



The communication requirements of efficient allocations and supporting prices[☆]

Noam Nisan^a, Ilya Segal^{b,*}

^a*Institute of Computer Science, Hebrew University, Jerusalem, Israel*

^b*Department of Economics, Stanford University, Stanford, CA 94305, USA*

Received 10 February 2004; final version received 12 October 2004

Available online 23 March 2005

Abstract

We show that any communication finding a value-maximizing allocation in a private-information economy must also discover supporting prices (in general personalized and nonlinear). In particular, to allocate L indivisible items between two agents, a price must be revealed for each of the $2^L - 1$ bundles. We prove that all monotonic prices for an agent must be used, hence exponential communication in L is needed. Furthermore, exponential communication is needed just to ensure a higher share of surplus than that realized by auctioning all items as a bundle, or even a higher *expected* surplus (for some probability distribution over valuations). When the utilities are submodular, efficiency still requires exponential communication (and fully polynomial approximation is impossible). When the items are identical, arbitrarily good approximation is obtained with exponentially less communication than exact efficiency.

© 2005 Elsevier Inc. All rights reserved.

JEL classification: C61; D02; D44

Keywords: Communication complexity; Message space dimension; Preference elicitation; Informational efficiency of prices; Combinatorial auctions; Submodular valuations; Substitutable items; Homogenous items; Approximation; Distributional complexity

[☆] An earlier version of this paper was circulated under the title “The Communication Complexity of Efficient Allocation Problems.”

* Corresponding author. Fax: +1 650 725 5702.

E-mail addresses: noam@cs.huji.ac.il (N. Nisan), ilya.segal@stanford.edu (I. Segal).

1. Introduction

We have recently seen great interest in so-called *combinatorial auctions*, designed to allocate L heterogeneous indivisible items among N bidders whose valuations for the different items can be interdependent. Recent important applications include auctions of FCC spectrum licenses and online procurement (see [47] for an overview). The objective of an auction is to elicit enough information about bidders' preferences so as to realize an efficient or approximately efficient allocation. The mechanism design literature has used the Revelation Principle to ensure the bidders' incentives to reveal their preferences truthfully (e.g., using the Vickrey–Groves–Clarke transfers). However, full revelation of a bidder's preferences requires naming a willingness to pay for each of the $2^L - 1$ bundles of items. Already with $L = 30$, this would involve the communication of more than 1 billion numbers, which is beyond the capabilities of any human or machine.

Recognition of the communication problem has prompted researchers to propose mechanisms in which valuations need not be fully revealed. For example, in many proposed iterative auction designs, at each stage bidders only need to describe their preferences over a small number of bundles [4,9,17,26,36,38]. The hope is that such designs could achieve or at least approximate efficiency, while allowing bidders to communicate much less than their complete preferences. Is this hope justified?

We evaluate the communication requirements of allocation problems by establishing the indispensable role of prices in finding efficient allocations in economies with distributed information. Intellectual origins of this idea lie in the early 20th century debate on central planning alternatives to the market system. Hayek [15] argued that prices succinctly summarize the knowledge of “particular circumstances of time and place,” which is too enormous to be communicated to a central planner. Hurwicz [18] and Mount and Reiter [35] formalized Hayek's intuition by showing that in classical convex economies, the Walrasian price mechanism is “informationally efficient,” i.e., verifies Pareto efficient allocations with the least amount of communication (as measured by the dimensionality of the message space) among all mechanisms satisfying a continuity restriction. Furthermore, Jordan [21] showed that the Walrasian mechanism is a unique individually rational mechanism having this property. These results were extended to convex economies with public goods and externalities [42,44].

However, the existing literature on the “informational efficiency” of price equilibria does not apply to allocation problems in which the agents' preferences may be nonconvex and the goods may be indivisible, such as the combinatorial auction setting. First, the literature restricts attention to mechanisms that satisfy a continuity restriction, which precludes the communication of discrete allocations. It does not rule out the possibility that mechanisms that do not satisfy this restriction could achieve efficiency without revealing supporting prices. Second, the literature does not define a price equilibrium concept for allocation problems in which a linear-price equilibrium need not exist.¹ Finally, the literature only

¹ The literature on public-good economies (e.g., [42]) allows personalized (“Lindahl”) prices for the public goods, but still restricts them to be linear, and ensures equilibrium existence by assuming convex preferences over divisible public goods.

considers exact efficiency, and does not examine the potential communication savings from allowing approximation.

The present paper addresses these questions in an allocation problem with privately known preferences. We show that in any such problem, *any* communication mechanism that finds a surplus-maximizing allocation must also discover a price equilibrium supporting it, which in general involves nonlinear and personalized prices.² This holds even if the agents are truthful (follow the prescribed reporting strategies), thus the necessity of discovering prices has nothing to do with the problem of providing agents with appropriate incentives. The result also implies that the simpler “nondeterministic” problem of *verifying* the efficiency of a proposed allocation is exactly that of announcing supporting prices along with the allocation.

In the classical convex economy one could restrict attention to prices that are Walrasian (i.e., anonymous and linear in consumption), which are much easier to communicate than the agents’ preferences. In the combinatorial allocation problem, on the other hand, we show that a huge price space must be used in order to ensure equilibrium existence. Specifically, we show that any possible $(2^L - 1)$ -dimensional price vector listing the prices of all bundles of items is a unique equilibrium price vector for a given agent for some valuation profile. Since any efficient mechanism must communicate such a vector, it must use at least as much communication as a full revelation of one agent’s preferences.

Our approach can also be extended to the problem of *approximating* the maximum total surplus within a constant. For this purpose, note that in the “discretized” problem in which the valuations are restricted to be multiples of $\delta > 0$, any misallocation loses at least surplus δ . Therefore, approximating the maximum surplus within less than δ is at least as hard as realizing exact efficiency in the discretized problem. Since the discretized problem can always be solved with finite communication, the relevant measure of the communication burden is the number of transmitted bits. Such discrete problems have been examined in the computer science field of communication complexity, pioneered by Yao [49] and surveyed in [25].³ Since exact efficiency in the discretized problem still requires the communication of (discrete) prices, we are able to show that guaranteeing a better approximation of efficiency than that achieved by auctioning off all items as a bundle requires communicating a number of bits that grows exponentially with L .

The concept of approximation used above required *uniform* approximation of maximum surplus across all states. Similarly, the communication burden was defined as the *maximum* number of bits transmitted across states. Instead of using such “worst-case” measures, one may assume a probability distribution over possible valuations, and ask how many bits must be transmitted *on expectation* to realize a given *expected* surplus. We show that, for *some* joint probability distribution over the agents’ valuations, achieving a higher expected surplus

² A related result is obtained by Parkes [37]. He considers the combinatorial auction problem and shows the necessity of revealing supporting prices by those communication languages that verify efficiency while revealing so-called “outcome-independent” information. Our result is more general in that it does not impose any restrictions on the communication language.

³ The general communication complexity problem is to compute a function (in our case, efficient or approximately efficient allocation) whose inputs (in our case, agents’ utilities) are distributed among agents. For more detailed comparisons between the communication complexity literature and the economic literature on real-valued communication, see [31,45].

than that from the bundled auction still requires communicating an exponential expected number of bits.⁴

These results imply that the only hope to achieve or approximate efficiency without enormous communication is by focusing on cases in which the agents' preferences (or probability distribution over them) are known a priori to lie in a certain class. One example is given by valuations satisfying the “(gross) substitute property” of Kelso and Crawford [23]. With such valuations, a Walrasian equilibrium with L per item prices exists, and as we show here, it can be found with (truly) polynomial communication. However, the substitute property is very restrictive. We show that for the somewhat larger class of “submodular” valuations (i.e., those exhibiting diminishing marginal utility of items), efficiency still requires very extensive communication, and a fast (so-called “fully polynomial”) approximation is impossible.

Finally, we consider the case where the items are known to be homogeneous, and so agents only care about the number of items consumed. This case exhibits a drastic difference between the communication requirements of exact and approximate efficiency. Namely, exact efficiency again requires at least as much communication as a full description of one agent's preferences, which in this case takes L numbers. On the other hand, approximation within any given ε (more generally, fully polynomial approximation) is achieved with only $O(\log L)$ bits. In the setting considered by Calsamiglia [10], the homogeneous good to be allocated is *divisible*, and exact efficiency requires infinitely dimensional communication, yet we construct a fully polynomial approximation that allocates the good in small discrete units (provided that the valuations satisfy a weak smoothness condition). Thus, in this particular case an enormous savings in communication is achieved with only a slight sacrifice in economic efficiency.

The paper is organized as follows. In Section 2 we describe the general allocation problem, the model of communication, and the measures of the communication burden. In Section 3 we characterize efficient communication as that discovering a price equilibrium. In Section 4 we use this characterization to derive a lower bound on the burden of efficient communication. In Section 5 we define the concepts of approximation and relate them to the analysis of discretized problems. In Section 6 we apply the results to the combinatorial allocation problem. In Section 7 we examine the problem with several restricted classes of valuations. Section 8 presents average-case analysis. Section 9 discusses the relation of our results to the computational complexity literature. Section 10 discusses how agents could be given the incentives to obey the suggested communication protocols. Section 11 concludes.

2. The allocation problem and communication

2.1. The allocation problem

Let K be the set of allocations and N be the finite set of agents, with $N \geq 2$ to avoid the trivial case. (With a slight abuse of notation, we use the same letter to denote a set

⁴ We prove this particular result by reducing the problem from that of approximate set packing, which is analyzed in the appendix.

and its cardinality when this causes no confusion.) An agent's valuation assigns real values to all allocations, and is therefore represented with a vector in \mathbb{R}^K . The class of possible valuations of agent $i \in N$ is denoted by $U_i \subset \mathbb{R}^K$. Agent i 's valuation $u_i \in U_i$ is assumed to be his privately observed "type." A *state* is a valuation profile $(u_1, \dots, u_N) \in U \equiv U_1 \times \dots \times U_N \subset \mathbb{R}^{NK}$.

The objective is to implement an *allocation rule*, which is a nonempty-valued correspondence $F : U \rightarrow K$. For each state $u \in U$, the allocation rule describes the nonempty subset $F(u) \subset K$ of "desirable" allocations. For example, we may be interested in the *efficient* allocation rule, which selects the allocations maximizing the sum of the agents' valuations (*total surplus*):

$$F^*(u) = \underset{k \in K}{\text{Arg max}} \sum_{i \in N} u_{ik} \quad \text{for all } u \in U.$$

Identification of efficiency with surplus-maximization is based on the ability to compensate agents with monetary transfers, and the quasilinearity of their payoffs in these transfers; approximately efficient allocation rules will be defined in Section 5 below.

2.2. Communication

We now describe communication procedures that solve the allocation problem. It is well known that communication can be shortened by letting agents send messages sequentially rather than simultaneously. For example, an agent need not report his valuation for allocation k if previous messages have made it clear that k stands no chance of being efficient. Therefore, we must consider multi-stage communication protocols.

In the language of game theory, a general communication protocol is described with an extensive-form message game as well as each agent's strategy in this game (complete action plan contingent on his type and observed history). Instead of payoffs, the game assigns allocations to terminal nodes (and so is more properly called a "game form," or "mechanism"). The agents are assumed to follow the prescribed strategies (their incentives to do so will be discussed in Section 10). Such communication protocols are called "deterministic," because the message sent by an agent at a given information set is fully determined by his type and the preceding messages. A protocol *realizes* allocation rule F if in every state $u \in U$ it achieves a terminal node to which an allocation from $F(u)$ is assigned.

Dealing with deterministic communication protocols is quite cumbersome. Analysis can be simplified by considering what is known as "nondeterministic communication" in computer science and as the "verification scenario" in economics. Imagine an omniscient oracle who knows the true state u , and consequently knows a "desirable" allocation $k \in F(u)$, but has to prove to an ignorant outsider that k is indeed desirable. He carries out the proof 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 u_i . The set of messages acceptable to all agents in state u is described by the message correspondence $\mu(u)$. The acceptance of message m by all agents must prove to the outsider that allocation k is desirable. This communication is called "nondeterministic" because it does not

find out (determine) a desirable allocation, but simply verifies that a candidate allocation is desirable.⁵

A famous economic example of nondeterministic communication is the 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 allocated consumption bundle maximizes his utility under the announced prices. The classical Welfare Theorems say that Walrasian equilibria can be used to verify Pareto efficiency in convex exchange economies.

Although we are ultimately interested in deterministic communication, studying nondeterministic is useful because it is much simpler to analyze and yet still relevant, for the following reasons:

1. Any deterministic communication can be represented as nondeterministic by letting all the messages be sent by the oracle instead of the agents, and having each agent accept the message sequence if and only if all the messages sent in his stead are consistent with his strategy given his type. The oracle’s message space M is thus identified with the set of the protocol’s possible message sequences (terminal nodes), and the message correspondence μ is single-valued. Therefore, any lower bound on the communication requirements of nondeterministic protocols applies to deterministic protocols as a particular case.
2. A nondeterministic protocol realizing allocation rule F can often be used as a basis for an iterative *deterministic* protocol approximating F . At each stage of the iteration, a message $m \in M$ is announced, and each agent reports the direction in which the message should be adjusted to become “more acceptable” to him. Such adjustment processes approximating Walrasian equilibria are known as “tâtonnement.” Nondeterministic communication can then be viewed as a steady state of the deterministic adjustment process.

Formally, we define nondeterministic communication as follows:

Definition 1. A *nondeterministic communication protocol* is a triple $\Gamma = \langle M, \mu, h \rangle$, where M is the message set, $\mu : U \rightarrow M$ is the message correspondence, and $h : M \rightarrow K$ is the outcome function, and the message correspondence μ has the following properties:

- *Existence:* $\mu(u) \neq \emptyset$ for all $u \in U$,
- *Privacy Preservation:* $\mu(u) = \cap_i \mu_i(u_i)$ for all $u \in U$, where $\mu_i : U_i \rightarrow M$ for all $i \in N$.

Protocol Γ *realizes* allocation rule $F : U \rightarrow K$ if $h(\mu(u)) \subset F(u)$ for all $u \in U$.

Existence means that an acceptable message exists in each state. *Privacy Preservation* follows from the fact that each agent does not observe other agents’ types when making his acceptance decision, thus the set of messages acceptable to him is a function $\mu_i(u_i)$ of

⁵ For the sake of economists, we note that the computer science term “nondeterministic” used here does not mean randomness. We do not consider randomized communication protocols since they would be useless given the worst-case measure of the communication burden defined in the next subsection, or given the average-case measure against a given probability distribution used in Section 8.

his own type u_i only.⁶ Finally, the definition of realization says that the acceptance of a message m by all agents proves to the outsider that $h(m) \in F(u)$.

Definition 1 has a nice interpretation in terms of geometric properties of the subsets $\mu^{-1}(m)$ of the state space U , each such subset being the event in which a given message m occurs [25, Chapter 2]. In this interpretation, *Existence* requires that the collection $\{\mu^{-1}(m)\}_{m \in M}$ of such events cover the state space U . *Privacy Preservation* requires that each element of the covering be a product set $\mu_1^{-1}(m) \times \dots \times \mu_N^{-1}(m)$ —a “rectangle” in computer science parlance. Realization requires that for each rectangle $\mu^{-1}(m)$ from the covering, outcome $h(m)$ be “desirable” on the whole rectangle (so that, in computer science parlance, the rectangle must be “monochromatic”).

2.3. The burden of discrete and continuous communication

The deterministic (nondeterministic) communication burden of an allocation rule is defined as the minimum communication burden of a deterministic (respectively, nondeterministic) protocol realizing it. The communication burden of a protocol is naturally defined as the length of the realized message sequence, i.e., the number of messages sent in the course of the protocol. Since this number may vary across states, we focus on the “worst-case” communication burden—the maximum length of the message sequence over all states (though see Section 8 below for some average-case results). For this measure to be interesting, we must require all communication to be encoded with “elementary” messages, bounding the amount of information conveyed with an elementary message.

The computer science literature on communication complexity considers discrete communication, in which the elementary messages convey a bit of information [25].⁷ In particular, in the nondeterministic case, the minimal binary encoding of the oracle’s message from set M takes $\log |M|$ bits. (log will stand for the binary logarithm to simplify notation).

In the case of continuous communication, agents are allowed to send real-valued elementary messages. We also want to allow finite-valued messages (say, to communicate discrete combinatorial allocations), but not count them toward the communication burden. Thus, the worst-case burden of continuous communication is defined as the maximum number of real-valued elementary messages sent in the course of the protocol. In the nondeterministic case, we can identify the communication burden with the dimension of the oracle’s message space M , for which purpose we must define a topology on M .

A well-known problem in continuous communication is the possibility of “smuggling” multidimensional information in a one-dimensional message space (e.g., using the inverse Peano function). Note, however, that with such “smuggling,” a small error in the message would yield a huge error in its meaning. To avoid this, we define a metric on messages based on their meaning: the distance between messages m and m' is defined as the Hausdorff

⁶ This is an established, if perhaps unfortunate, term in the economic literature. Privacy Preservation reflects the fact that the agents do not observe each other’s types *initially*, but does not constrain the revelation of types in the course of communication.

⁷ This is merely a normalization, because an elementary message in *any* finite alphabet could be coded with a fixed number of bits.

distance between the events $\mu^{-1}(m)$ and $\mu^{-1}(m')$ in which they occur.⁸ The metric dimension of space M in this metric, denoted by $\dim M$, then serves as a measure of the communication burden. ($\dim M$ can be defined as the Hausdorff dimension, the box-counting dimension, or the packing index [11]—this would not affect the results in this paper.) With this definition, any coding of messages from M with d real numbers that prevents small errors in the code from drastically distorting the meaning (formally, the inverse of the code is Lipschitz continuous) must have $d \geq M$ [11, Exercise 6.1.9(1)]. Thus, $\dim M$ is the relevant measure of communication burden if the communication must be robust to small transmission errors, due either to analog noise or to discretization (“quantization”). (See Proposition 13 below for a formal result along these lines.)

Our approach to measuring continuous communication stands in contrast to the previous literature on continuous communication, in which the topology on the message space is taken as given, and dimension smuggling is ruled out by imposing a continuity restriction on the communication protocol [1,30,35,48]. For example, Mount and Reiter [35] and Walker [48] require the message correspondence μ be “locally threaded”—i.e., have a continuous selection on a neighborhood of any point. This requirement rules out a priori some important communication protocols, e.g., the communication of discrete allocations (μ cannot have a continuous selection in a neighborhood in which the allocation switches).

3. Efficient communication and supporting prices

One way to verify the efficiency of an allocation is by announcing supporting prices:

Definition 2. A pair $(p, k) \in \mathbb{R}^{NK} \times K$, where $k \in K$ is the proposed allocation and $p \in \mathbb{R}^{NK}$ is a list of personalized allocation prices, is a *price equilibrium* in state $u \in U$ if it satisfies the following inequalities:

$$u_{ik} - p_{ik} \geq u_{ik'} - p_{ik'} \quad \text{for all } i \in N, k' \in K. \tag{1}$$

$$\sum_{i \in N} p_{ik} \geq \sum_{i \in N} p_{ik'} \quad \text{for all } k' \in K. \tag{2}$$

Eq. (1) says that the proposed allocation maximizes each agent’s utility net of the announced prices. Eq. (2) can be interpreted as requiring that the proposed allocation maximize the designer’s revenue given the announced prices. Denote the set of price equilibria in state u by $E(u)$, and let $E : U \rightarrow \mathbb{R}^{NK} \times K$ denote the price equilibrium correspondence.

A price equilibrium with $N = 2$ agents is illustrated in Fig. 1. Without loss of generality, we normalize both agents’ utilities and prices for the equilibrium allocation k to zero, and graph agent 1’s valuations and prices for all allocations in the downward direction, and those of agent 2 in the upward direction. (The economic interpretation of Fig. 1 is as an “Edgeworth box” whose vertical dimension represents the split of money between the two

⁸ Formally, the distance is defined as $\rho_M(m, m') = \max \left\{ d_M(\mu^{-1}(m), \mu^{-1}(m')), d_M(\mu^{-1}(m'), \mu^{-1}(m)) \right\}$, where $d_M(A, B) = \sup_{R \in A} \inf_{R' \in B} \rho_U(R, R')$ for $A, B \subset U$, with ρ_U describing the given metric on the state space U .




Fig. 1. Price equilibrium.

agents, and whose horizontal dimension represents the allocations. In this interpretation, u_1 and u_2 depict the agents' respective indifference curves passing through the equilibrium point, and p_1 and p_2 depict the boundaries of their respective budget sets.) Eq. (2) says that the curve representing p_1 must lie above that representing p_2 , and (1) that the curve representing u_1 is above that for p_1 and the curve representing u_2 is below that for p_2 . Thus, the four curves must be ordered vertically as in Fig. 1.

This general notion of price equilibrium was first suggested by Mas-Colell [33] and later used by Bikhchandani and Mamer [7] and Bikhchandani and Ostroy [8] for combinatorial allocation problems. These papers have made the observation that all price equilibrium allocations are efficient, often referred to as the *First Welfare Theorem*. In Fig. 1, efficiency simply means that u_1 must lie above u_2 .

The First Welfare Theorem allows to verify efficiency with a *price protocol*, whose message space is $M \subset \mathbb{R}^{NK} \times K$, whose message correspondence is $\mu(u) = E(u) \cap M$, and whose outcome function is $h(p, k) = k$. Note that we allow many different price protocols described by different message spaces, i.e., different feasible price-allocation pairs. Any price protocol satisfies *Privacy Preservation* by construction, since each agent verifies his inequality in (1) using his own type only, and any agent can verify (2). If we find a message space M large enough to satisfy *Existence* (i.e., that a price equilibrium from M exists in all states), the size of M provides an upper bound on the burden of efficient nondeterministic communication. (Namely, it is $\log |M|$ bits for discrete communication, and $\dim M$ for continuous communication.)

One might think that among all possible efficient communication protocols, there would be some that are distinct from, and perhaps much simpler than, price protocols. However, we establish that this is not the case—any efficient communication protocol must reveal supporting equilibrium prices:⁹

Proposition 1. *Communication protocol $\Gamma = \langle M, \mu, h \rangle$ realizes the efficient allocation rule F^* if and only if there exists an assignment $p : M \rightarrow \mathbb{R}^{NK}$ of prices to messages such that protocol $\langle M, \mu, \langle p, h \rangle \rangle$ realizes the price equilibrium correspondence E .*

⁹ A similar result is established by Parkes [37], but only for a restricted communication language.

Proof. The “if” statement obtains by adding up, for any given allocation $k' \in K$, inequalities (1) for $i \in N$ and inequality (2) to obtain $\sum_{i \in N} u_{ik} \geq \sum_{i \in N} u_{ik'}$. For the “only if” statement, suppose protocol $\langle M, \mu, h \rangle$ realizes F^* . For each $m \in M$, let $\hat{k} = h(m)$, and let $p_{ik} = \sup_{u_i \in \mu_i^{-1}(m)} [u_{ik} - u_{i\hat{k}}]$ for all $i \in N, k \in K$. By construction, (p, \hat{k}) satisfies (1) for each $u \in \mu^{-1}(m)$. At the same time, for all $k \in K$ we can write

$$\sum_i [p_{ik} - p_{i\hat{k}}] = \sum_i \sup_{u_i \in \mu_i^{-1}(m)} [u_{ik} - u_{i\hat{k}}] = \sup_{u \in \mu^{-1}(m)} \sum_i [u_{ik} - u_{i\hat{k}}] \leq 0,$$

using *privacy preservation* (which means that $\mu^{-1}(m) = \mu_1^{-1}(m) \times \dots \times \mu_1^{-1}(m)$) and the fact that $\hat{k} \in F^*(u)$ for each $u \in \mu^{-1}(m)$. Thus, (p, \hat{k}) also satisfies (2). Therefore, (p, \hat{k}) is a price equilibrium in every state $u \in \mu^{-1}(m)$. \square

The “if” statement of the proposition is the First Welfare Theorem. The “only if” statement can be thought of as a strengthening of the traditional Second Welfare Theorem, which says only that for any efficient allocation we can construct supporting prices *given full information about the economy*. For general price equilibria, the construction is trivial: for example, we can simply take prices $p^i = u^i$ for all agents i . The Second Welfare Theorem is not a very useful result, because with full information, an efficient allocation can be implemented directly, without using prices. In contrast, Proposition 1 says not only that supporting prices exist, but also that they must be revealed by *any* efficient communication, not just by full revelation.

The proof of the “only if” part of Proposition 1 for two agents is illustrated in Fig. 2, which depicts the valuations in the same way as Fig. 1. Suppose that a message m has verified that allocation k is efficient. Consider the two agents’ valuations that are consistent with message m . Due to *Privacy Preservation*, k must be efficient in any state in which agent 1’s valuation is consistent with m and agent 2’s valuation is consistent with m . Graphically, this means that any valuation curve of agent 1 consistent with m (in Fig. 2, u_1 and u'_1) must lie above any valuation curve of agent 2 consistent with m (in Fig. 2, u_2 and u'_2). Therefore, letting agent 1’s price curve p_1 be the lower envelope of his valuation curves consistent with m , and letting agent 2’s price curve p_2 be the *upper* envelope of his valuation curves consistent with m , p_1 will lie above p_2 , thus the prices will satisfy condition (2). Also, by construction, the prices will satisfy condition (1) in all states consistent with m (states (u_1, u_2) , (u'_1, u_2) , (u_1, u'_2) , and (u'_1, u'_2) in Fig. 1). Thus, on the basis of a message m verifying k we have constructed a price equilibrium supporting k in all states in which m occurs.

Proposition 1 implies that the nondeterministic communication burden of the efficient allocation rule F^* is *exactly* the minimum size of a price-allocation space M ensuring the existence of an equilibrium from $M \subset \mathbb{R}^{NK} \times K$. While finding such a minimal space is in general a hard problem, below we derive some upper and lower bounds for it in specific cases (and in some of these cases the two bounds match).




Fig. 2. Revelation of prices.

4. A lower bound on efficient communication

Note that two utility functions that differ by a constant describe the same preferences. In counting the states of the world, we need to avoid double-counting such utility functions. Similarly, in measuring the dimension of (subsets of) the state space, we need to ensure that the dimensionality of utility functions coincides with the dimensionality of preferences. For this purpose, we assume that any two preferences that are “close” to each other are described by utility functions that are close to each other:

Definition 3. The state space U is *normalized* if there exists $C > 0$ such that for all $u, u' \in U$,¹⁰

$$\|u' - u\| \leq C \sup_{i \in N, j, k \in K} \left| [u'_{ij} - u_{ij}] - [u'_{ik} - u_{ik}] \right|.$$

In particular, a normalized state space U cannot contain two distinct states between which each agent’s utility differs only by a constant. For example, normalization holds (with $C = 1$) when each agent is always assigned zero utility to one of the allocations.

Now we can show that the subset of the state space on which all allocations are equally efficient constitutes a “fooling set” for the price equilibrium correspondence. Namely, letting

$$U^* = \{u \in U : F^*(u) = K\},$$

we show that two distinct states from U^* cannot have the same supporting prices:

¹⁰ For definiteness, we let $\|\cdot\|$ represent the sup-norm on \mathbb{R}^{NK} , though in a finite-dimensional space U all norms are known to be equivalent.

Lemma 1. *In a normalized allocation problem, for any $u, u' \in U^*$ with $u \neq u'$, $E(u) \cap E(u') = \emptyset$.*

Proof. Suppose in negation that $(p, \hat{k}) \in E(u) \cap E(u')$. $u \in U^*$ means that the total surplus $\sum_i u_{ik}$ does not depend on $k \in K$, and therefore (2) can be written in state u as

$$\sum_i [u_{i\hat{k}} - p_{i\hat{k}}] \leq \sum_i [u_{ik} - p_{ik}] \quad \text{for all } k \in K.$$

On the other hand, (1) means that

$$u_{i\hat{k}} - p_{i\hat{k}} \geq u_{ik} - p_{ik} \quad \text{for all } i \in N, k \in K.$$

Comparing the two displays, we see that for each i , $u_{ik} - p_{ik}$ does not depend on $k \in K$. Since by the same argument the same is true for $u'_{ik} - p_{ik}$, we see that $u'_{ik} - u_{ik}$ does not depend on k for all i . By normalization, this implies that $u = u'$ —a contradiction. \square

The proof of Lemma 1 for the case $N = 2$ can be illustrated using Fig. 1, by observing that in any state from U^* , the two agents' valuation curves coincide, and the price curves are squeezed in-between. Therefore, both agents' prices must coincide with their utilities (up to a constant). Different states from U^* will give rise to different indifference curves and will therefore pin down different price curves.

We can now use Proposition 1 to conclude that any efficient communication protocol must yield distinct messages in distinct states from U^* . This implies that the protocol uses at least $|U^*|$ distinct messages. A similar statement can be made for continuous communication, though extra care should be taken to rule out “dimension smuggling”:

Proposition 2. *In a normalized allocation problem, if U^* is a finite set, then any efficient protocol transmits at least $\log |U^*|$ bits, and otherwise the dimension of its message space is at least $\dim U^*$.*

Proof. Lemma 1 and Proposition 1 together imply that the restriction of the message correspondence μ to U^* has an injective selection $\sigma : U^* \rightarrow \mu(U^*)$. This implies the first statement of the proposition. For the second statement, we will also show that σ^{-1} is Lipschitz continuous. For this purpose, take any $m, m' \in \mu(U^*)$, and let $u = \sigma^{-1}(m)$, and $u' = \sigma^{-1}(m')$. Since $u' \in U^*$, as shown in the proof of Lemma 1, any equilibrium price vector p in state u' coincides with u' up to agent-dependent constants. By Proposition 1, price vector u' must also support allocation $h(m') = \hat{k}$ in any state $u'' \in \mu^{-1}(m')$. Eq. (1) then implies that for all $i \in N, k \in K$,

$$u''_{ik} - u''_{i\hat{k}} \leq u'_{ik} - u'_{i\hat{k}} = u'_{-i,\hat{k}} - u'_{-i,k} \leq u''_{-i,\hat{k}} - u''_{-i,k}$$

(where $u_{-i,k} = \sum_{j \neq i} u_{jk}$). Subtracting $u_{ik} - u_{i\hat{k}} = u_{-i,\hat{k}} - u_{-i,k}$, and taking the absolute value, we see that

$$\begin{aligned} \left| (u'_{ik} - u_{ik}) - (u'_{i\hat{k}} - u_{i\hat{k}}) \right| &\leq \max \left\{ \left| (u''_{ik} - u_{ik}) - (u''_{i\hat{k}} - u_{i\hat{k}}) \right|, \right. \\ &\quad \left. \left| (u''_{-i,k} - u_{-i,k}) - (u''_{-i,\hat{k}} - u_{-i,\hat{k}}) \right| \right\} \\ &\leq 2(N - 1) \|u'' - u\|. \end{aligned}$$

Now, using normalization and the triangle inequality,

$$\begin{aligned} \|u' - u\| &\leq C \sup_{i \in N, j, k \in K} \left| (u'_{ij} - u_{ij}) - (u'_{ik} - u_{ik}) \right| \\ &\leq 2C \sup_{i \in N, k \in K} \left| (u'_{ik} - u_{ik}) - (u'_{i\hat{k}} - u_{i\hat{k}}) \right| \\ &\leq 4C(N - 1) \|u'' - u\|. \end{aligned}$$

By the construction of the metric on messages (see footnote 8), this implies

$$\|u' - u\| \leq 4C(N - 1) \inf_{u'' \in \mu^{-1}(m')} \|u'' - u\| \leq 4C(N - 1) \rho(m', m).$$

Thus, σ^{-1} is Lipschitz continuous, which implies that $\dim M \geq \dim \sigma(U^*) \geq \dim U^*$ [11, Exercise 6.1.9(1)]. \square

The simplest application of Proposition 2 is to the case $N = 2$:

Corollary 1. *Suppose that in a normalized allocation problem with $N = 2$, for each $u_1 \in U_1$ there exists a “dual utility” $u_2 \in U_2$ such that $u_{1k} + u_{2k}$ does not depend on k . Then, if U_1 is a finite set, any efficient protocol transmits at least $\log |U_1|$ bits, and otherwise the dimension of its message space is at least $\dim U_1$.*

5. Approximation and discretized problems

One may hope that approximate efficiency could be achieved with less communication than exact efficiency. In this section we discuss how to analyze the communication burden of approximate efficiency. To be consistent with the computer science literature on approximation (e.g. [46]), we use an approximation measure that is invariant to the units of measurement. Namely, defining the maximum surplus available in state u by

$$S(u) = \max_{k \in K} \sum_i u_{ik},$$

we define the allocation rule F_r^* realizing approximation ratio $r \in [0, 1]$ as follows:¹¹

$$F_r^*(u) = \left\{ k \in K : \sum_i u_{ik} \geq rS(u) \right\}.$$

By construction, $F_1^* = F^*$ (the efficient allocation rule), and the set $F_r^*(u)$ is nonincreasing in r for all u .

For the study of approximation, we need to assume¹²

$$\bar{S} \equiv \sup_{u \in U} S(u) < \infty, \quad \underline{S} \equiv \inf_{u \in U: S(u) \neq 0} S(u) > 0. \tag{3}$$

The communication burden of F_r^* can be evaluated by examining the communication burden of the exactly efficient allocation rule F^* in a discretized problem. In computer science, discrete problems are usually considered in their own right (the inputs are restricted to be integers), but we will use them as a stepping stone for the analysis of approximation in the continuous problem.

For each state $u \in U$, let u^δ denote the state in which all utilities are rounded off to multiples of $\delta > 0$. Define the *upper δ -discretized problem* as the problem with the discrete state space $U^\delta = \{u^\delta : u \in U\}$. Suppose that we have a protocol Γ realizing exact efficiency for problem U^δ . We can then ask the agents to round off their utilities to multiples of δ and follow protocol Γ . Since the sum of rounded-off utilities for every allocation is within $N\delta/2$ from the true surplus at this allocation, the maximization of this sum results in a surplus loss of at most $N\delta$. Since the maximum available surplus is bounded below by \underline{S} , we realize approximation ratio $1 - N\delta/\underline{S}$, using as much communication as in Γ . In particular, full revelation of valuations rounded off with a sufficiently fine precision achieves an approximation ratio arbitrarily close to 1.

Examination of a discretized problem also allows to bound the communication burden of approximation *from below*. For this purpose, define the *lower δ -discretized problem* as the problem with the discrete state space $U^\delta \cap U$. This state space consists of those valuations from U that are multiples of δ . (In most applications considered in this paper, $U^\delta \subset U$, hence the lower and upper discretized problems coincide, in which case we will simply call them the discretized problem.) Since any misallocation in problem $U^\delta \cap U$ loses at least surplus δ , and the maximum available surplus is bounded above by \bar{S} , realizing an approximation ratio higher than $1 - \delta/\bar{S}$ in problem $U^\delta \cap U$ is equivalent to realizing exact

¹¹ This is a “worst-case” definition, since it requires uniform approximation across all states. The weaker requirement of “average-case” approximation given some probability distribution over states is considered in Section 8 below.

¹² Both inequalities are needed to ensure that approximation can be achieved with finite communication. For example, consider the problem of allocating one object between two agents whose valuations lie in $[0, 1]$, and so $\underline{S} = 0$. Pick $r \in (0, 1)$, and consider the restricted problem in which both agents’ valuations lie in the set $\{r^n\}_{n=0}^\infty \subset [0, 1]$. In this restricted problem, realizing an approximation ratio higher than r is equivalent to exact efficiency, and Corollary 1 implies that this requires a countable message space. Since arbitrary $r \in (0, 1)$ can be chosen, this implies that no positive approximation can be achieved with finite communication. The same conclusion is reached when the agents’ valuations lie in $[1, \infty]$ (and so $\bar{S} = \infty$), by considering the restricted problem $\{r^{-n}\}_{n=0}^\infty$.

efficiency. Realizing the same approximation ratio in problem U requires at least as much communication. This observation will allow us to bound below the communication burden of approximating efficiency in problem U by applying Corollary 1 to the upper discretized problem. To summarize:

Proposition 3. *In an allocation problem in which (3) holds, (i) realizing approximation ratio $1 - \delta N/\underline{S}$ does not require more communication than realizing exact efficiency in the upper discretized problem U^δ , and (ii) realizing an approximation ratio higher than $1 - \delta/\bar{S}$ requires at least as much communication as realizing exact efficiency in the lower discretized problem $U^\delta \cap U$.*

We want to characterize the dependence of the communication burden on the desired approximation, as well as on the parameters of the problem. For this purpose, we use the following three concepts, listed from weaker to stronger, which are standard in the computer science literature on approximation algorithms [46]:

- A *polynomial approximation scheme* (PAS) in some parameters is a protocol that for any given $\varepsilon > 0$ realizes approximation ratio $1 - \varepsilon$ using a number of bits that is polynomial in the parameters.
- A *fully polynomial approximation scheme* (FPAS) in some parameters is a protocol that for any $\varepsilon > 0$ realizes approximation ratio $1 - \varepsilon$ using a number of bits that is polynomial in ε^{-1} and the parameters.
- A *truly polynomial approximation scheme* (TPAS) in some parameters is a protocol that for any $\varepsilon > 0$ realizes approximation ratio $1 - \varepsilon$ using a number of bits that is polynomial in $\log \varepsilon^{-1}$ and the parameters.

PAS achieves an arbitrarily close approximation with polynomial communication, but does not stipulate how the communication burden depends on the approximation error. In contrast, in FPAS and TPAS, the error must shrink sufficiently fast with the number of bits transmitted. An economic example of FPAS is an ascending-bid auction for one unit with a minimum bid increment ε . Suppose that the agents' valuations for the unit lie in $[1, 2]$. At each price level starting from $p = 1$, the auction asks each agent to send one bit—"in" or "out." If at least one agent sends "in," the price is incremented by ε . The auction stops when all agents send "out," assigning the item to (one of) the agent(s) who sent "in" in the previous stage. Suppose that the agents behave "truthfully"—send "in" if and only if their valuations exceed the current price. Then the auction is exactly efficient for the discretized problem U^ε , and therefore, by Proposition 3(i), it realizes approximation ratio $1 - \varepsilon$. Since the maximum number of price increments is ε^{-1} , the auction's worst-case complexity is $N\varepsilon^{-1}$.

TPAS requires a much faster approximation than FPAS—the error must now shrink exponentially with the number of bits transmitted. An economic example of TPAS is a sealed-bid auction of a single unit, in which the agents submit their valuations rounded off to a multiple of ε . Suppose the agents' valuations for the unit lie in $[1, 2]$. Since the auction is exactly efficient for the upper discretized problem U^ε , by Proposition 3(i), it realizes approximation ratio $1 - \varepsilon$. Since it takes $\log \varepsilon^{-1}$ bits to transmit a valuation rounded off to a multiple of ε , the total number of bits transmitted is $N \log \varepsilon^{-1}$.

Note that our TPAS example was obtained by taking a fully efficient continuous protocol and asking the agents to round off their messages. This technique can be generalized: given a d -dimensional continuous protocol realizing approximation ratio r , rounding off the messages yields a TPAS to approximation ratio r that is linear in d . Intuitively, a message round-off error that is small in the Hausdorff metric defined in footnote 8 yields an allocation that is desirable for some state that is not too far from the true state, and therefore a small efficiency loss. Formally, we have

Proposition 4. *In an allocation problem satisfying (3), if there exists a protocol realizing approximation ratio r with a message space of upper box dimension $d > 0$,¹³ then for any $\varepsilon > 0$ there exists a protocol realizing approximation ratio $r - \varepsilon$ using $C(\varepsilon)$ bits, with $C(\varepsilon) \sim d \log \varepsilon^{-1}$ as $\varepsilon \rightarrow 0$.*

Proof. Suppose that $\Gamma = \langle M, \mu, h \rangle$ is a protocol realizing approximation ratio r . By the definition of upper box dimension, for each $\delta > 0$ there exists a covering of M by balls of radius δ centered at points $M_\delta \subset M$, with $\log |M_\delta| \sim d \log \delta^{-1}$. For each $m \in M$, we can choose a point $r_\delta(m) \in M_\delta$ such that $\rho(m, r_\delta(m)) \leq \delta$. Let $\mu_\delta(u) = r_\delta(\mu(u))$.

Consider the protocol $\Gamma_\delta = \langle M_\delta, \mu_\delta, h \rangle$. Note that Γ_δ inherits Existence and Privacy Preservation from Γ . Take any state $u \in U$ and any $m_\delta \in \mu_\delta(u)$. By construction, $m_\delta = r_\delta(m)$ for some $m \in \mu(u)$. For any $u' \in \mu^{-1}(m_\delta)$ we can write

$$\frac{\sum_i u_i(h(m_\delta))}{S(u)} \geq \frac{\sum_i u'_i(h(m_\delta))}{S(u')} - \frac{|\sum_i u_i(h(m_\delta)) - \sum_i u'_i(h(m_\delta))| + |S(u) - S(u')|}{S(u')}.$$

Since Γ realizes approximation ratio r , and so $h(m_\delta) \in F_r^*(u')$, and since arbitrary $u' \in \mu^{-1}(m_\delta)$ can be chosen, this implies

$$\frac{\sum_i u_i(h(m_\delta))}{S(u)} \geq r - \frac{2N}{S} \inf_{u' \in \mu^{-1}(m_\delta)} \|u' - u\|.$$

By the construction of the Hausdorff metric ρ on messages (see footnote 8), $\inf_{u' \in \mu^{-1}(m_\delta)} \|u' - u\| \leq \rho(m, m_\delta) \leq \delta$. Therefore, taking $\delta(\varepsilon) = \varepsilon S / (2N)$ ensures that $\Gamma_{\delta(\varepsilon)}$ realizes approximation ratio $r - \varepsilon$. The number of bits communicated by $\Gamma_{\delta(\varepsilon)}$ is

$$\log |M_{\delta(\varepsilon)}| \sim d \log \delta(\varepsilon)^{-1} = d \log \varepsilon^{-1} + d \log (2N/S) \sim d \log \varepsilon^{-1} \quad \text{as } \varepsilon \rightarrow 0. \quad \square$$

In particular, the proposition implies that if we discretize two efficient continuous protocols with message spaces of dimension d_1 and d_2 , respectively, to guarantee approximation error ε , then the worst-case complexities of the two discretized protocols are related as d_1/d_2

¹³ The upper box dimension coincides with the Hausdorff dimension for most “well-behaved” sets, but could exceed it for some sets [11].

asymptotically as $\varepsilon \rightarrow 0$.¹⁴ Thus, the dimension of the continuous message space based on the Hausdorff metric over messages is indeed a relevant measure of communication even if real-life communication is discrete. The result also implies that if we have continuous communication that is polynomial in some parameters and realizes approximation ratio r , then the discretized protocol is a TPAS for r in the same parameters. In such a case we will simply say that the discretized protocol is a “polynomial protocol realizing approximation ratio r .”

6. The Combinatorial Allocation problem

We now specialize to the *Combinatorial Allocation problem*, in which the allocation is that of L items among the N agents. Formally, the allocation space is $K = N^L$, and $k(l)$ denotes the agent holding item $l \in L$ in allocation $k \in K$.

We impose several standard restrictions on valuations (which can only reduce the communication burden):

- *No externalities* (NE): For each i and each $u_i \in U_i$, there exists $v_i : 2^L \rightarrow \mathbb{R}$ such that $u_{ik} = v_i(k^{-1}(i))$ for each $k \in K$.

In words, each agent i 's utility is a function v_i of the bundle $k^{-1}(i)$ allocated to him. We will call v_i the agent's *valuation*, and let $V_i \subset \mathbb{R}^{2^L}$ denote the class of his possible valuations. The state space can then be represented by $V = V_1 \times \cdots \times V_N$.

Each $v_i \in V_i$ is also assumed to satisfy the following restrictions:

- *Normalization* (N): $v_i(\emptyset) = 0$.
- *Monotonicity* (M): $v_i(S)$ is nondecreasing in $S \subset L$.
- *Boundedness* (B): Either $v_i(L) = 0$ or $v_i(L) \in [\gamma, 1]$ (where $\gamma > 0$ is viewed as a parameter of the problem).

Restriction (N), which is without loss of generality, rules out distinct valuations that differ only by a constant, and ensures the normalization property in Definition 3. Restrictions (M) and (B) are not needed for the analysis of exact efficiency, but will be used in the analysis of approximation (in particular, (B) ensures condition (3) above). Given that the analysis is invariant to scale, the choice of the upper bound as 1 is arbitrary, and the value of the lower bound γ will be irrelevant for most results. Let $V_{\text{gen}} \subset \mathbb{R}^{2^L}$ denote the class of all valuations satisfying (NE), (N), (M), and (B). (We will consider additional restrictions on valuations in Section 7.)

We obtain lower bounds for combinatorial allocation problems using Corollary 1 as a workhorse. While formulated for $N = 2$ agents, we can also apply it with $N > 2$ by letting

¹⁴ Conversely, if a given large number of bits C is to be transmitted in both cases, then the ratio of the approximation errors of the two protocols will grow exponentially with C when $d_2 < d_1$. This is consistent with the findings of Hurwicz and Marschak [19,20].

agents $i > 2$ have valuations that are identically zero.¹⁵ To apply Corollary 1, for each valuation $v \in V$ define the *dual* valuation by

$$\tilde{v}(S) = v(L) - v(L \setminus S) \quad \text{for all } S \subset L. \tag{4}$$

When the two agents' valuations are v and \tilde{v} , all allocations have the same surplus $v(L)$. Since $v \in V_{\text{gen}}$ implies that $\tilde{v} \in V_{\text{gen}}$, Corollary 1 yields

Proposition 5. *In the combinatorial allocation problem with general valuations, the dimension of the message space in any efficient protocol is at least $\dim V_{\text{gen}} = 2^L - 1$.*

Therefore, the communication burden of efficiency is at least as large as a full description of one agent's valuations. Recall from the discussion in Section 3 that this lower bound is essentially tight for $N = 2$. A more general upper bound is given by $(N - 1)(2^L - 1)$ numbers (all agents but one fully reveal their valuations), but the exact communication burden for $N > 2$ remains an open problem.¹⁶

To bound below the communication burden of approximation, we apply the same logic to the discretized problem $V_{\text{gen}}^\delta \subset V_{\text{gen}}$ with $\delta = 1$ (i.e., the problem in which the agents' valuations for all bundles are either 0 or 1). Corollary 1 implies that any efficient protocol for the discretized problem transmits at least $\log |V_{\text{gen}}^1|$ bits. Only counting those valuations in V_{gen}^1 that have $v(S) = 0$ for $|S| < L/2$ and $v(S) = 1$ for $|S| > L/2$, we see that $\log |V_{\text{gen}}^1| \geq \binom{L}{L/2}$.^{17,18} Thus, already with $N = 2$, by Proposition 3(ii) we obtain a lower bound on the communication burden of realizing a higher approximation ratio than $1 - \delta/N = 1/2$:

Proposition 6. *In the combinatorial allocation problem with general valuations, any protocol realizing an approximation ratio higher than $1/2$ communicates at least $\log |V_{\text{gen}}^1| \geq \binom{L}{L/2}$ bits.*

Observe that approximation ratio $1/N$ can be realized by auctioning off all items as a bundle to the highest bidder. (Indeed, the bundled auction realizes surplus $\max_i v_i(L)$,

¹⁵ We do not apply Proposition 2 directly with $N > 2$ agents, since then restriction (NE) dramatically reduces the fooling set U^* .

¹⁶ To see that the upper bound is not tight either, consider the case where $N > L$. In this case, in any allocation at most L agents receive nonempty bundles, therefore it cannot be efficient to allocate a bundle to an agent who does not hold one of the top L valuations for it. Thus, efficiency can be verified by announcing only the L highest valuations for each bundle (and the agents holding them), which uses only $L(2^L - 1)$ real numbers. We are obliged to Moshe Babaioff for bringing this to our attention.

¹⁷ For ease of exposition let L be even whenever it is divided by 2. If L is odd, one item can be assigned zero marginal benefit to both players and ignored, and so $L/2$ can be replaced with its integer part.

¹⁸ Note that $|V_{\text{gen}}^1|$ is the number of monotone boolean functions of L boolean variables. There is no closed-form expression for this number (this is known as "Dedekind's problem"), but its asymptotic behavior is described in [24].

while no individual agent can have a higher utility than that at any allocation.) Thus, Proposition 6 means that for $N = 2$, any improvement upon the bundled auction still requires very extensive communication, which still grows exponentially with L .¹⁹ In fact, we can prove a similar statement for $N > 2$, though using a different proving technique:

Proposition 7. *In the combinatorial allocation problem with general valuations, realizing an approximation ratio higher than $1/N$ requires communicating at least $\ln 2 \cdot \exp\{L/(2N^2) - 2 \ln N\}$ bits.*

Proof. In the appendix we consider the following *set packing problem*: Each of the N agents holds a collection of subsets of L , and the goal is to approximate the maximum *packing number*—the number of subsets in the union of their collections that are packed together, i.e., are pairwise disjoint. The set packing problem is reduced to the combinatorial allocation problem by letting, for each agent i , $v_i(S)$ be the maximum number of subsets in his collection that can be packed into S . We prove a lower bound on the communication complexity of any protocol realizing an approximation ratio higher than $1/N$ for the maximum packing number, which applies even when each agent i 's individual packing number (corresponding to $v_i(L)$) is restricted to be at most 1. This lower bound is given in the proposition. \square

This result should be contrasted with the findings of Lehmann et al. [29] and Holzman et al. [16], who suggest “simple” protocols improving upon the bundled auction. For example, Holzman et al. [16] notes that auctioning off the items in two equal-sized bundles achieves approximation ratio $r(L) = 2/L$ for any N , thus improving upon the bundled auction when $N > L/2$ (splitting L into more bundles allows further improvement). Lehmann et al. [29] proposes a polynomial approximation algorithm that can be adapted to the following protocol: At each stage of the protocol, each agent i who is not yet allocated any items announces a subset S_i of yet unallocated items that maximizes the ratio $v_i(S_i)/\sqrt{|S_i|}$, along with the maximum ratio itself. The agent who announces the highest ratio receives the requested bundle and quits. In the course of this protocol, agents announce $N(N+1)/2$ valuations and bundles. Lehmann et al. [29] shows that this algorithm realizes approximation ratio $r(L) = 1/\sqrt{L}$, which is higher than that realized by the bundled auction when $N > \sqrt{L}$.

Observe that these improvements over the bundled auction do not contradict Proposition 7. Intuitively, the proposition implies that in large problems in which the number N of agents is “substantially smaller” than the number L of items (e.g., smaller than $L^{1/2-\epsilon}$), “simple” protocols (e.g., polynomial in L) cannot improve over the bundled auction. When N is either comparable with or larger than L , simple protocols *can* improve over bundled auctions, though both bundled auctions and all other simple protocols realize a vanishing share of the available surplus as $N, L \rightarrow \infty$.

¹⁹ Indeed, by Stirling's formula, $\binom{L}{L/2} \sim \sqrt{2/(\pi L)} \cdot 2^L$ as $L \rightarrow \infty$.

7. Restricted valuations

7.1. Submodular valuations

Here each agent’s valuation space $V_i = V_{sm}$ is the set of all valuations $v \in V_{gen}$ for which the marginal benefit of each item $l \in L$, $v(S \cup l) - v(S)$, is nonincreasing in $S \subset L$.²⁰ Corollary 1 cannot be applied to this case directly, since dual (4) of a submodular valuation is not submodular unless both are additive (as defined in Section 7.3). We get around this problem by considering a subproblem in which all efficient allocations involve even splits of items, i.e., lie in

$$\tilde{K} = \left\{ k \in K : \left| k^{-1}(1) \right| = \left| k^{-1}(2) \right| = L/2 \right\}.$$

Namely, consider the set \tilde{V} of valuations $v \in \mathbb{R}^{2^L}$ satisfying

$$v(S) = 2|S|/L \text{ for } |S| < L/2, \tag{5a}$$

$$v(S) = 1 \text{ for } |S| > L/2, \tag{5b}$$

$$v(S) \in [1 - 1/L, 1] \text{ for } |S| = L/2, \tag{5c}$$

$$v(\{1, \dots, L/2\}) + v(\{L/2 + 1, \dots, L\}) = 1 - 1/(2L). \tag{5d}$$

One can easily verify that $\tilde{V} \subset V_{sm}$.

Note that in any state $(v^1, v^2) \in \tilde{V} \times \tilde{V}$, all efficient allocations lie in \tilde{K} . (Indeed, any other allocation has at most surplus $2 - 2/L$, while (5d) ensures that a surplus of at least $2 - 1/L$ can be obtained by giving items 1 to $L/2$ to one agent and items $L/2 + 1$ to L to the other.) Eq. (5d) ensures that thus constructed valuation class with allocations restricted to \tilde{K} is normalized.

For each $v \in \tilde{V}$, define its “quasi-dual” $\hat{v} \in \tilde{V}$ as follows:

$$\hat{v}(S) = \begin{cases} 2 - 1/L - v(L \setminus S) & \text{for } |S| = L/2, \\ 2|S|/L & \text{for } |S| < L/2, \\ 1 & \text{for } |S| > L/2. \end{cases}$$

By construction, the set of efficient allocations in any state $(v, \hat{v}) \in \tilde{V} \times \tilde{V}$ is exactly \tilde{K} . Thus, we can apply Corollary 1 with the two agents’ valuations restricted to \tilde{V} and the allocations restricted to \tilde{K} , which yields

Proposition 8. *In the combinatorial allocation problem with submodular valuations, the dimension of the message space in any efficient protocol is at least $\dim \tilde{V} = \left| \tilde{K} \right| - 1 = \binom{L}{L/2} - 1$.*

Thus, the communication burden of exact efficiency is still exponential in L (recall footnote 19). To examine approximation, consider the discretized problem $\tilde{V}^\delta \subset \tilde{V}$ with

²⁰ An equivalent definition of submodularity is that $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$ for all $S, T \subset L$.

$\delta = 1/L$ and $N = 2$. Applying Corollary 1 with the two agents' valuations restricted to \tilde{V}^δ and the outcomes restricted to \tilde{K} implies that the number of bits communicated by an efficient protocol for the discretized problem is at least $\log |\tilde{V}^\delta| = |\tilde{K}| - 1$. Proposition 3(ii) then implies

Proposition 9. *With submodular valuations, realizing an approximation ratio higher than $1 - 1/(2L)$ requires communicating at least $\binom{L}{L/2} - 1$ bits.*

Note that if we had a FPAS in L , then it could be used to realize approximation ratio $1 - 1/(2L)$ using polynomial communication in L , contradicting Proposition 9. Therefore, we have

Corollary 2. *With submodular valuations, FPAS in L is impossible.*

This result implies, for example, that an ascending-bid auction with L per-item prices and bid increment ε cannot approximate efficiency within ε , because the auction's worst-case complexity would be $NL\varepsilon^{-1}$, and so it would be a FPAS. Yet, we have been unable to rule out PAS—i.e., achieving any approximation ratio with polynomial communication in N, L . We do know from [28] that approximation ratio $1/2$ is realized by the (deterministic) “greedy” protocol that allocates the items in a fixed order to the agents who announce the highest current marginal benefit for them, and so communicates only NL numbers.

7.2. Homogeneous valuations

Here each agent's valuation space $V_i = V_h$ is the set of valuations $v \in V_{\text{gen}}$ that satisfy $v(S) = \phi(|S|)$ for all $S \subset L$, where $\phi : \{0, \dots, L\} \rightarrow \mathbb{R}$. That is, agents care only about the number of items they receive.

Since dual (4) of a homogeneous valuation is homogeneous, Corollary 1 implies

Proposition 10. *In the combinatorial allocation problem with homogeneous valuations, the dimension of the message space in any efficient protocol is at least $\dim V_h = L$.*

Now consider the discretized problem $V_h^\delta \subset V_h$, with $\delta = 1/R$ for some R . By Corollary 1, exact efficiency in this problem requires communicating at least $\log |V_h^\delta|$ bits. Note that $|V_h^\delta|$ is the number of monotone functions $\phi : \{1, \dots, L\} \rightarrow \{0, \dots, R\}$ such that $\phi(L) \geq \gamma R$. If we fix $\phi(L) = R$ for simplicity, then the number of such functions is exactly the number of ways that R indistinguishable balls (corresponding to the function's unit jumps) can be partitioned into L urns (corresponding to the jump points). This number is $\binom{R+L-1}{L-1}$.²¹ Therefore, exact efficiency in problem V_h^δ requires communicating at least

²¹ This is proven by putting all the R balls in a row with $L - 1$ dividers, and letting urn $l = 1, \dots, L$ contain the balls lying between dividers $l - 1$ and l . Thus, the different allocations of balls into urns correspond to the different positions that the $L - 1$ dividers can occupy in a row of $R + L - 1$ objects.

$\log \binom{\delta^{-1}+L-1}{L-1}$ bits. By Proposition 3(ii), at least as much communication is needed to realize an approximation ratio higher than $1 - \delta/2$ when $N = 2$.

On the other hand, the full revelation protocol in the discretized problem V_h^δ uses at most $N\delta^{-1} \log(L + 1)$ bits. Indeed, each agent needs only to communicate δ^{-1} jump points in $\{0, \dots, L\}$, each of which is communicated with $\log(L + 1)$ bits. By Proposition 3(i), this communication realizes approximation ratio $1 - \delta/\gamma$ in the continuous problem (note that the round-off error arises only on agents with non-null valuations). Summarizing the results, we have

Proposition 11. *In the combinatorial allocation problem with homogeneous valuations, for any $\delta > 0$, (i) realizing an approximation ratio higher than $1 - \delta/2$ requires communicating at least $\log \binom{\delta^{-1}+L-1}{L-1}$ bits, and (ii) approximation ratio $1 - \delta/\gamma$ is realized by full revelation of valuations rounded off to multiples of δ , which takes at most $N\delta^{-1} \log(L + 1)$ bits.*

Corollary 3. *In the combinatorial allocation problem with homogeneous valuations, (i) TPAS in $\log L$ is impossible even for $N = 2$ and $\gamma = 1$, but (ii) full revelation of rounded off valuations is a FPAS in parameters $\log L, N$, and γ^{-1} .*

Proof. (i) If there existed a TPAS in $\log L$, then realizing approximation ratio $1 - 1/(4L)$ would take only polynomial communication in $\log L$. By Proposition 11(i), however, it requires communicating at least $\log \binom{2L+L-1}{L-1} = \log \frac{(3L-1)!}{(2L)!(L-1)!} \geq L - 1$ bits.

(ii) By Proposition 11(ii), full revelation of valuations rounded off to multiples of $\varepsilon\gamma$ realizes approximation ratio $1 - \varepsilon$ using at most $N\varepsilon^{-1}\gamma^{-1} \log(L + 1)$ bits. \square

Proposition 11(ii) means, in particular, that we can achieve any given positive approximation error using $O(\log L)$ bits. As L grows, this amount of communication is proportional to that of simply announcing an allocation (which takes roughly $N \log(L + 1)$ bits). On the other hand, Corollary 3(i) means that the extra communication burden needed to guarantee halving the approximation error (regardless of the starting error) is exponential in $\log L$. Intuitively, this means that when the number L of items is large, and we have a protocol that achieves a close approximation of efficiency, a small reduction in inefficiency requires an enormous increase in communication.

Our analysis can also be related to the model of Calsamiglia [10], in which instead of L indivisible goods there is one unit of an infinitely divisible good. In this case, V is the space of nondecreasing functions $v : [0, 1] \rightarrow [0, 1]$, and so $\dim V = \infty$. Corollary 1 then implies that efficiency requires an infinite-dimensional message space, re-deriving Calsamiglia’s result.^{22,23} However, Calsamiglia’s model allows a FPAS, provided that

²² Calsamiglia [10] restricts the valuation of agent 1 to be concave and that of agent 2 to be convex. Since the dual of a concave valuation is convex, the analysis goes through without modification. Similarly, the agents’ valuations can be restricted to be arbitrarily smooth, since smoothness is preserved under duality.

²³ In contrast, when both agents’ valuations are known to be concave, a Walrasian equilibrium with a single real-valued price exists and realizes efficiency (regardless of whether the good is divisible or not).

the agents' valuation functions satisfy

$$|v(x') - v(x)| \leq (-\log |x' - x|)^{-A} \quad \text{for all } x, x' \in [0, 1]$$

for some $A > 0$. Under this mild strengthening of continuity (for example, implied by Hölder continuity of any degree), restricting agents to consume the good in $L = 2^{(\varepsilon\gamma)^{-1/A}}$ identical discrete units reduces the surplus by at most $N (\log L)^{-A} = N\varepsilon\gamma$. Running the protocol described in Corollary 3(ii) on this discretized allocation space will approximate the maximum surplus within $2N\varepsilon\gamma$, and so realize approximation ratio $1 - 2\varepsilon$, while communicating only $N\varepsilon^{-1} \log(L + 1) \approx N\varepsilon^{-1} (\varepsilon\gamma)^{-1/A}$ bits. Thus, we have a FPAS, even though exact efficiency in this case requires infinite-dimensional communication.

7.3. Substitute valuations

Here each agent's valuation space $V_i = V_{\text{sub}}$ is the set of valuations $v \in V_{\text{gen}}$ whose indirect utility function $w(p) = \max_{S \subseteq L} (v(S) - \sum_{l \in S} p_l)$ is submodular in $p \in \mathbb{R}_+^L$. This is one of the many equivalent definitions of the substitute property [13,34].²⁴

Since $V_{\text{sub}} \subset V_{\text{sm}}$ (see [13]), dual (4) of a substitute valuation is not one, except when both are *additive*, i.e., take the form $v(S) = \sum_{l \in S} \phi_l$ for some $\phi \in \mathbb{R}^L$. Let V_{add} denote the class of additive valuations. Since $V_{\text{add}} \subset V_{\text{sub}}$, and the dual of an additive valuation is itself, Corollary 1 yields

Proposition 12. *In the combinatorial allocation problem with additive or substitute valuations, the dimension of the message space in any efficient protocol is at least $\dim V_{\text{add}} = L$.*

This lower bound is attained by the Walrasian equilibrium with per-item prices, which always exists with substitute valuations [13,23].

A major disadvantage of the Walrasian protocol is that it is nondeterministic, leaving open the question of how to *find* an equilibrium. Deterministic protocols achieving this were proposed in [3,9,14,34]. These protocols are variations on the ascending-bid auction with prices quoted for individual items, and so they are only FPAS.

We improve upon the proposed auction designs by describing a TPAS for this setting. For this purpose, we write the efficient allocation problem as an integer programming problem, letting $x_{iS} = 1$ if agent i 's allocation $k^{-1}(i) = S$ and $x_{iS} = 0$ otherwise. As shown in [7], if a Walrasian equilibrium exists then any efficient allocation must also solve the relaxed surplus-maximization program in which fractional allocations x_{iS} are allowed:

$$\begin{aligned} & \max_{x \in \mathbb{R}_+^{N \cdot 2L}} \sum_{i \in N, S \subseteq L} x_{iS}^i v_i(S) & (P) \\ \text{s.t.} & \sum_{i \in N, S \subseteq N: l \in S} x_{iS} \leq 1 \quad \text{for all } l \in L, & \sum_{S \subseteq L} x_{iS} \leq 1 \quad \text{for all agents } i \in N. \end{aligned}$$

²⁴ Kelso and Crawford [23] introduced the property under the name of "gross substitutes," but in the absence of wealth effects the concepts of gross and net substitutability coincide.

Linear program (P) has only $N + L$ constraints but an exponential number of variables, so it would be hard to solve it directly. It is easier to solve the dual program

$$\begin{aligned} \min_{p \in \mathbb{R}_+^L, w \in \mathbb{R}_+^N} \quad & \sum_{l \in L} p_l + \sum_{i \in N} w_i & (D) \\ \text{s.t.} \quad & w_i - \left[v_i(S) - \sum_{S \ni l} p_l \right] \geq 0 \quad \text{for all } i \in N, S \subset L, \end{aligned}$$

where p_l and w_i denote the Lagrange multipliers with the constraints in (P) associated with item $l \in L$ and agent $i \in N$, respectively. Examination of the complementary slackness condition shows that a vector $(x, p, w) \in \mathbb{R}^{N \cdot 2^L} \times \mathbb{R}^L \times \mathbb{R}^N$ is comprised of solutions to (P) and (D) if and only if it describes a Walrasian equilibrium, with x being the (possibly fractional) equilibrium allocation, p the price vector, and w the vector of agents' utilities. We proceed under the assumption that an integral Walrasian equilibrium allocation exists.

While (D) has an exponential number of constraints, each of the constraints depends on the valuation of a single agent. This allows to solve (D) with a separation-based linear programming algorithm, such as the ellipsoid method [22]. The method uses an oracle, who, presented with a candidate solution, produces a violated inequality whenever one exists. Consider a protocol running a separation-based algorithm, but instead of each oracle query, asking each agent i to report a bundle S that gives him a higher net utility than his "utility target" w_i at the current price vector p . If such a report is made by one of the agents, the protocol continues. It is known that when the inputs (valuations) are discrete multiples of δ , the separation-based algorithm produces a solution within a number of steps that is polynomial in the number of variables (in our case $N + L$) and $\log \delta^{-1}$ (the size of each "input" number). Since at each step there are at most N numbers and bundles announced by the agents, the whole protocol uses polynomial communication in N, L , and $\log \delta^{-1}$. Thus, we have a polynomial procedure to calculate the value of δ -discretized program (P), which approximates the true value of (P) to within $N\delta/2$. An approximate integer solution to (P) can then be deduced using standard computational techniques of self-reduction, yielding a TPAS.²⁵

8. Average-case analysis

Given a probability distribution over the states of the world, we can relax the notion of approximation to the requirement that only the expected surplus be close to optimal. At the same time, we can count the expected rather than worst-case number of bits transmitted,

²⁵ Indeed, since (P) has an integer solution, there exists an allocation of item 1 that does not reduce its value, and so does not reduce the value of the δ -discretized (P) by more than $N\delta$. Thus, let us find an agent such that upon allocating item 1 to him, the value of the δ -discretized (P) does not fall by more than $N\delta$. Such allocation of item 1 may not be exactly optimal, but it will not reduce the value of (P) by more than $2N\delta$. Then allocate item 2 in the same fashion, then item 3, etc. Since we try allocating each item to each agent, we will use NL calls to the polynomial procedure solving the δ -discretized (P). Since the accumulated loss of surplus is at most $L \cdot 2N\delta$, we have a TPAS.

which allows a savings from coding more frequent messages with fewer bits (as in Shannon's information theory [43]). The communication measure that uses such average-case counting is called "distributional communication complexity," since the results clearly depend on the assumed joint distribution of the states of the world. For example, if the distribution puts all weight on a single state, then an efficient allocation is known and can be implemented with no communication. Thus, it is only interesting to consider distributions that are sufficiently diffuse so that no allocation has a high a priori probability of being efficient. It turns out that for *some* such distributions (which we do not claim to be realistic), the communication complexity of approximating efficient combinatorial allocations still grows exponentially with the number of items:

Proposition 13. *In the combinatorial allocation problem, there exists a sequence of joint probability distributions over valuation profiles (v^1, \dots, v^N) for each N and L such that for any $\varepsilon > 0$, realizing fraction $1/N + \varepsilon$ of the maximum expected surplus requires transmitting an expected number of bits that is at least $c \exp \{L / (2N^2) - 5 \ln N\}$, for some fixed $c > 0$.*

Proof. Consider the set packing problem described in the proof of Proposition 7. In the appendix we prove a lower bound on the communication complexity of distinguishing between the states in which N subsets can be packed from those in which only one subset can be packed (and each agent's individual packing number is at most 1), which applies to randomized protocols with any bounded error. Using the equivalence of randomized complexity and distributional complexity (which follows from the Minimax Theorem—see [25, Section 3.4]), it follows that for some sequence of distributions over states, our lower bound applies to the "Discrimination Problem" of distinguishing the states with surplus $S(v) = N$ from those with $S(v) = 1$ correctly with probability at least $1/2 + \varepsilon/4$. In particular, we must have $\Pr \{S(v) = N\} \leq 1/2 + \varepsilon/4$, for otherwise declaring " N " would solve the Discrimination Problem.

Now consider the conditional distribution on states with $S(v) = N$ (assigning probability zero to all other states), so that the maximum expected surplus is N . Any protocol Γ that achieves fraction $1/N + \varepsilon$ of it on the conditional distribution must realize a surplus greater than 1 with probability at least ε . We can adapt Γ to solve the Discrimination Problem for the original distribution as follows: Run Γ and ask the agents to announce their utilities at the realized allocation. Declare " N " if the sum of the announcements exceeds 1, declare " 1 " otherwise. The probability of error is at most $\Pr \{S(v) = N\} (1 - \varepsilon) \leq (1/2 + \varepsilon/4) (1 - \varepsilon) < 1/2 - \varepsilon/4$, hence the protocol solves the Discrimination Problem. This implies that Γ must satisfy our lower bound, which is stated in the proposition. \square

Since the bundled auction guarantees share $1/N$ of the expected surplus, the proposition implies that for some joint distribution over the agents' valuations, achieving a higher expected surplus than the bundled auction still requires expected communication that is exponential in L .

In the distribution constructed in the above proposition, the valuations are not necessarily independently distributed. We can obtain a (weaker) lower bound on approximation for independently distributed valuations using the distributional lower bounds of Babai et al. [5]:

Proposition 14. *In the combinatorial allocation problem with $N = 2$ agents, there exists a sequence of probability distributions pairs D_1, D_2 over valuations for each L such that realizing fraction c of the maximum expected surplus (for some fixed $c < 1$) when the agents' valuations are distributed independently according to D_1, D_2 , respectively, requires communication of an expected number of bits that is exponential in L .*

Proof. We will use a reduction from the “disjointness problem” from communication complexity theory (see [25]). In this problem, each agent $i = 1, 2$ is given a subset X_i a set M , and the objective is to decide whether $X_1 \cap X_2 = \emptyset$. Babai et al. [5] proves a lower bound on the distributional complexity of disjointness for product distributions:

Theorem 1 (Babai et al. [5]). *There exists a distribution D on subsets of M with $|M| = m$, and a fixed $d > 0$, such if the two agents' sets are drawn according to D , then any protocol that communicates in expectation at most $d\sqrt{m}$ bits must err with at least 1% probability when attempting to solve the disjointness problem.*

We will now show that any protocol for combinatorial allocation that achieves 99.5% expected efficiency when the agents' valuations are drawn independently according to distributions D_1, D_2 (to be defined below) can be used to obtain a protocol for the disjointness problem for $m = \binom{L}{L/2}$ that errs on at most 1% of inputs drawn according to D . Thus the lower bound of $d\sqrt{m} = d\sqrt{\binom{L}{L/2}}$ (which is exponential in L) communication applies to the combinatorial allocation problem.

Distributions D_1, D_2 of the two agents' valuations are defined as follows: Let $M = \{S \subset L : |S| = L/2\}$, hence $|M| = m$. For each $i = 1, 2$, the valuation v_i of agent i is chosen by first choosing a random subset $X_i \subset M$ according to the distribution D in the above theorem. We define $v_i(S) = 0$ for $|S| < L/2$; $v_i(S) = 1$ for $|S| > L/2$. For agent 1, we define for $|S| = L/2$, $v_1(S) = 1$ if $S \in X_1$ and $v_1(S) = 0$ otherwise. For agent 2, we define for $|S| = L/2$, $v_2(S) = 1$ if $L \setminus S \in X_2$ and $v_2(S) = 0$ otherwise. Now in order to solve the disjointness problem on X_1 and X_2 , the two parties can each create a valuation according to the rule specified above, solve the combinatorial allocation problem, and declare that X_1 and X_2 are not disjoint when the obtained allocation has value 2, and disjoint otherwise. Observe that finding an allocation with surplus 2 means finding a partition of L into two sets $(S, L \setminus S)$ of size $L/2$ each such that $S \in X_1$ and $S \in X_2$, thus proving that X_1 and X_2 are not disjoint. Thus, the declaration of nondisjointness is always correct. On the other hand, since any inefficient allocation loses at least half the available surplus, an allocation protocol that loses at most 0.5% of expected surplus cannot produce an inefficient allocation with probability more than 1%. Therefore, the probability of falsely declaring disjointness is at most 1%. \square

The proof is done for $c = 99.5\%$, which is derived from a constant quoted in [5]. No optimization of the constant was attempted and it seems likely that a substantial strengthening is possible.

9. Comparison with computational complexity

The communication problem examined here is fundamentally different from the previously considered problem of *computing* an efficient or approximately efficient allocation when all the valuations are known. The computational complexity of a problem is defined relative to its input size, but in the combinatorial allocation problem the size of the input—a description of the valuations—is itself exponential in the number L of items. For this reason, the computational complexity literature has tended to focus on cases in which the input size is small, such as that of “single-minded preferences,” in which each agent values only a single bundle of items.²⁶ Even in such simple cases from the viewpoint of communication, the efficient combinatorial allocation problem has been shown to be NP-complete [47].

Nevertheless, we believe that the communication bottleneck is more severe in practice than the computational one. Recall that NP-completeness only indicates that the computational complexity is exponential asymptotically as the number L of items goes to infinity (and that only if $P \neq NP$, which is considered likely but not proven). In practice, computational complexity can be handled for up to hundreds of items (and thousands of bids) optimally [41,47] and thousands of items (with tens of thousand of bids) near-optimally [50]. In contrast, we derive *exact* lower bounds on communication complexity for any given L . For example, with general valuations, Proposition 5 establishes that exact efficiency requires communicating at least one price for each of the $2^L - 1$ possible bundles of items, and Corollary 6 establishes that with two bidders any improvement over the bundled auction still requires communicating at least $\binom{L}{L/2}$ bits. For example, with $L = 40$ items, any improvement would require the bidders to send at least $\binom{40}{20}$ bits $\simeq 16$ Gigabytes of data, which would take 164 years of continuous typing (the average typist enters 200 keystrokes per minute, and each keystroke transmits 8 bits). Considering the average case does not help matters much for *some* probability distributions over valuations. It is possible that the average-case communication complexity is lower for some realistic probability distributions, yet all simulations so far have not been able to handle the communication problem with more than a dozen items (see, e.g. [17]).²⁷

10. Incentives

So far we have ignored the agents’ incentives to follow the prescribed strategies. If the agents behave in their self-interest, the designer faces additional “incentive-compatibility” constraints requiring that the agents’ strategies constitute an equilibrium of

²⁶ Alternatively, one could handle large inputs by allowing the algorithm to ask a “black box” (oracle) for each agent queries about the agent’s valuation. (Such models typically restrict the queries that can be used—e.g., only to ask the agent to report his valuation for a particular allocation, or his demand at given prices.) The number of queries in such a model is bounded below by the deterministic communication complexity of the problem divided by the amount of information in a single query.

²⁷ Note also that while the computational burden may be distributed among the bidders, e.g., by asking them to suggest allocations or matches to their bundle [6], all the communication in a combinatorial auction passes through the auctioneer.

the communication game. In this section we show how, when the agents' utilities are quasi-linear in monetary transfers, such transfers may be used to satisfy their incentive constraints without substantially raising the communication burden.

Suppose that after running the protocol, we ask each agent to report his payoff $\pi_i = u_{ik}$ at the resulting allocation k , and pay each agent i a transfer $t_i = \sum_{j \neq i} \pi_j$. Under this transfer scheme (first proposed by Reichelstein [39, pp. 45–46]), 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 [32], the agents become a "team"). If the protocol is efficient, then obeying the prescribed strategies constitutes an ex post *equilibrium* under the described transfer scheme: no unilateral deviation by an agent can increase his payoff in any state.²⁸

But what if the protocol is not exactly but only approximately efficient? The behavior of rational agents in such a protocol will depend on their beliefs. Let us make the standard assumption in economics that all the agents have a correct prior over the states of the world. Since the proposed transfer scheme converts the game played by the agents into one of common interest, the strategy profile that maximizes the expected surplus will also constitute a Bayesian–Nash Equilibrium of the game.²⁹

A possible criticism of the proposed transfer scheme is that it is very costly—the designer must pay out a total subsidy equal to $n - 1$ times the total surplus. The cost can be covered without affecting incentives by charging each agent a participation fee that depends on other agents' messages. The largest fee f_i that ensures agent i 's participation is the surplus that could be achieved in his absence, and the resulting net transfer to the agent, called the *Vickrey–Groves–Clarke (VCG) transfer*, is $t_i - f_i$ (this equals to the externality imposed by agent i externality on the other agents).

Since f_i can be calculated by solving the maximization problem without agent i , VCG transfers can be revealed while multiplying the communication burden of efficiency by at most $N + 1$. However, it is often possible to calculate VCG transfers without solving the N additional subproblems. Lahaie and Parkes [27], building on Proposition 1, characterize efficient communication protocols that reveal VCG transfers as those that reveal *universal equilibrium prices*, which are defined as supporting prices that also remain equilibrium prices upon removing any single agent. Such revelation may require more communication than simply finding an efficient allocation (and *some* supporting prices), but less than the efficient communication burden multiplied by $N + 1$. Examples of protocols that yield a universal price equilibrium (and thus VCG payments) in specific allocation problems

²⁸ Furthermore, truthfulness becomes a dominant strategy if message sequences that are not consistent with any type are banned. (Then an agent knows that other agents' messages will be consistent with *some* state, and in any such state his truthfulness will maximize the total surplus, and hence the agent's payoff.)

²⁹ Taken to an extreme, this observation implies that agents need not be offered any protocol at all—it suffices to make them internalize the communication costs, along with others' utilities, using ex post payments. Then 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 (being able to coordinate on it). But if agents are not fully rational, it is not clear how to consider their incentives in the first place.

without solving the subproblems obtained by removing a single agent can be found in [9,12,38].

11. Conclusion

Price mechanisms are the most commonly observed and the best studied economic allocation mechanisms. However, until now there has not been a complete understanding of their role. To be sure, the Welfare Theorems show that the Walrasian price mechanism produces efficient allocations in convex economies. Still, the possibility remained that other mechanisms also produce efficient allocations in convex economies, or that nonprice mechanisms perform better than price mechanisms in nonconvex economies. The designers of combinatorial auctions have proposed many designs that purport to find efficient allocations without discovering the prices supporting them (e.g. [17]).

The present paper has shown that in fact, *any* efficient mechanism is “essentially” a price mechanism, in the sense that it must reveal supporting prices along with the efficient allocation itself. Thus, the indispensable role of prices for implementing efficient allocations is made clear. This result holds regardless of the agents’ incentives, even if the agents report truthfully.

In the combinatorial allocation problem, efficient communication must name (at least) one price for each of the $2^L - 1$ possible bundles, where L is the number of items. We demonstrate that all possible monotonic price vectors must be used, hence the required communication is at least as extensive as full revelation of one agent’s preferences. Even if we only require a better approximation of efficiency than that obtained by auctioning off all items as a bundle, bidders in a two-bidder auction would still have to communicate at least $\binom{L}{L/2}$ bits. For $L = 40$ (a realistic number), this would take more than a 100 years of continuous typing. Even if approximation is required only on expectation, it would still require exponential communication in L for *some* probability distribution over valuations. These results imply that for realistic values of L , any combinatorial auction design or “preference elicitation” scheme suggested in the literature would either take a prohibitively long time, or, if terminated after some reasonable time, would not be guaranteed to produce an efficient or approximately efficient allocation.

Our results should not be taken to imply that all real-life combinatorial auctions are useless, any more than Arrow’s impossibility theorem implies that all real-life institutions are useless. Rather, by showing that no institution is guaranteed to achieve good results on the universal preference domain, Arrow’s theorem has led researchers to examine the performance of specific institutions on restricted domains. Similarly, we hope to motivate auction designers to focus on specific classes of preferences or probability distributions over them. The burden of characterizing the environments on which a proposed auction design works well should lie on the proposer. We expect the tools developed in the present paper to be useful for such characterizations, as they proved to be for analyzing the cases of submodular, homogeneous, and substitute valuations.

Finally, we have clarified the validity of measuring the communication burden of efficiency with the dimension of the required message space, as is common in the economic literature. The key question is whether this measure accurately reflects the difficulty of

approximating efficiency with a discretized mechanism. We find that the dimensionality of message space is indicative of the complexity of achieving a “truly polynomial” approximation of efficiency. On the other hand, a somewhat slower but still practical “fully polynomial” approximation is sometimes achieved with much less communication. In such cases, the economic measure may seriously overstate the “hardness” of the communication problem. A dramatic example of this is offered by Calsamiglia’s [10] model of allocating a homogeneous divisible good, in which exact efficiency requires infinite-dimensional communication, but we demonstrate a fully polynomial approximation mechanism.

Acknowledgments

We are grateful to Paul Milgrom, Luis Rayo, Azeem Shaikh, Yoav Shoham, Steve Tadelis, Moshe Tennenholtz, and participants of seminars at Harvard-MIT, Princeton, Stanford, Toronto, and the 2002 DIMACS Workshop on Computational Issues in Game Theory and Mechanism Design for useful comments. N.N. was supported by a grant from the Israeli Academy of Sciences. I.S. was supported by the Guggenheim Foundation and the National Science Foundation (Grant SES-0214500), and was hosted by the Institute for Advanced Study in Princeton during his work on the project. Research assistance by Ronald Fadel and Hui Li is gratefully acknowledged.

Appendix. A set packing lower bound

Our lower bound for the set packing problem will use a reduction from the following “approximate-disjointness” problem which was studied in [2]:

The Approximate Disjointness Problem. Each player $i \in N$ holds a subset $B_i \subset T$ of a finite set T . We are required to distinguish between the following two extreme cases:

- Negative: $\bigcap_{i \in N} B_i \neq \emptyset$,
- Positive: for every $i \neq j$, $B_i \cap B_j = \emptyset$.

A lower bound on the required communication of cT/N^4 for some fixed constant $c > 0$ was given in [2] for randomized protocols (with two-sided error). The lower bound was improved by Jaikumar Radhakrishnan and Venkatesh Srinivasan to $\ln 2 \cdot T/N$ for deterministic and nondeterministic protocols.

Let us now define the approximate set packing problem. Each agent $i \in N$ holds a collection $A_i \subset 2^T$. The objective is to approximate the *packing number*—the maximum number of sets in $\bigcup_{i \in N} A_i$ that can be packed together, i.e., are pairwise disjoint. We restrict attention to cases in which each individual agent’s collection A_i contains no two pairwise disjoint sets. We will prove a lower bound on the communication complexity of distinguishing between the case where the packing number is 1 (i.e. any two sets $S_i \in A_i$ and $S_j \in A_j$ for $i, j \in N$ intersect) and the case where there exist N disjoint sets $S_i \in A_i$ for $i \in N$, and so the packing number is N .

We reduce this problem from the approximate-disjointness problem on a set T that consists of partitions of set L into N subsets. That is, each $t \in T$ is a partition (t_1, \dots, t_N) of L . We require set T to have the following property:

Definition A.1. A set T of partitions of set L among N agents has the pairwise-intersection property if for every $t', t'' \in T$ such that $t' \neq t''$ we have $t'_i \cap t''_j \neq \emptyset$ for every $i, j \in N$, i.e., any two elements of different partitions intersect.

Lemma A.1. *There exists a set T of partitions of set L among N agents with the pairwise-intersection property of size $|T| = \lfloor e^{L/(2N^2)}/N \rfloor$.*

Proof. We use the probabilistic method by constructing a random set T of partitions as follows: Choose each partition in T independently at random, by placing each element of L independently in one of the N parts of the partition with equal probabilities. Now fix $i, j \in N$ and two partitions $t', t'' \in T$ such that $t' \neq t''$. We can calculate

$$\Pr\{t'_i \cap t''_j = \emptyset\} = (1 - 1/N^2)^L < e^{-L/N^2}.$$

Since there are at most $N^2 T^2$ such choices of indices, we get that as long as $N^2 T^2 \leq e^{L/N^2}$, the probability that we obtain a set of partitions satisfying the pairwise-intersection property is positive. Hence, such a set of partitions exists. \square

We can now specify the reduction of approximate disjointness on set T to the approximate set packing problem. Player i receiving as input a set $B_i \subset T$ constructs the collection $A_i = \{t_i | t \in B_i\}$. Now, if there exists $t \in \cap_i B_i$, then (t_1, \dots, t_N) constitutes an N -packing. If, on the other hand, $B_i \cap B_j = \emptyset$ for all $i \neq j$, then for any two sets $t'_i \in A_i$ and $t''_j \in A_j$ we have $t' \neq t''$ and thus by the pairwise intersection property $t'_i \cap t''_j \neq \emptyset$ for every $i, j \in N$, hence the packing number is at most 1.

From the lower bounds described above for the approximate disjointness problem we obtain the following lower bounds for the set packing problem:

Theorem A.1. *Any N -player protocol (deterministic or nondeterministic) realizing an approximation ratio higher than $1/N$ for the set packing problem communicates at least $\ln 2 \cdot e^{L/(2N^2)} - 2 \ln N$ bits. A randomized protocol achieving this communicates, on expectation, at least $ce^{L/(2N^2)} - 5 \ln N$ bits, for some fixed $c > 0$.*

References

- [1] H. Abelson, Lower bounds on information transfer in distributed computations, *J. Assoc. Comput. Mach.* 27 (1980) 384–392.
- [2] N. Alon, Y. Matias, M. Szegedy, The space complexity of approximating the frequency moments, *J. Comput. System Sci.* 58 (1999) 137.
- [3] L. Ausubel, An efficient dynamic auction for heterogeneous commodities, Mimeo, 2002.
- [4] L. Ausubel, P. Milgrom, Ascending auctions with package bidding, *Frontiers Theoret. Econ.* 1 (2002).
- [5] L. Babai, P. Frankl, J. Simon, Complexity classes in communication complexity theory, *Proc. FOCS* (1986) 337–347.

- [6] J. Banks, J. Ledyard, D. Porter, Allocating uncertain and unresponsive resources: an experimental approach, *RAND J. Econ.* 20 (1989) 1–25.
- [7] S. Bikhchandani, J.W. Mamer, Competitive equilibrium in an exchange economy with indivisibilities, *J. Econ. Theory* 74 (1997) 385–413.
- [8] S. Bikhchandani, J. Ostroy, The package assignment model, *J. Econ. Theory* 107 (2002) 377–406.
- [9] S. Bikhchandani, R. Vohra, S. de Vries, J. Schummer, Linear programming and Vickrey auctions, in: B. Dietrich, R. Vohra (Eds.), *Mathematics of the Internet: E-Auction and Markets*, Springer, Berlin, New York, 2002.
- [10] X. Calsamiglia, Decentralized resource allocation and increasing returns, *J. Econ. Theory* 14 (1977) 262–283.
- [11] G.E. Edgar, *Measure, Topology, and Fractal Geometry*, Springer, Berlin, New York, 1990.
- [12] J. Feigenbaum, A. Krishnamurthy, R. Sami, S. Shenker, Hardness results for multicast cost sharing, *Theoret. Comput. Sci.* 304 (2003) 215–236.
- [13] F. Gul, E. Stacchetti, Walrasian equilibrium with gross substitutes, *J. Econ. Theory* 87 (1999) 95–124.
- [14] F. Gul, E. Stacchetti, The english auction with differentiated commodities, *J. Econ. Theory* 92 (2000) 66–95.
- [15] F. Hayek, The use of knowledge in society, *Amer. Econ. Rev.* 35 (1945) 519–530.
- [16] R. Holzman, N. Kir-Dahav, D. Monderer, M. Tennenholtz, Bundling equilibrium in combinatorial auctions, *Games Econ. Behav.* 47 (2004) 104–123.
- [17] B. Hudson, T. Sandholm, Effectiveness of query types and policies for preference elicitation in combinatorial auctions, in: *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, New York, 19–23 July, 2004, pp. 386–393.
- [18] L. Hurwicz, On the dimensional requirements of informationally decentralized Pareto-satisfactory processes, in: K.J. Arrow, L. Hurwicz (Eds.), *Studies in Resource Allocation Processes*, Cambridge University Press, New York, 1977, pp. 413–424.
- [19] L. Hurwicz, T. Marschak, Finite allocation mechanisms: approximate Walrasian versus approximate direct revelation, *Econ. Theory* 21 (2003) 545–572.
- [20] L. Hurwicz, T. Marschak, Comparing finite mechanisms, *Econ. Theory* 21 (2003) 783–841.
- [21] J.S. Jordan, The competitive allocation process is informationally efficient uniquely, *J. Econ. Theory* 28 (1982) 1–18.
- [22] H. Karloff, *Linear Programming*, Birkhäuser Verlag, Basel, 1991.
- [23] A.S. Kelso Jr., V.P. Crawford, Job matching, coalition formation and gross substitutes, *Econometrica* 50 (1982) 1483–1504.
- [24] A.D. Korshunov, O Chisle Monotonnykh Bulevykh Funktsii, *Problemy Kibernetiki* 38 (1981) 5–108 (in Russian).
- [25] E. Kushilevitz, N. Nisan, *Communication Complexity*, Cambridge University Press, Cambridge, 1997.
- [26] A. Kwasmica, J. Ledyard, D. Porter, C. DeMartini, A new and improved design for multi-object iterative auctions, *Management Sci.*, forthcoming.
- [27] S. Lahaie, D. Parkes, Applying learning algorithms to preference elicitation in combinatorial auctions, Working paper, Harvard University, 2004.
- [28] B. Lehmann, D. Lehmann, N. Nisan, Combinatorial auctions with decreasing marginal utilities, in: *Proceedings of the Third ACM Conference on Electronic Commerce*, ACM Press, New York, 2001, pp. 18–28.
- [29] D. Lehmann, L. O’Callaghan, Y. Shoham, Truth revelation in approximately efficient combinatorial auctions, *J. Assoc. Comput. Mach.* 49 (2002) 577–602.
- [30] Z.-Q. Luo, J.N. Tsitsiklis, Communication complexity of algebraic computation, *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science*, 1991, pp. 758–765.
- [31] T. Marschak, On economies of scope in communication, *Econ. Design* 2 (1996) 1–30.
- [32] J. Marschak, R. Radner, *Economic Theory of Teams*, Yale University Press, New Haven, 1972.
- [33] A. Mas-Colell, Efficiency and decentralization in the pure theory of public goods, *Quart. J. Econ.* 94 (1980) 625–641.
- [34] P. Milgrom, Putting auction theory to work: the simultaneous ascending auction, *J. Polit. Economy* 108 (2000) 245–272.
- [35] K. Mount, S. Reiter, The information size of message spaces, *J. Econ. Theory* 28 (1974) 1–18.

- [36] D.C. Parkes, *Ibundle: an efficient ascending price bundle auction*, Proceedings of the First ACM Conference on Electronic Commerce (1999), pp. 148–157.
- [37] D.C. Parkes, *Price-based information certificates for minimal-revelation combinatorial auctions*, in: J. Padget, et al. (Eds.), *Agent-mediated Electronic Commerce-Designing Mechanisms and Systems: Revised Papers from the AAMAS 2002 Workshop on Agent Mediated Electronic Commerce (Lecture Notes in Computer Science)*, Springer, Berlin, New York, 2002, pp. 103–122.
- [38] D.C. Parkes, L. Ungar, *An ascending-price generalized Vickrey auction*, Working paper, Harvard University, 2002.
- [39] S. Reichelstein, *Incentive compatibility and informational requirements*, *J. Econ. Theory* 34 (1984) 32–51.
- [40] M.H. Rothkorf, A. Pekeč, R.M. Harstad, *Computationally manageable combinatorial auctions*, *Management Sci.* 44 (1998) 1131–1147.
- [41] T. Sandholm, S. Suri, A. Gilpin, D. Levine, *Cabob: a fast optimal algorithm for combinatorial auctions*, 17th International Joint Conference on Artificial Intelligence, Morgan Kaufmann, 2001.
- [42] F. Sato, *On the informational size of message spaces for resource allocation processes in economies with public goods*, *J. Econ. Theory* 24 (1981) 48–69.
- [43] C.E. Shannon, *A mathematical theory of communication*, *Bell System Tech. J.* 27 (1948) 379–423, 623–656.
- [44] G. Tian, *A unique informationally efficient allocation mechanism in economies with consumption externalities*, *Internat. Econ. Rev.* 45 (2004) 79–111.
- [45] T. Van Zandt, *Decentralized information processing in the theory of organizations*, in: Murat Sertel (Ed.), *Contemporary Economic Issues, Vol. 4: Economic Design and Behavior*, MacMillan Press Ltd., London, 1999, pp. 125–160.
- [46] V.V. Vazirani, *Approximation Algorithms*, Springer, Berlin, New York, 2001.
- [47] R. Vohra, S. de Vries, *Combinatorial auctions: a survey*, *INFORMS J. Comput.* 15 (2003) 284–309.
- [48] M. Walker, *On the informational size of message spaces*, *J. Econ. Theory* 15 (1977) 366–375.
- [49] A.C.-C. Yao, *Some complexity questions related to distributive computing*, 34th Annual ACM Symposium on Theory of Computing, 1979, pp. 209–213.
- [50] E. Zurel, N. Nisan, *An efficient approximate allocation algorithm for combinatorial auctions*, ACM Conference on Electronic Commerce, 2001.