. If agents can only communicate bits but have real-valued utilities, they can report their rounded-off utilities, in which case the proposed transfer scheme would make the protocol approximately incentive-compatible.

³⁹Even if the protocol is not exactly efficient but maximizes expected surplus given some common-knowledge prior subject to a constraint on communication costs, the proposed strategy profile will satisfy BIC, since no agent would be able to increase expected surplus by deviating. Furthermore, if agents are also made to internalize the communication costs through ex post transfers, then they need to be given any protocol at all — the protocol that maximizes the expected surplus net of communication costs will emerge as a Bayesian-Nash equilibrium of the “free-form” game in which agents can send any messages and implement an allocation. To be sure, this argument relies heavily on the agents’ rationality—both individual (being able to calculate an optimal protocol) and collective (having a common prior and being able to coordinate on a protocol). But if agents are not fully rational, it is not clear how to model their incentives in the first

Another literature on incentive-compatible communication studies a “dual” question: instead of asking how much communication is needed to achieve a given goal, it asks how to maximize a given objective function subject to a fixed communication constraint. The objective is typically to maximize the profits of one of the agents subject to other agents’ participation constraints. See, e.g., Green and Laffont (1987), Melumad et al. (1992), and a recent survey by Mookherjee (2006).

6.3 Average-Case Goals: Prices vs. Authority and Coercion

We have examined the problem of achieving a given social goal with certainty. However, given a probability distribution over states, we could allow probabilistic goals—e.g., require approximating the *probability* of finding an efficient outcome, or *expected* surplus.⁴⁰ Is it still necessary or desirable to find supporting prices to achieve such approximation? We show that the answer is “no,” by giving two examples in which (a) an efficient outcome can be found with a high probability with little or no communication, while (b) verifying efficiency of the outcome by describing supporting prices would require enormously more communication.⁴¹ In one example, the low-communication approximately efficient mechanism can be interpreted as coercion, and in the other, as authority.

Example 5 (Coercion): We need to decide whether to provide an indivisible public good to N agents whose valuations u_i for the good are drawn i.i.d. from $\{0, 1\}$, with $\Pr\{u_i = 1\} = \rho \in (0, 1)$. Let the cost of provision be between $k - 1$ and k , hence efficiency requires providing the good if and only if $\sum_i u_i \geq k$. Observe that when N is large and $k/N < \rho - \alpha$ for some fixed $\alpha > 0$, by the Law of Large Numbers, providing the good without any communication is efficient with a high probability. On the other hand, to *verify* that provision is efficient, by Proposition 2 we need to

place.

⁴⁰In contrast, to, say, approximating the maximum surplus within ε across all states, which was considered in Subsection 4.3.

⁴¹While for simplicity we show this for worst-case number of bits, the same results extend to the *expected* number of bits using Shannon’s (1948) entropy lower bound.

describe a supporting valuation (“Lindahl”) equilibrium, i.e., describe k agents willing to pay price 1 for the good. The probability that any such valuation equilibrium is indeed an equilibrium is ρ^k . Therefore, to find supporting prices with probability $\varepsilon > 0$ we need to use at least ε/ρ^k different equilibria, which requires sending at least $\log_2(\varepsilon/\rho^k) = k \log_2 \rho^{-1} + \log_2 \varepsilon$. Thus, as $N, k \rightarrow \infty$ so that $k/N < \rho - \alpha$, finding a supporting price equilibrium to verify efficiency with any fixed probability ε requires unbounded communication, while providing the good without any communication is efficient with probability approaching 1.

Example 6 (Authority): Two agents have utilities in $\{0, 1\}$ for allocations from set K .

The probability distribution is as follows: Each agent for each allocation draws utility 1 with probability ρ_K and 0 with probability $1 - \rho_K$, and the draws are independent across allocations and between the agents. Assume that as $K \rightarrow \infty$, (i) $\rho_K K \rightarrow \infty$, and (ii) $\rho_K^2 K \rightarrow 0$. By (i), the asymptotic probability that there is no surplus-1 allocation for an agent is $(1 - \rho_K)^K \sim e^{-\rho_K K} \rightarrow 0$. By (ii), the asymptotic probability that there is no surplus-2 allocation is $(1 - \rho_K^2)^K \sim e^{-\rho_K^2 K} \rightarrow 1$. Thus, the “authority protocol” in which one agent names the best allocation for him achieves efficiency with probability approaching 1, communicating only $\log_2 K$ bits. On the other hand, to *verify* that there is no allocation with a higher surplus, by Proposition 2 we need to announce a supporting valuation equilibrium. The probability that a given valuation equilibrium is an equilibrium conditional on the random state having maximal surplus 1 (which asymptotically occurs with probability 1) can be bounded above by $(1 - \rho_K)^K \sim e^{-\rho_K K}$.⁴² Thus, any protocol announcing a supporting price equilibrium with a fixed probability $\varepsilon > 0$ must asymptotically use at least $\varepsilon e^{\rho_K K}$ distinct messages, and so communicate $\log_2(\varepsilon e^{\rho_K K}) \sim \rho_K K \log_2 e$ bits. This communication cost could

⁴²To see this, recall first that in the binary-utility setting we could use the valuation equilibria (p_1, p_2, k) with prices $p_1, p_2 \in \{0, 1\}^K$, normalized so that $p_1, p_2 \neq (1, \dots, 1)$, and $p_1(k') + p_2(k') = 1$ for all $k' \in K$. For such (p_1, p_2, k) to be an equilibrium, the agents’ equilibrium utilities $u_i(k) - p_i(k)$ must be nonnegative, and therefore, in a surplus-1 state, both agents’ equilibrium utilities must be zero. This can only be an equilibrium in states (u_1, u_2) in which for all k' with $p_1(k') = 0$, $u_1(k') = 0$, and for all other k' , $u_2(k') = p_2(k') = 0$. This implies the upper bound.

be exponentially higher than the $\log_2 K$ bits used by authority (e.g., when $\rho_K = K^{-\alpha}$ with $\alpha \in (1/2, 1)$, which satisfies (i),(ii)).⁴³

Example 5 may be interpreted as justifying government provision of public goods when the provision is likely to be efficient, but the communication cost of using Lindahl markets with a large number of agents would be prohibitive. Example 6 may be interpreted as formalizing the view of Coase (1937) and Simon (1951) of firms as “islands of conscious power” in which the price mechanism is superseded by decision-making by authority. In the example, as suggested by Coase, the cost of “discovering what the relevant prices are” proves to be prohibitively high, while the benefit is vanishingly small.

The recent work on understanding the allocation of authority in firms (e.g., Aghion and Tirole 1997, Dessein 2002) has arbitrarily restricted attention mechanisms that allocate formal authority, accompanied by more extensive informal communication. If incentives were the only concern, then it would be optimal to use an extensive formal mechanism. Example 6 offers a potential explanation for the use of formal authority: if the costs of formal communication are higher than that of informal, it could be optimal to use only extremely simple formal communication such as authority, supplemented with extensive informal communication.

6.4 Interdependent Values

We have assumed that each agent knows his own preferences, which are not affected by other agents’ private information except through the implemented allocation. A more gen-

⁴³In the same setting, Nisan and Segal (forth., Proposition 14) show that there exists a probability distribution over the two agents’ binary utilities for which a surplus-2 allocation is guaranteed to exist, but finding it requires exponential communication in K , and so authority is optimal among subexponential mechanisms. Nisan and Segal (forth.) apply this result to showing the uselessness of practical combinatorial auctions, where “authority” allocation is achievable by giving all the objects to one agent. A shortcoming of this example is that the probability distribution over utilities needed for it to hold may not be a “natural” one. Segal (1995) obtained the same result for the probability distribution over utilities described as in the example, but under the restriction that communication cannot use a common “language” (labeling of allocations).

eral formulation would allow *interdependent values*, i.e., direct dependence of one agent’s preferences on other agents’ private information. One example is when other agents have private information about the quality of the goods allocated to the agent. Another example is when an agent is acquiring assets for future resale, and other agents have relevant private information for predicting the future resale price of the assets.

The performance of price mechanisms in such interdependent-value settings has been extensively studied. The most widely used price equilibrium concept for such settings is *Rational Expectations Equilibrium* (REE), in which agents infer information through the announced prices, and make choices from their budget sets to maximize their expected utilities given the inferred information (see, e.g., Mas-Colell et al. 1995 Section 19.H, Radner 1979, Grossman 1981). Can we offer a normative foundation for rational expectations price equilibria in the interdependent-value setting akin to the Budget Equilibrium Revelation Property for the private-value setting?

Note that in the interdependent-value setting, social goals such as Pareto efficiency may be achieved without revealing supporting REE prices. The simplest example is allocating an object among agents who have a “pure common value” for it, which depends on the agents’ private signals. In this example, any allocation would be efficient, and so could be achieved without any communication, but the REE would typically depend on private information.

One may argue that statistically efficient aggregation of private information may be desirable for reasons other than allocational efficiency (e.g., to guide investment decisions). Thus, many papers have examined the validity of the (strong form) of the “Efficient Market Hypothesis,” which says that REE prices form a sufficient statistic for the value of a security given all the private information. Contrary to the hypothesis, there exist cases in which an REE reveals *no* information about the value of a security, even though pooling agents’ private information would reveal the value fully:

- **Example 7 (Feigenbaum et al. (2005)):** There are two risk-neutral agents, each of whom privately observes a fair coin toss. The agent can trade a security whose value is 1 if the two agents’ coins fall on the same side and 0 otherwise.⁴⁴ There exists an

⁴⁴This construction is known in game theory as a “jointly controlled lottery.” For an economic example, let

REE with price $1/2$ that does not depend on the agents' private information. Since the price is uninformative, each agent continues to believe the security has value 1 with probability $1/2$, and so is willing to trade any amount at price $1/2$. On the other hand, if pooling both agents' information would reveal the exact value of the security.

One might argue that when agents' type spaces are finite, a continuous price would “generically” be fully revealing. However, when agents' type spaces are continuous, “generically” prices cannot be a sufficient statistic for private signals if the total dimension of the signals exceeds the dimension of the price space (which is realistic when agents observe complex signals or when the number of agents is large). Formally, we are facing a communication (verification) problem, whose solution may require a large message space than the available price space (a formal point along these lines is made by Jordan (1983)).

These arguments bring into question the recent popularity of “prediction markets” as means of aggregating dispersed private information to forecast various events, from sales at Hewlett Packard to election outcomes to terrorist attacks (see, e.g., Wolfers and Zitzewitz 2004). While some special communication problems may be solved efficiently with a prediction market, the general applicability of price mechanisms for aggregating common-value information is unclear. In particular, the recent proposals to use prediction markets to replace the managerial task of information aggregation and decision making⁴⁵ do not have a theoretical foundation.

agent 1 be the marketing manager of an auto company, who knows which car body will be in high demand next year, and let agent 2 be the company's manufacturing manager, who knows which car body will be cheap to produce next year. The security is contingent on the company's profits.

⁴⁵E.g., “With employees in the trading pits betting on the future, who needs the manager in the corner office?” – Times Magazine (2004).

7 Conclusion

In the past 30 years, economists have focused on the issue of incentives.⁴⁶ However, consider a thought experiment in which everybody is honest, and ask whether the fundamental economic institutions, such as markets and firms, would still be recognizable in this hypothetical world. It is our conjecture that the answer is “yes”: The primary function of these institutions is to process information and make decisions, and their fundamental features are explained by this function (even though incentives may be important for understanding many of their aspects).

This chapter has focused on one kind of economic institutions—price-based mechanisms. We have shown that, contrary to widespread belief, prices are needed not in order to incentivize the agents, but in order to aggregate distributed information about their preferences into a socially desirable decision. Thus, we have provided a justification for and characterized the scope of the price-based “market design” approach (as opposed to more general mechanism design), and characterized the form of “prices” that must be discovered to solve a given social choice problem.

Some of our extensions also offer promising avenues for understanding non-price allocation mechanisms such as firms and governments. For example, as noted by Coase (1937) and Simon (1951), communication in firms differs fundamentally from that in markets: Decisions in firms are usually made by the authority of managers, without “discovering what the relevant prices are.” We indeed find an example where authority may emerge as an optimal communication mechanism (Example 6 in Subsection 6.3): it finds an efficient allocation with a high probability, while the communication cost of verifying this efficiency by describing prices for all possible allocations is exponentially higher. Another notable aspect of communication in firms is that much of it is done by professional managers who specialize in aggregating information and making decisions. In Subsection 5.2 we noted how hiring such managers may economize on individual communication costs. Thus, while “theories

⁴⁶For example, consider the statements “Most of economics can be summarized in four words: ‘People respond to incentives.’ The rest is commentary” (Landsburg 1993) and “Economics is, at root, the study of incentives” (Levitt and Dubner 2005).

of the firm” based on incentives or incomplete contracts take managerial tasks as given, a theory based on communication may explain what it is that managers actually do.⁴⁷

References

- [1] Abelson, H. (1980): “Lower bounds on information transfer in distributed computations.” *Journal of the Association for Computer Machinery*, 27, 384-392.
- [2] Aghion, Philippe, and Jean Tirole (1997): “Formal and Real Authority in Organizations,” *Journal of Political Economy*, 105, 1-29.
- [3] Athey, S., and I. Segal (2005): “Efficient Dynamic Mechanisms,” working paper, Stanford University
- [4] Babaioff, Moshe, and Liad Blumrosen (2004): “Computationally Feasible Truthful Auctions for Convex Bundles.” *Proceedings of the 8th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, Lecture Notes in Computer Science, Springer-Verlag.
- [5] Bergemann, D., and J. Valimaki (2005): “Information Acquisition in Mechanism Design.” *Advances in Economic Theory: 9th World Congress*, forthcoming.
- [6] Bikhchandani, Sushil, and John Mamer (1997): “Competitive Equilibrium in an Exchange Economy with Indivisibilities,” *Journal of Economic Theory* 74, 385-413.
- [7] Bikhchandani, Sushil, and Joel Ostroy (2002): “The Package Assignment Model,” *Journal of Economic Theory* 107, 377-406.
- [8] Brandt, F., and T. Sandholm (2005): “Unconditional Privacy in Social Choice,” In Ron van der Meyden, ed., *Proceedings of Theoretical Aspects of Rationality and Knowledge X*, National University of Singapore, pp. 207-218.

⁴⁷Such a theory may complement existing models in which managers are hired to perform computations, such as addition of numbers — see, e.g., van Zandt (1998).

- [9] Calsamiglia, X. (1977): “Decentralized Resource Allocation and Increasing Returns,” *Journal of Economic Theory*, 14, 262-283.
- [10] Calsamiglia, X. (1987): “Informational Requirements of Parametric Resource Allocation Processes.” In: Theodore Groves, Roy Radner, and Stanley Reiter, eds., *Information, Incentives, and Economic Mechanisms*, Minneapolis: University of Minnesota Press.
- [11] Coase, R. (1937): “The Nature of the Firm,” *Economica*, 4, 386-405.
- [12] Conitzer, Vincent, and Tuomas Sandholm (2005). “Communication Complexity of Common Voting Rules.” In *Proceedings of the ACM Conference on Electronic Commerce 2005*.
- [13] Cramton, P., Y. Shoham, and R. Steinberg (eds.) (2006): *Combinatorial Auctions*, MIT Press.
- [14] Debreu, G. (1952): “A Social Equilibrium Existence Theorem”, *Proceedings of the National Academy of Sciences* 38(10), 886-893.
- [15] Debreu, G. (1983): “Neighboring Economic Agents,” in *Mathematical Economics: Twenty Papers of Gerard Debreu*, New York: Cambridge University Press, 173-178
- [16] Dessein, W. (2002): ‘Authority and Communication in Organizations.’ *Review of Economic Studies*, 69(4), pp. 811-38.
- [17] Dobzinski, S., N. Nisan, and M. Schapira (2005). Approximation Algorithms for Combinatorial Auctions with Complement-free Bidders. *Symposium on Theory of Computing* 2005.
- [18] Echenique, F. (2005): “Counting Combinatorial Choice Rules,” *Games and Economic Behavior*, forthcoming.
- [19] Edgar, G.E. (1990): *Measure, Topology, and Fractal Geometry*. New York: Springer-Verlag.

- [20] Fadel, Ronald, and Ilya Segal (2005a): “Communication Cost of Selfishness: Ex Post Implementation,” In Ron van der Meyden, ed., *Proceedings of Theoretical Aspects of Rationality and Knowledge X*, National University of Singapore, pp. 165-176.
- [21] Fadel, R., and I. Segal (2005b): “Communication Cost of Selfishness,” working paper, Stanford University
- [22] Feigenbaum, J., L. Fortnow, D. Pennock, and R. Sami (2005): “Computation in a Distributed Information Market,” *Theoretical Computer Science*, 343, 114-132.
- [23] Feigenbaum, J., A. Krishnamurthy, R. Sami, and S. Shenker (2003): “Hardness Results for Multicast Cost Sharing,” *Theoretical Computer Science*, 304, 215-236.
- [24] Forges, F. (1990), “Universal Mechanisms,” *Econometrica*, 58, 1341-1364.
- [25] Green, J. and Laffont J. (1987): “Limited Communication and Incentive Compatibility”, *Information, Incentives, and Economic Mechanisms: Essays in honor of Leonid Hurwicz*, T. Groves, Radner, R. and Reiter, S. (ed.), Minneapolis: University of Minnesota Press
- [26] Greenberg, J. (1990): *The Theory of Social Situations: An Alternative Game-Theoretic Approach*. Cambridge: Cambridge University Press.
- [27] Grossman, S. (1981): “An Introduction to the Theory of Rational Expectations under Asymmetric Information.” *Review of Economic Studies* 48, pp. 541-559.
- [28] Gusfield, D., and R.W. Irving (1989). *The Stable Marriage Problem: Structure and Algorithms*. Cambridge: MIT Press.
- [29] Hammond, P. (1997): “Game Forms versus Social Choice Rules as Models of Rights,” in K.J. Arrow, A.K. Sen, and K. Suzumura (eds.) *Social Choice Re-examined*, Vol. II (IEA Conference Volume No. 117) ch. 11, 82-95, London: Macmillan.
- [30] Hatfield, J.F., and P.R. Milgrom (2005): “Matching with Contracts,” *American Economic Review* 95, 913-935.

- [31] Hayek, F.A. (1945): “The Use of Knowledge in Society,” *American Economic Review* 35, 519-30.
- [32] Hurwicz, L. (1977): “On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes,” in K.J. Arrow and L. Hurwicz, eds., *Studies in Resource Allocation Processes*, 413-424, New York: Cambridge University Press.
- [33] Hurwicz, L., and T. Marschak (2003a): “Finite allocation mechanisms: Approximate Walrasian versus approximate direct revelation,” *Economic Theory* 21, 545-572.
- [34] Hurwicz, L., and T. Marschak (2003b): “Comparing finite mechanisms,” *Economic Theory* 21 (2003), 783-841.
- [35] Hurwicz, L., and S. Reiter (2006): *Designing Economic Mechanisms*, Cambridge University Press.
- [36] Ishikida, T., and T. Marschak (1996): “Mechanisms That Efficiently Verify the Optimality of a Proposed Action,” *Economic Design* 2(1), 33-68.
- [37] Jordan, J.S. (1982): “The Competitive Allocation Process is Informationally Efficient Uniquely,” *Journal of Economic Theory* 28(1), 1-18.
- [38] Jordan, J.S. (1983), “On the Efficient Market Hypothesis,” *Journal of Economic Theory* 51(5), 1325-1343.
- [39] Ju, Biung-Ghi (2001). “Nash Implementation and Opportunity Equilibrium,” working paper, University of Kansas.
- [40] Karloff, H. (1991). *Linear Programming*. Basel: Birkhäuser Verlag.
- [41] Kelso, Alexander S. Jr., and Vincent P. Crawford (1982). “Job Matching, Coalition Formation, and Gross Substitutes.” *Econometrica* 50, 1483-1504.
- [42] Kushilevitz, E., and N. Nisan (1997). *Communication Complexity*. Cambridge University Press.

- [43] Landsburg, S.E. (1993). *The Armchair Economist: Economics and Everyday Life*. New York: The Free Press.
- [44] Levitt, S., and S. Dubner (2005). “Freakonomics: A Rogue Economist Explores the Hidden Side of Everything,” New York: Harper Collins.
- [45] Lipton, R. J., E. Markakis, E. Mossel, and A. Saberi (2004). “On Approximately Fair Allocations of Indivisible Goods,” *Proceedings of the 5th ACM Conference on Electronic Commerce*, pp,125-131.
- [46] Luo, Z.-Q., and J.N. Tsitsiklis, Communication Complexity of Algebraic Computation,” *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science*, 758-765, 1991.
- [47] Marschak, T., and S.Reichelstein (1998): “Network Mechanisms, Informational Efficiency, and Hierarchies,” *Journal of Economic Theory*, 79, 106-141.
- [48] Mas-Colell, A. (1980): “Efficiency and Decentralization in the Pure Theory of Public Goods,” *Quarterly Journal of Economics* 94, 625-641.
- [49] Mas-Colell, A., M.D. Whinston, and J. Green (1995). *Microeconomic Theory*. New York: Oxford University Press.
- [50] Maskin, E. (1999): “Nash Equilibrium and Welfare Optimality,” *Review of Economic Studies* 66, 23-38.
- [51] Melumad, N., D. Mookherjee and S. Reichelstein (1992): “A Theory of Responsibility Centers,” *Journal of Accounting and Economics*, 15, 445-484.
- [52] Milleron, J.-C. (1972): “Theory of Value with Public Goods: A Survey Article,” *Journal of Economic Theory* 5, 419-477.
- [53] Miyagawa, E. (2002): “Reduced-Form Implementation,” Columbia University Working Paper.
- [54] Mookherjee, D. (2006): “Decentralization, Hierarchies and Incentives: A Mechanism Design Perspective,” *Journal of Economic Literature*, forthcoming.

- [55] Mount, K., and S. Reiter (1974): “The Information Size of Message Spaces,” *Journal of Economic Theory* 28, 1-18.
- [56] Myerson, R.B. (1991): *Game Theory: Analysis of Conflict*. Cambridge: Harvard University Press.
- [57] Nisan, N., and I. Segal (2004): “The Communication Requirements of Efficient Allocations and Supporting Prices,” forthcoming, *Journal of Economic Theory*.
- [58] Nisan, N., and I. Segal (2005): “Exponential Communication Inefficiency of Demand Queries,” In Ron van der Meyden, ed., *Proceedings of Theoretical Aspects of Rationality and Knowledge X*, National University of Singapore, pp. 158-164.
- [59] Parkes, D.C. (2000): “Optimal Auction Design for Agents with Hard Valuation Problems,” *Agent Mediated Electronic Commerce* (IJCAI Workshop)
- [60] Parkes, D.C. (2002): “Price-Based Information Certificates for Minimal-Revelation Combinatorial Auctions,” in *Agent-Mediated Electronic Commerce IV*, Padget et al. (eds), LNAI 2531, 103-122, Springer-Verlag.
- [61] Radner, R. (1979): “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices.” *Econometrica*, 47, 655-678.
- [62] Reichelstein, S. (1984): “Incentive Compatibility and Informational Requirements,” *Journal of Economic Theory*, 34, 32-51.
- [63] Reichelstein, S., and S. Reiter (1988): “Game Forms with Minimal Message Spaces,” *Econometrica* 56(3), 661-692.
- [64] Roth, A.E., and M.A.O. Sotomayor (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge: Cambridge University Press.
- [65] Salanie, B. (1997): *The Economics of Contracts: A Primer*. Cambridge: MIT Press.
- [66] Sato, F. (1981). “On the Informational Size of Message Spaces for Resource Allocation Processes in Economies With Public Goods,” *Journal of Economic Theory*, 24, 48-69.

- [67] Segal, I. (1995): “Communication Complexity and Communication by Authority,” working paper
- [68] Segal, I. (2005): “Communication Requirements of Social Choice Rules and Supporting Budget Sets,” working paper, Stanford University.
- [69] Sen, A.K. (1970): “The Impossibility of a Paretian Liberal,” *Journal of Political Economy* 78, 152-157.
- [70] Serrano, R. and O. Volij (2000): “Walrasian Allocations without Price-Taking Behavior,” *Journal of Economic Theory* 95, 79-106.
- [71] Simon, H. (1951): “A Formal Theory of the Employment Relationship”, *Econometrica* 19, 293-305.
- [72] Time Magazine (2004), “The End of Management?” by Barbara Kiviat, July 6.
- [73] van Zandt, T. (1998): “Organizations that Process Information with an Endogenous Number of Agents,” in Mukul Majumdar, ed., *Organizations with Incomplete Information*, Cambridge: Cambridge University Press. Chapter 7, pages 239-305.
- [74] Wolfers, J., and E. Zitzewitz (2004): “Prediction Markets,” *Journal of Economic Perspectives* 18, 107-126.
- [75] Walker, M. (1977): On the Informational Size of Message Spaces,” *Journal of Economic Theory* 15, 366-375.
- [76] Williams, S.R. (1986): “Realization and Nash Implementation: Two Aspects of Mechanism Design,” *Econometrica* 54, 139-152.




Figure 1.




Figure 2.




Figure 3.




Figure 4.

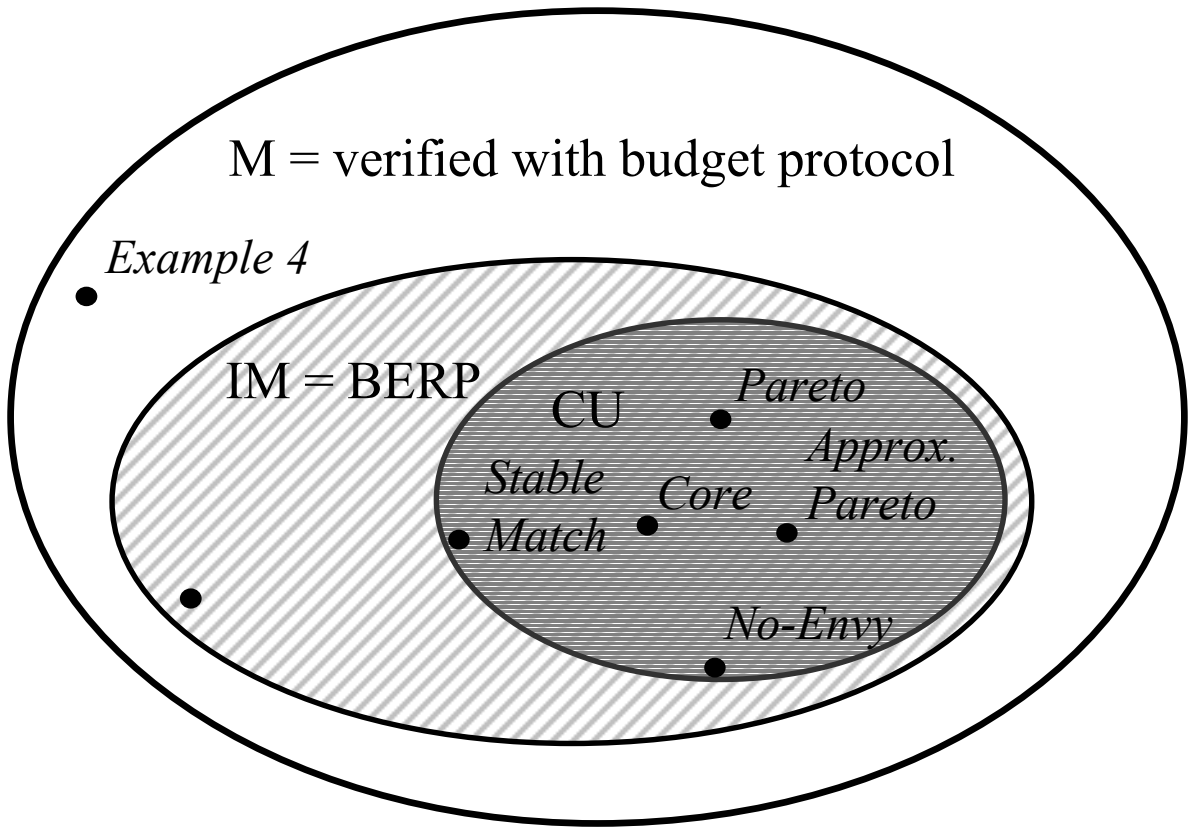


Figure 5.




Figure 6.




Figure 7.




Figure 8.




Figure 9.