The Communication Requirements of Efficient Allocations and Supporting Lindahl Prices[*]

Noam Nisan[†]   Ilya Segal[‡]

March 23, 2003

Abstract

We show that any communication finding a Pareto efficient allocation in a private-information economy must also discover supporting Lindahl prices. In particular, efficient allocation of L indivisible objects requires naming a price for each of the $2^L - 1$ bundles. Furthermore, exponential communication in L 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 valuations are submodular, efficiency still requires exponential communication (and fully polynomial approximation is impossible). When the objects are homogeneous, arbitrarily good approximation is obtained using exponentially less communication than that needed for exact efficiency.

[†]Institute of Computer Science, Hebrew University, Jerusalem, Israel, email: noam@cs.huji.ac.il

[‡]Department of Economics, Stanford University, Stanford, CA 94305; email: ilya.segal@stanford.edu

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 Vohra and de Vries (2002) 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 one billion numbers, which is beyond the capabilities of any human or machine.

Recognition of the communication problem has prompted researchers to propose simpler mechanisms, in which valuations are not 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 (see, e.g., Parkes (1999), Bikhchandani et al. (2001), Parkes and Ungar (2001), Ausubel and Milgrom (2002), Kwasnica et al. (2002), Zinkevich et al. (2003)). The hope was that such designs could achieve or at least approximate efficiency, while allowing bidders to communicate much less than their complete preferences.

This paper demonstrates that the hope is not justified. We prove this by showing the crucial role of prices in finding efficient allocations. This role was discussed already in the early 20th-century debate on central planning alternatives to the market system (which is an intellectual antecedent to the current literature on auction design). Hayek (1945) argued that prices succinctly summarize the knowledge of "particular circumstances of time and place" that is too enormous to be communicated to a central planner. Hurwicz (1960) and Mount and Reiter (1974) formalized Hayek's intuition by showing that in

classical convex economies, the Walrasian price mechanism indeed verifies Pareto efficient allocations with the least amount of communication (as measured by the dimensionality of the message space). However, to our knowledge, nobody has shown that prices play an indispensable role even in *nonconvex* allocation problems (such as the combinatorial allocation problem). It has only been known that in such economies, a Walrasian equilibrium need not exist, and more extensive communication may be needed (see Calsamiglia (1977)).

The main insight of the present paper is that prices play an indispensable role in *any* social choice problem with privately known preferences. Namely, we show that in any such problem, *any* communication mechanism that finds a weakly Pareto efficient allocation must also discover a Lindahl equilibrium supporting it. In such an equilibrium, each agent is given a "budget set" (which could be delineated by quoting personalized prices for all allocations), and the agents' optimal choices are consistent with each other. Our result 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.[1]

In the same spirit, Parkes (2002) shows that prices must be revealed when a particular restricted communication language is used to certify efficiency in the combinatorial allocation problem. Unlike Parkes's result, our result allows fully general communication and holds in a general social choice problem.

In the classical convex economy one can restrict attention to Lindahl prices that are Walrasian (i.e., anonymous and linear in consumption). Such prices are easy to communicate—easier even than the allocation itself. In the combinatorial allocation problem, on the other hand, the space of Lindahl prices required to ensure equilibrium existence is huge. Specifically, we show that at least one price should be named for each of the $2^L - 1$ possible bundles of objects. Thus, unlike in the classical economy, the communication of supporting prices is at least as hard as a full revelation of one agent's

---

[1]Our result also implies that the simpler "nondeterministic" problem of *verifying* the efficiency of a proposed allocation is exactly that of announcing supporting Lindahl prices along with the allocation.

preferences, and exponentially harder than the communication of the allocation itself.

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 (1979) and surveyed in Kushilevitz and Nisan (1997).[2] Since exact efficiency in the discretized problem still requires the communication of (discrete) Lindahl prices, we are able to show that guaranteeing a better approximation of efficiency than that achieved by auctioning off all objects as a bundle still requires communicating a very large number of bits, which grows exponentially with L.

The concept of approximation used above requires *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 than that from the bundled auction still requires communicating an exponential expected number of bits.[3]

These results imply that the only hope to achieve or approximate efficiency without

---

[2]The general communication complexity problem is to compute a function (in our case, desired allocation) whose inputs (in our case, agents' preferences) are distributed among agents. The communication complexity literature has developed in parallel with, and shares many techniques with, the economic literature on real-valued communication. For more detailed comparisons of the two literatures, see Marschak (1996) and Van Zandt (1999).

[3]We prove this particular result using a different technique. The proof uses a lower bound on the communication complexity of approximate set packing, which is derived in the Appendix.

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 (1982) and Gul and Stacchetti (1999). 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 objects 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 (1977), 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 can be achieved with only a slight sacrifice in economic efficiency.

The paper is organized as follows. In Section 2 we formulate the general social choice problem and the particular case of the combinatorial allocation problem. In Section 3 we discuss the model of communication and the measures of the communication burden. In Section 4 we characterize efficient communication as that discovering a Lindahl equilibrium. In Section 5 we use this characterization to derive a lower bound on the burden of efficient communication. In Section 6 we define the concepts of approximation and relate them to the analysis of discretized problems. In Section 7 we apply the results to the combinatorial allocation problem. In Section 8 we examine the problem with several

4

restricted classes of valuations. Section 9 presents average-case analysis. Section 10 discusses the relation of our results to the computational complexity literature. Section 11 discusses how agents could be given the incentives to obey the suggested communication protocols. Section 12 concludes.

2 The Social Choice Problem

First we describe the general *social choice problem*. Let N be the finite set of agents, and X be the set of social alternatives. Let $N \geq 2$ to avoid the trivial case. (With a slight abuse of notation, we will use the same letter to denote a set and its cardinality when this causes no confusion.) Let R_i denote the set of possible preference relations of agent $i \in N$ over set X. (We can restrict attention to preference relations that are complete and transitive, though this is not needed for the formal analysis.) Agent i's preference relation $\succeq_i \in R_i$ is assumed to be his privately observed *type*. A *state* is a preference profile $\succeq = (\succeq_1, \ldots, \succeq_N) \in R_1 \times \ldots \times R_N \equiv R$.

Let $\bar{X} \subset X$ denote the set of *feasible* alternatives.[4] The goal is to implement a choice correspondence $F : R \twoheadrightarrow \bar{X}$. In particular, we will be interested in implementing the *(weakly Pareto) efficient* choice correspondence F^*. Letting $L(x, \succeq_i) = \{y \in X : x \succeq_i y\}$ denote the "lower contour set" of preference $\succeq_i$ at alternative x, F^* can be defined as

$$F^*(\succeq) = \left\{ x \in \bar{X} : \bar{X} \subset \cup_i L(x, \succeq_i) \right\}.$$

In words, a feasible alternative x is efficient if any other feasible alternative makes at least one agent weakly worse off.

Now we describe several progressive specializations of the social choice problem. In a *problem with numeraire*, the set of alternatives takes the form $X = \bar{X} = K \times \left\{ t \in \mathbb{R}^N : \sum_i t_i = 0 \right\}$, with $t_i \in \mathbb{R}$ representing the *(monetary) transfer* of numeraire to agent i and $k \in K$ the *(non-monetary) allocation*, and each agent i's preferences $\succeq_i \in R_i$

---

[4]We allow for $\bar{X} \neq X$ to allow, for example, for the Walrasian mechanism, in which an agent may be asked to contemplate allocations that are not feasible for the whole economy.

are continuous, monotonic, and locally nonsatiated in his own transfer t_i, and do not depend on other agents' transfers t_{-i}, for each $k \in K$.

A *quasilinear problem* is a problem with numeraire in which the agents' preferences are quasilinear in monetary transfers. Formally, for each preference $\succeq_i \in R_i$ there exists a non-monetary utility function $u_i : K \to \mathbb{R}$ such that $(k', t') \succeq_i (k, t)$ if and only if $u_i(k') + t'_i \geq u_i(k) + t_i$. In a quasilinear problem, agents' types can be identified with their non-monetary utility functions, and the state space with the space $U = U_1 \ldots \times U_N \subset \mathbb{R}^{NK}$ of possible utility profiles.

In a quasilinear problem, the Pareto efficient allocations (both weakly and strictly) are those that maximize the *surplus* (the sum of the agents' utilities). Formally, letting

$$S(u) = \max_{k \in K} \sum_i u_i(k)$$

denote the maximum surplus available in state u, the efficient correspondence is

$$F^*(u) = \left\{ (k, t) \in K \times \mathbb{R}^N : \sum_i u_i(k) = S(u) \right\}.$$

Finally, the *combinatorial allocation problem* is a quasilinear problem with numeraire in which the nonmonetary allocation describes the allocation of L items among the N agents. Thus, the allocation set is $K = N^L$, where $k(l)$ denotes the agent holding object $l \in L$ in allocation $k \in K$.

3   Communication

We now describe the communication procedures used to solve the social choice problem. It is well known that communication can be shortened by letting agents send messages sequentially rather than simultaneously. For example, agent i does not need to report his preference between alternatives x and y if it is clear from the preceding messages that y is strictly dominated by z by all agents, and so does not stand a chance of being efficient. Therefore, we must consider general multi-stage communication protocols, which we describe below (see also Kushilevitz and Nisan (1997)).

In the language of game theory, a communication protocol is an extensive-form message game along with the agents' strategies in this game (complete action plans contingent on their types and observed history). Instead of payoffs, the game assigns alternatives to terminal nodes (and so it is more properly called a "game form" or "mechanism"). The agents are assumed to follow the prescribed strategies (see Section 11 for a discussion of incentive compatibility). 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* choice correspondence F if in every state $\succeq$ it achieves a terminal node to which an alternative from $F(\succeq)$ is assigned.

Dealing with deterministic communication protocols is quite cumbersome. Analysis is drastically simplified by considering what is called "nondeterministic communication" in computer science and the "verification scenario" in economics. Imagine an omniscient oracle who knows the state of the world $\succeq$ and consequently knows a "desirable" alternative $x \in F(\succeq)$. However, he has to prove to an ignorant outsider that alternative x is indeed desirable. He does this by publicly announcing a message $m \in M$. Each agent i either accepts or rejects the message, doing this on the basis of his own type $\succeq_i$. The set of messages acceptable to all agents in state $\succeq$ is described by the message correspondence $\mu(\succeq)$. The acceptance of message m by all agents must prove to the outsider that alternative x is desirable. (This communication is called "nondeterministic" because if the oracle is not omniscient, he has to "guess" a message that is acceptable to all agents, and there may be more than one such message in a given state.) Formally, nondeterministic communication is defined as follows:

Definition 1 *A nondeterministic communication protocol is a triple* $\Gamma = \langle M, \mu, h \rangle$, *where M is the message set, $\mu : R \twoheadrightarrow M$ is the message correspondence, and $h : M \rightarrow \bar{X}$ is the outcome function, and the message correspondence μ has the following two properties:*

- Existence: $\mu(\succeq) \neq \emptyset$ *for all* $\succeq \in R$,

- Privacy Preservation: $\mu(\succeq) = \cap_i \mu_i(\succeq_i)$ *for all* $\succeq \in R$, *where* $\mu_i : R_i \twoheadrightarrow M$ *for all* $i \in N$.

Protocol Γ realizes *choice correspondence* $F : R \twoheadrightarrow \bar{X}$ *if* $h(\mu(\succeq)) \subset F(\succeq)$ *for all* $\succeq \in R$.

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 he accepts is a function $\mu_i(\succeq_i)$ of his own type $\succeq_i$ only.[5] Finally, the definition of realization says that the acceptance of message m by all agents proves to the outsider that $h(m) \in F(\succeq)$.[6]

While nondeterministic communication is patently unrealistic, there are at least three reasons to examine it:

1. Every deterministic communication can be represented as nondeterministic simply 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 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 will apply to deterministic protocols as a particular case.

---

[5] This term is established in the economic literature, but differs from the everyday concept of "privacy." Privacy Preservation captures the fact that the agents do not observe each other's types *initially*, but imposes no constraints on the revelation of information in the course of communication.

[6] Definition 1 has a nice interpretation in terms of the geometric properties of the subsets $\mu^{-1}(m)$ of the state space $R = R_1 \times \ldots R_N$, describing the event in which message m occurs. In this interpretation, Existence means that the collection $\{\mu^{-1}(m)\}_{m \in M}$ is a covering of R. Privacy Preservation means that each element of the covering is a product set $\mu_1^{-1}(m) \times \ldots \times \mu_N^{-1}(m)$—a "rectangle" in computer science parlance (see Kushilevitz and Nisan (1997)). The definition of realization means that for each rectangle $\mu^{-1}(m)$ from the covering there exists a single outcome $h(m)$ that is "desirable" on the whole rectangle (in computer science parlance, the rectangle is "monochromatic").

2. 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 quantities.[7] Each agent accepts the announcement if and only if his announced allocation constitutes his optimal choice from the budget set given by the announced prices. This nondeterministic protocol realizes the Pareto efficient correspondence in classical economies.

3. A nondeterministic protocol realizing choice correspondence 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 "tatonnement." Nondeterministic communication can then viewed as a steady state of the deterministic adjustment process.

The (nondeterministic) communication burden of a choice correspondence is defined as the minimum communication burden of a (nondeterministic) protocol realizing it. The communication burden of a protocol is naturally defined as the number of messages transmitted in it. Since this number may differ across states, we will focus on the "worst-case" communication burden—the maximum number of messages sent over all states.[8] For this measure to be interesting, the amount of information that can be sent with each message should be bounded.

The computer science literature considers discrete communication, in which case it is natural to restrict each message to take only two possible values, i.e., convey a bit of information.[9] In particular, in the nondeterministic case, the oracle who encodes his

---

[7]To better capture Hayek's (1945) idea of "decentralization," the choice of quantities can be delegated to the agents, provided that the uniqueness of choices is ensured by the strict convexity of preferences. Such communication processes have been called "parametric" in the economic literature.

[8]One may instead be interested in the *expected* number of messages transmitted given some probability distribution over states. Such *distributional* complexity is considered in Section 9 below.

[9]Any other finite-valued message could be coded using bits.

message from set M optimally takes $\log|M|$ bits to transmit it.[10]

In the case of continuous communication, agents are allowed to send real-valued messages. We also want to allow finite-valued messages (say, to communicate discrete allocations), but they will not be counted towards the communication burden. Thus, the worst-case burden of continuous communication can be defined as the maximum number of real-valued messages sent in the course of the protocol. In the nondeterministic case, we can let the oracle's message space M be an arbitrary metric space, and measure the communication burden with a metric dimension of M, such as the Hausdorff dimension, the box dimension, or the packing index (Edgar (1990)). This dimension measure will conveniently count the real-valued messages and ignore the discrete messages.[11] It is well known that to avoid the "smuggling" of multidimensional information into a one-dimensional message space (using, e.g., the inverse Peano function), a smoothness restriction must be imposed on the message correspondence. In Section 5 we will formulate such a restriction and motivate it by the protocol's robustness to noise or discretization in the communication channel.

---

[10]We will use log to mean $\log_2$.

[11]Other papers have examined various *topological* notions of message space dimension, ruling out dimension smuggling by requiring the message correspondence to have a locally continuous selection (Mount and Reiter (1974, 1977), Walker (1977), Abelson (1980), Luo and Tsitsiklis (1993)). We cannot use this approach because finite-valued messages (such as allocations) must be discontinuous in the state, and therefore may not preserve topological dimension. For example, in the problem of allocating one object between two agent, consider the efficient protocol announcing an irrational supporting price p when giving the object to agent $i = 1$ and a rational supporting price when giving it to agent $i = 2$. The set of possible messages $(p, i) \in \mathbb{R}^2$ has topological dimension zero, but its metric dimension is 1, which is consistent with our intuition that the protocol sends one real number. In general, metric dimension either equals or exceeds topological dimension (Edgar (1990, Theorems 6.2.9, 6.5.15)).

4 Efficient Communication and Lindahl Equilibria

In this section we characterize efficient communication as that producing a Lindahl equilibrium.

Definition 2 $(B, \hat{x}) \in 2^{XN} \times \bar{X}$ *is a Lindahl equilibrium in state* $\succeq \in R$ *if*

$$B_i \;\; \subset \;\; L(\hat{x}, \succeq_i) \;\; \text{for all } i \in N, \tag{L1}$$

$$\bar{X} \;\; \subset \;\; \cup_i B_i. \tag{L2}$$

In words, the designer announces a feasible alternative $\hat{x} \in \bar{X}$ and a *budget set* $B_i \subset X$ for each agent i. (L1) means that no agent i can improve upon $\hat{x}$ by choosing from his budget set. (L2) means that the designer must cover the feasible set with budget sets, i.e., make sure that every feasible alternative is in someone's budget set. Let $E : U \twoheadrightarrow 2^{XN} \times \bar{X}$ denote the Lindahl equilibrium correspondence.

A Lindahl equilibrium for two agents is depicted in Figure 1. The budget sets of agents 1 and 2, represented by the shaded and crossed area, respectively, cover the feasible set, represented by the box. Without loss of generality we can include $\hat{x}$ in all agents' budget sets (though this is not needed for the formal arguments). Given the depicted indifference curves, $\hat{x}$ is each agent's preferred choice from his respective budget set.

A budget set can always be described by assigning a price to each alternative, and including those alternatives whose price is no higher than that of $\hat{x}$. Furthermore, in a problem with numeraire, prices can be defined in terms of the numeraire. For this purpose, we can without loss of generality restrict attention to budget sets that are closed and allow free disposal of numeraire. Any such set can be described as $B_i = \{(k, t) \in X : t_i + p_i(k) \leq \hat{t}_i\}$, where

$$p_i(k) = \inf \{\hat{t}_i - t_i : (k, t) \in B_i\} \tag{1}$$

—the minimum amount of numeraire that agent i must sacrifice to achieve non-monetary allocation k. Thus, finding a Lindahl equilibrium is synonymous with finding supporting

Lindahl prices $p_i(\cdot)$.[12] This can be seen graphically in Figure 1, if we interpret the feasible box as an Edgeworth box whose vertical dimension represents the allocation of numeraire, and the horizontal dimension is the non-monetary allocation k. The prices $p_1(\cdot), p_2(\cdot)$ then describe the "price curves"—the frontiers of the agents' budget sets. (L2) then means that the price curve of agent 1 cannot fall below that of agent 2. More generally, (L2) amounts to requiring that $\sum_i p_i(k) \leq 0$ for all $k \in K$. Since (L1) also implies that $p_i(\hat{k}) \geq 0$ for each i, we must have $p_i(\hat{k}) = 0$ for all i. Therefore, conditions (L1),(L2) can be rewritten in terms of prices as

$$\left(\hat{k}, \hat{t}_i\right) = \left(\hat{k}, \hat{t}_i - p_i\left(\hat{k}\right)\right) \succeq_i \left(k, \hat{t}_i - p_i(k)\right) \text{ for all } i \in N, k \in K, \quad \text{(L1')}$$

$$\sum_i p_i(k) \leq \sum_i p_i\left(\hat{k}\right) \text{ for all } k \in K. \quad \text{(L2')}$$

(L1') says that no agent i can be better off by choosing another allocation given the prices, and (L2') says that the same is true for the designer whose objective is the revenue extracted from the agents.

We now state the main result of the paper:

Proposition 1 *Protocol* $\langle M, \mu, h\rangle$ *realizes the weakly Pareto efficient correspondence* F^* *if and only if there exists an assignment* $B : M \to 2^{XN}$ *of budget sets to messages such that protocol* $\langle M, \mu, (B, h)\rangle$ *realizes the Lindahl equilibrium correspondence* E.

Proof. The "if" part": by the definition of Lindahl equilibrium, for any $m \in \mu(u)$, $\bar{X} \subset \cup_i B_i(m) \subset \cup_i L(h(m), \succeq_i)$, therefore $h(m) \in F^*(\succeq)$.

For the "only if" part, suppose protocol $\langle M, \mu, h\rangle$ realizes F^*. For each $m \in M$, let $\hat{x} = h(m)$, and let $B_i(m) = \cap_{\succeq_i \in \mu_i^{-1}(m)} L(\hat{x}, u_i)$ for each i. By construction, (L1) holds for each $\succeq \in \mu^{-1}(m)$. At the same time, using Privacy Preservation, we can write

$$\cup_i B_i(m) = \cup_i \cap_{\succeq_i \in \mu_i^{-1}(m)} L(\hat{x}, \succeq_i) = \cap_{\succeq \in \mu^{-1}(m)} \cup_i L(\hat{x}, \succeq_i) \supset \cap_{\succeq \in \mu^{-1}(m)} \bar{X} = \bar{X},$$

---

[12] If agent i's utility is known to depend only on component k_i of the allocation (say, his own consumption of private goods), then without loss of generality we can restrict attention to a price p_i for him that has the same property.

because $\hat{x} \in F^*(\succeq)$ for each $\succeq \in \mu^{-1}(m)$. Thus, (L2) also holds. Therefore, $(B(m), h(m))$ is a Lindahl equilibrium in every state $\succeq \in \mu^{-1}(m)$. ∎

The "if" part of Proposition 1 is the First Welfare Theorem for Lindahl equilibria: it says that any equilibrium is efficient. The "only if" part is a strengthening of the Second Welfare Theorem: it says not only that supporting budget sets exist for any efficient alternative $\hat{x}$ (which is trivial because we can take $B_i = L(\hat{x}, \succeq_i)$ for all i), but also that they must be revealed in the course of any efficient communication. This result significantly generalizes a result of Parkes (2002) that shows the necessity of using prices when certifying efficiency using a particular restricted communication language in the combinatorial allocation problem.

The proof of Proposition 1 is illustrated in Figure 2. Suppose that we have observed an efficient protocol producing message m and resulting in allocation $\hat{x} = h(m)$. Due to Privacy Preservation, $\hat{x}$ must be efficient in any state in which agent 1's preferences are consistent with m and agent 2's preferences are consistent with m. Graphically, *any* indifference curve of agent 1 consistent with m that passes through $\hat{x}$ must be separated from *any* indifference curve of agent 2 consistent with m that passes through $\hat{x}$. This implies that if we let agent 1's price curve be the lower envelope of his indifference curves consistent with m passing though $\hat{x}$ (in Figure 2, $\succeq_1$ and $\succeq_1'$), and let agent 2's price curve be the *upper* envelope of his indifference curves consistent with m passing though $\hat{x}$ (in Figure 2, $\succeq_2$ and $\succeq_2'$), these price curves will be separated. In other words, the two budget sets delineated by these price curves will cover the feasible box. At the same time, by construction of the budget sets, $\hat{x}$ is an optimal choice for both agent 1 and agent 2 in all states consistent with m. Thus, on the basis of a message m we have constructed a Lindahl equilibrium supporting $\hat{x}$ in all states in which m occurs.

Proposition 1 means that the communication burden of efficiency is exactly that of finding a Lindahl equilibrium. This applies to both deterministic and nondeterministic communication. The nondeterministic version is particularly simple because to verify that we have a Lindahl equilibrium it suffices to announce it. Formally, we can take a

message space $M \subset 2^{XN} \times \bar{X}$, outcome function $h(B, x) \equiv x$, and message correspondence $\mu = E$. Since E satisfies Privacy Preservation (each agent can verify his own (L1) and any agent can verify (L2)), we only have to choose message space M large enough to ensure that a Lindahl equilibrium from M exists in all states. The minimum size of such space gives exactly the nondeterministic communication burden of efficiency.

The simplest example of efficient communication is the protocol in which all agents reveal their preferences and an efficient alternative $\hat{x}$ is chosen. This protocol discovers the Lindahl equilibrium with budget sets $B_i = L(\hat{x}, \succeq_i)$ for all i. A small communication savings is achieved by having only agents $i = 1, \ldots, N-1$ fully reveal their preferences and letting agent N choose an efficient alternative $\hat{x}$. This protocol discovers the Lindahl equilibrium with budget sets $B_i = L(\hat{x}, \succeq_i)$ for $i < N$ and $B_N = \bar{X} \setminus \cup_{i \leq N-1} L(\hat{x}, \succeq_i)$.

An example of far less extensive efficient communication obtains in the *classical economy* with L divisible goods. Here $X = \mathbb{R}_+^{NL}$, the agents' preferences are convex, monotonic, and locally non-satiated, and the feasible set is[13]

$$\bar{X} = \left\{ x \in \mathbb{R}_+^{NL} : \sum_i x_i = \bar{x}, \, x_i \neq 0 \text{ for all } i \right\}.$$

A Lindahl equilibrium $(B, \hat{x}) \in 2^{XN} \times \bar{X}$ is a *Walrasian equilibrium* if for some price vector $p \in \mathbb{R}_+^L$, each budget set takes the form $B_i = \{x \in X : p \cdot x_i \leq p \cdot \hat{x}_i = p \cdot \omega\}$ for some endowment allocation $\omega \in X$. (See Figure 3 for an illustration, and observe that (L2) holds). The classical Welfare Theorems say that an allocation is efficient if and only if it is supported with a Walrasian equilibrium. Furthermore, for the classical economy with smooth utility functions, Proposition 1 can be strengthened to say that any communication producing an interior efficient allocation must also produce a *Walrasian equilibrium* (the price vector for message m could be taken to be the gradient of *any* agent's utility function in *any* state consistent with m—all such gradients are collinear by the first-order conditions for efficiency). Thus, the nondeterministic communication

---

[13] If some agents could be allocated zero bundles, then doing so would reduce the communication burden of Pareto efficiency. For example, efficiency could be achieved trivially without any communication by giving all goods to one agent.

burden of efficiency is exactly the minimum dimension of a price-allocation subspace ensuring the existence of a Walrasian equilibrium from this subspace. This dimension can be bounded below by considering the "fooling set" of *Cobb-Douglas* states, in which each agent i's utility takes the form $u_i(x_i) = \prod_l x_{il}^{\alpha_{il}}$ with $\sum_l \alpha_{il} = 1$.[14] Since no two distinct Cobb-Douglas states can share a Walrasian equilibrium, the required dimension of the price-allocation subspace must be at least as large as the dimension of the fooling set, which is $N(L-1)$. This lower bound is achieved by the Walrasian protocol with a fixed endowment allocation ω: a normalized price vector is communicated with $L-1$ real numbers, and a feasible allocation at which all agents' budget constraints bind is communicated with $(N-1)(L-1)$ real numbers.

5 A Lower Bound on Efficient Communication

In this section we consider nonconvex allocation problems, in which a Walrasian equilibrium need not exist. We construct a fooling set to bound below the dimensionality of the required Lindahl price space, and therefore, by Proposition 1, the communication burden of efficiency. To simplify analysis, from now on we restrict attention to quasilinear problems.

Note that in a quasilinear problem, two utility functions that differ by a constant describe the same preferences. In counting the states of the world, we will need to ensure that we do not count two different states describing the same preferences. Also, in measuring the metric dimension of (subsets of) the state space, we need to ensure that the dimensionality of utility functions coincides with the dimensionality of preferences. For these purposes, we assume that any two preferences that are "close" to each other

---

[14] A "fooling set" in computer science terminology is a set of instances of the problem which can already be shown to be hard to solve (the problem can only become harder when all possible instances are allowed). A corresponding term in the economic literature on the dimensionality of message spaces is "a set with the uniqueness property." The fooling set consisting of Cobb-Douglas states was first proposed by Hurwicz (1960).

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 $u, u' \in U$,*[15]

$$\|u'_i - u_i\| \leq C \sup_{i \in N, j, k \in K} |[u'_i(j) - u_i(j)] - [u'_i(k) - u_i(k)]|.$$

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 feasible alternatives are equally efficient constitutes a "fooling set" for the Lindahl equilibrium correspondence. Namely, letting

$$U^* = \left\{ u \in U : F^*(u) = \bar{X} \right\},$$

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

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

Proof. Suppose in negation that $(B, \hat{x}) \in E(u) \cap E(u')$. Let $p \in \mathbb{R}^{NK}$ denote the Lindahl prices corresponding to budget sets B according to (1). Letting $\hat{x} = \left(\hat{k}, \hat{t} \right)$, and using the fact that $u \in U^*$, we can write (L2′) in state u as

$$\sum_i \left[u_i \left(\hat{k} \right) - p_i \left(\hat{k} \right) \right] \leq \sum_i [u_i(k) - u_i(k)] \text{ for all } k \in K.$$

On the other hand, (L1′) means that

$$u_i \left(\hat{k} \right) - p_i \left(\hat{k} \right) \leq u_i(k) - p_i(k) \text{ for all } i \in N, k \in K.$$

Comparing the two displays, we see that for each i, $u_i(k) - p_i(k)$ does not depend on $k \in K$. Since by the same argument the same is true for $u'_i(k) - p_i(k)$, we see that

---

[15]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 .

$u_i'(k) - u_i(k)$ does not depend on k for all i. By normalization, this implies that $u = u'$ —a contradiction. ∎

The proof of Lemma 1 for the case $N = 2$ is illustrated in Figure 4. In any state from U^*, the two agents' indifference curves coincide, and both agents' supporting Lindahl 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. The graphical proof can be extended to $N > 2$ by observing that in any state from U^*, the indifference curves of an agent i coincide with the indifference curves of the aggregate utility of the other agents, and therefore again pin down the prices facing agent i.

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. For the case of continuous messages, we would like to say similarly that the dimension of the message space is at least that of U^*. For this purpose, we need to rule out "dimension smuggling," which we do as follows:[16]

Definition 4 *Protocol* $\langle M, \mu, h \rangle$ *satisfies* Lipschitz Lower HemiContinuity (LLHC) *if there exists* $A > 0$ *such that for any* $m, m' \in M$ *and* $u \in \mu^{-1}(m)$, *there exists* $u' \in \mu^{-1}(m')$ *satisfying* $\|u' - u\| \leq A\rho(m', m)$ *(where* ρ *is the metric on* M*).*

Intuitively, LLHC ensures that a small error in the message does not result in a large misrepresentation of the state of the world. This property guarantees that a small noise or discretization error in the communication channel would not degrade the performance of the mechanism drastically (Propositon 4 below states a formal result along these lines).

Proposition 2 *In a normalized quasilinear problem, any efficient protocol transmits at least* $\log |U^*|$ *bits. Moreover, if the protocol is LLHC, the dimension of its message space is at least* $\dim U^*$.

---

[16]A similar Lipschitz restriction on the inverse message correspondence μ^{-1} was imposed by Hurwicz (1960).

Proof. Lemma 1 and Proposition 1 together imply that the restriction of the message correspondence μ to U^* has an injective selection $\sigma : U^* \to \mu(U^*)$. 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 $u' = \sigma^{-1}(m')$. By LLHC, there exists $u'' \in \mu^{-1}(m')$ such that $\|u'' - u\| \leq A\rho(m', m)$. As shown in the proof of Lemma 1, all Lindahl equilibrium prices in state u' coincide with u' up to agent-dependent constants. By Proposition 1, the same prices must also support allocation $h(m') = (\hat{k}, \hat{t})$ in state u''. (L1') then implies that for all $i \in N$, $k \in K$,

$$u_i''(k) - u_i''\left(\hat{k}\right) \leq u_i'(k) - u_i'\left(\hat{k}\right) = u_{-i}'\left(\hat{k}\right) - u_{-i}'(k) \leq u_{-i}''\left(\hat{k}\right) - u_{-i}''(k)$$

(where $u_{-i}(k) = \sum_{j \neq i} u_j(k)$). Subtracting $u_i(k) - u_i\left(\hat{k}\right) = u_{-i}\left(\hat{k}\right) - u_{-i}(k)$, and taking the absolute value implies:

$$
\begin{aligned}
&\left|[u_i'(k) - u_i(k)] - \left[u_i'\left(\hat{k}\right) - u_i\left(\hat{k}\right)\right]\right| \\
\leq\ & \max\left\{\left|[u_i''(k) - u_i(k)] - \left[u_i''\left(\hat{k}\right) - u_i\left(\hat{k}\right)\right]\right|, \left|[u_{-i}''(k) - u_{-i}(k)] - \left[u_{-i}''\left(\hat{k}\right) - u_{-i}\left(\hat{k}\right)\right]\right|\right\} \\
\leq\ & 2(N-1)\|u'' - u\| \leq 2(N-1)A\rho(m', m).
\end{aligned}
$$

Now, using normalization and the triangle inequality,

$$
\begin{aligned}
\|u' - u\| &\leq C \sup_{i \in N, j, k \in K} |[u_i'(j) - u_i(j)] - [u_i'(k) - u_i(k)]| \\
&\leq 2C \sup_{i \in N, k \in K} \left|[u_i'(k) - u_i(k)] - \left[u_i'\left(\hat{k}\right) - u_i\left(\hat{k}\right)\right]\right| \leq 4(N-1)AC\rho(m', m)
\end{aligned}
$$

Thus, σ^{-1} is Lipschitz continuous, which implies that $\dim M \geq \dim \sigma(U^*) \geq \dim U^*$ (see Edgar (1990, Exercise 6.1.9(1))). ∎

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

Corollary 1 *Suppose that in a normalized quasilinear problem with $N = 2$, for each $u_1 \in U_1$ there exists a "dual utility" $u_2 \in U_2$ such that $u_1(k) + u_2(k)$ does not depend on k. Then any efficient protocol transmits at least $\log|U_1|$ bits, and, under LLHC, has the message space dimension of at least $\dim U_1$.*

Corollary 1 will be a workhorse for the subsequent results on the combinatorial allocation problem. We will apply it also to $N > 2$ agents by letting agents $i > 2$ have constant utilities over all allocations k. Also, we will apply Corollary 1 to some cases in which the two agents can have dual utilities only on some subset $\tilde{K} \subset K$ of allocations, by restricting attention to a subclass of utilities for which allocations from $\tilde{K} \backslash K$ are never efficient (see Subsection 8.1). Corollary 1 has two other possible extensions that we do not pursue in this paper: First, it can be stated for non-quasilinear problems. Second, it could be shown not to be a "knife-edge" result: when the two agents' indifference curves cannot *exactly* coincide but can be close to each other, they would almost pin down Lindahl prices. Since two such states that are sufficiently different from each other must have distinct Lindahl prices, we would again obtain a lower bound on the communication burden.

6 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 (maintaining the focus on quasilinear problems).

To be consistent with the computer science literature on approximation (see, e.g., Vazirani (2001)), we use an approximation measure that is invariant to the units of measurement (see, e.g., Vazirani (2001)). Namely, we define the choice correspondence F_r^* realizing approximation ratio $r \in [0, 1]$ as follows:[17]

$$F_r^*(u) = \left\{ (k, t) \in K \times \mathbb{R}^N : \sum_i u_i(k) \geq rS(u) \right\}.$$

By construction, $F_1^* = F^*$ (the exactly efficient correspondence), and $F_r^* \supset F^*$ for $r \leq 1$.

---

[17]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 9 below.

For the study of approximation, we need to assume[18]

$$\overline{S} \equiv \sup_{u \in U} S\left(u\right) < \infty, \, \underline{S} \equiv \inf_{u \in U : S(u) \neq 0} S\left(u\right) > 0. \tag{2}$$

The communication burden of F_r^* can be evaluated by examining the communication burden of the exactly efficient correspondence 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 $U^\delta = \left\{u^\delta : u \in U\right\}$. 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,

---

[18]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 $\overline{S} = \infty$), by considering the restricted problem $\{r^{-n}\}_{n=0}^\infty$.

20

$U^\delta \subset U$, 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 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 a quasilinear problem in which (2) 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 (see, e.g., Vazirani (2001)):

- 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 object 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 V^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, LLHC ensures that rounding off a message results in an allocation that is desirable for some state that is not too far from the true state, and therefore the efficiency loss is not too large. Formally, we have

Proposition 4 *In a quasilinear problem satisfying (2), if there exists an LLHC protocol*

realizing approximation ratio r with a message space of upper box dimension $d > 0$,[19]

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 $\varepsilon \to 0$.

Proof. Suppose that $\Gamma = \langle M, \mu, h \rangle$ is an efficient protocol. By 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 $\Gamma_\delta = \langle M_n, \mu_n, h \rangle$. 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 $m \in \mu(u)$. By LLHC, there exists $u' \in \mu^{-1}(m_\delta)$ such that $\|u' - u\| \leq A\delta$. This implies that (i) $S(u) \leq S(u') + A\delta$, and (ii) $\sum_i u_i(h(m_\delta)) \geq \sum_i u_i'(h(m_\delta)) - A\delta$. Thus, we can write

$$\frac{\sum_i u_i(h(m_\delta))}{S(u)} \geq \frac{\sum_i u_i'(h(m_\delta)) - A\delta}{S(u') + A\delta} \geq \frac{\sum_i u_i'(h(m_\delta)) - 2A\delta}{S(u')} \geq r - 2A\delta/\underline{S},$$

where the last inequality uses the fact that Γ realizes approximation ratio r, hence $h(m_\delta) \in F_r^*(u')$. Therefore, taking $\delta(\varepsilon) = \varepsilon \underline{S}/(2A)$ ensures that $\Gamma_{\delta(\varepsilon)}$ realizes approximation ratio $r - \varepsilon$. The number of bits communicated by $\Gamma_{\delta(\varepsilon)}$ is

$$\log \left| M_{\delta(\varepsilon)} \right| \sim d \log \delta(\varepsilon)^{-1} = d \log \varepsilon^{-1} + d \log (2A/\underline{S}) \sim d \log \varepsilon^{-1} \text{ as } \varepsilon \to 0. \blacksquare$$

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 asymptotically as $\varepsilon \to 0$.[20] Thus, the dimensionality of the message space in a continuous protocol is indeed a relevant measure of communication even if

---

[19]The upper box dimension coincides with the Hausdorff dimension for most "well-behaved" sets, but could exceed it for some sets (see Edgar (1990)).

[20]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 (2003a,b).

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."

7   The Combinatorial Allocation Problem

The combinatorial allocation problem is a quasilinear problem in which the nonmonetary decision describes the allocation of L items among the N agents. Thus, the set of nonmonetary allocations is $K = N^L$, where $k(l)$ denotes the agent holding object $l \in L$ in allocation $k \in K$. We impose several standard restrictions on utilities (note that the restrictions can only reduce the communication burden).

- **No Externalities (NE):** For each i and each $u_i \in U_i$, there exists $v_i : 2^L \to \mathbb{R}$ such that $u_i(k) = v_i(k^{-1}(i))$ 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 $V = V_1 \times \ldots \times V_N$.

We also assume that each $v_i \in V_i$ satisfies 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 $v_i(L) \in [\gamma, 1]$.

(N) is without loss of generality, and it ensures the normalization assumption formulated in Section 5. (M) and (B) are not needed for the analysis of exact efficiency, but will play a role in the analysis of approximation. In particular, (B) ensures that (2) holds, so that approximation can be achieved with finite communication (see footnote

18 above). Given that the analysis so far has been invariant to scale, the choice of the upper bound as 1 is arbitrary. The lower bound γ can be viewed as a parameter of the problem along with N and L, but the value of γ will be irrelevant for most results. Let $V_{gen} \subset \mathbb{R}^{2^L}$ denote the class of all valuations satisfying (NE), (N), (M), and (B). (We will consider additional restrictions on valuations in Section 8.)

To apply Corollary 1, for each valuation $v \in V_{gen}$ define the *dual* valuation $\tilde{v} \in V_{gen}$ by

$$\tilde{v}(S) = v(L) - v(L \backslash S) \text{ for all } S \subset L. \tag{3}$$

When $N = 2$ and the two agents' valuations are v and $\tilde{v}$, all allocations have the same surplus $v(L)$. Thus, we can use Corollary 1 to obtain

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

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 4 that this lower bound is essentially tight for $N = 2$. A more general upper bound is given by $(N - 1)\left(2^L - 1\right)$ numbers (all agents but one fully reveal their valuations), but the exact communication burden for $N > 2$ remains an open problem.[21]

To bound below the communication burden of approximation, we apply the same logic to the discretized problem $V_{gen}^{\delta} \subset V_{gen}$ 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 $\log \left| V_{gen}^1 \right|$ bits. Only counting

---

[21] 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 non-empty 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), and having each agent i accept the communication if for each bundle $S \subset L$, either his valuation for S is announced correctly, or it does not exceed any of the L announced valuations. Therefore, when $N > L$, efficiency can be verified using only $L\left(2^L - 1\right)$ real numbers. We are obliged to Moshe Babayoff for bringing this to our attention.

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 $\log \left| V_{gen}^1 \right| \geq \binom{L}{L/2}$.[22] 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 $\log \left| V_{gen}^1 \right| \geq \binom{L}{L/2}$ bits.*

Observe that approximation ratio $1/N$ can be realized by auctioning off all objects as a bundle to the highest bidder. (Indeed, the bundled auction realizes surplus $\max_i v_i (L)$, 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.[23] 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 $\ln 2 \cdot \exp \left\{ L/ \left(2N^2 \right) - 2\ln N \right\}$ 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

[22]Note that $\left| V_{gen}^1 \right|$ is the number of monotone boolean functions of L boolean variables. The problem of counting these functions is known as "Dedekind's problem," which is unsolved, though its asymptotic behavior is obtained by Korshunov (1981).

[23]Indeed, by Stirling's formula, the communication burden $\binom{L}{L/2} \sim \sqrt{2/ (\pi L)} \cdot 2^L$ as $L \to \infty$.

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. ∎

This result should be contrasted with the findings of Lehmann et al. (1999) and Holzman et al. (2001), who suggest "simple" protocols improving upon the bundled auction. For example, Holzman et al. (2001) note that auctioning off the objects 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. (1999) propose 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. (1999) show 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 no 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 $L^{1/2-\varepsilon}$), "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 $N, L \to \infty$.

8 Restricted Valuations

8.1 Submodular Valuations

Here each agent's valuation space $V_i = V_{sm}$ is the set of all valuations $v \in V_{gen}$ that satisfy

$$v(S \cup T) + v(S \cap T) \leq v(S) + v(T) \text{ for all } S, T \subset L.$$

An equivalent definition of submodularity is that the marginal benefit of each item $l \in L$, $v(S \cup l) - v(S)$, is nonincreasing in $S \subset L$.

Corollary 1 cannot be applied to this case directly, since the dual (3) of a submodular valuation is typically not submodular (unless both are additive—see Subsection 8.3 below). We get around this problem by defining duality in such a way that the surplus is constant only on the allocations involving even splits of objects:

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

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

- $v(S) = 2|S|/L$ for $|S| < L/2$,

- $v(S) = 1$ for $|S| > L/2$,

- $v(S) \in [1 - 1/L, 1]$ for $|S| = L/2$,

- $\frac{1}{|\widetilde{K}|} \sum_{S \subset L : |S| = L/2} v(S) = 1 - 1/(2L)$.

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

Note that in any state $(v^1, v^2) \in \widetilde{V} \times \widetilde{V}$, all efficient allocations lie in $\tilde{K}$. (Indeed, the last bullet implies that the average surplus of all allocations from $\tilde{K}$ is $2 - 1/L$, while any other allocation has a surplus of at most $2 - 2/L$, and thus is dominated by at least one allocation from $\tilde{K}$.) The last bullet above ensures that thus constructed valuation class with allocations restricted to $\widetilde{K}$ is normalized.

For each $v \in \widetilde{V}$, define its "quasi-dual" $\widehat{v} \in \widetilde{V}$ as follows:

- $\widehat{v}(S) = 2 - 1/L - v(L\backslash S)$ for $|S| = L/2$,

- $\widehat{v}(S) = 2|S|/L$ for $|S| < L/2$.

- $\widehat{v}(S) = 1$ for $|S| > L/2$,

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

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

Consider now the discretized problem $\tilde{V}^\delta \subset \widetilde{V}$ with $\delta = 1/L$ and $N = 2$. Applying Corollary 1 with the two agents' valuations restricted to $\tilde{V}^\delta$ and the outcomes restricted to $\widetilde{K}$ implies that the number of bits communicated by an efficient protocol for the discretized problem is at least $\log\left|\tilde{V}^\delta\right| = \left|\widetilde{K}\right| - 1$. 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 Lehmann et al. (2001) that approximation ratio $1/2$ is realized

by the (deterministic) "greedy" protocol that allocates the objects in a fixed order to the agents who announce the highest current marginal benefit for them, and so communicates only NL numbers.

8.2   Homogeneous Valuations

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

Since the dual (3) 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 LLHC protocol is at least* $\dim V_h = L$.

Now consider the discretized problem $V_h^\delta \subset V_h$, with $\delta = 1/R$. 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 $\phi : \{1, \ldots, L\} \to \{0, \ldots, R\}$ 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 $\binom{R+L-1}{L-1}$.[24] Therefore, exact efficiency in problem V_h^δ requires communicating at least $\log \binom{\delta^{-1}+L-1}{L-1}$ 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 $\{0..L\}$, 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

---

[24]This is proven by putting all the R balls in a row with $L-1$ dividers, and letting urn $l = 1, .., L$ 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.

(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. ∎

Part (ii) of the Corollary implies, in particular, that we can achieve any given approximation error using $O(\log L)$ bits. On the other hand, part (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, when L 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 (1977), 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 $v : [0, 1] \to [0, 1]$, and so $\dim V = \infty$. Corollary 1 then implies that efficiency requires an infinite-dimensional message space, re-deriving

31

Calsamiglia's (1977) result.[25],[26] However, Calsamiglia's model allows a FPAS, provided that the agents' valuation functions satisfy

$$|v(x') - v(x)| \leq [-\log|x' - x|]^{-A} \text{ for some } A > 0.$$

Under this mild strengthening of continuity (for example, implied by Hölder continuity of any degree), restricting the 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 requires infinite-dimensional communication.

8.3   Substitute Valuations

Here each agent's valuation space $V_i = V_{sub}$ is the set of valuations $v \in V_{gen}$ whose indirect utility function $w(p) = \max_{S \subset L} \left(v(S) - \sum_{l \in S} p_l \right)$ is submodular in $p \in \mathbb{R}^L_+$. This is one of the many equivalent definitions of the substitute property—see Gul and Stacchetti (1999) and Milgrom (2000).[27]

Since $V_{sub} \subset V_{sm}$ (see Gul and Stacchetti (1999)), the dual (3) 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 $V^{add} \subset V^{sub}$, and the dual of an additive valuation is itself, Corollary 1 yields

---

[25]Calsamiglia (1977) 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.

[26]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).

[27]The property is more widely known as "gross substitutes," which is redundant because in the asbence of wealth effects the concepts of gross and net substitutability coincide.

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

This lower bound is attained by the Walrasian equilibrium with per-object prices, which always exists with substitute valuations (Kelso and Crawford 1982, Gul and Stacchetti 1999).

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 by Gul and Stachetti (2000) and Ausubel (2002). These protocols are variations on the ascending-bid auction with prices quoted for individual items, and so they are only FPAS. This is in fact true of all proposed approximation protocols based on the primal-dual schema (see, e.g., Bikhchandani et al. (2001)).

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 by Bikhchandani and Mamer (1997), 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:

$$\max_{x \in \mathbb{R}_+^{N \cdot 2^L}} \sum_{i \in N, S \subset L} x_S^i v_i(S) \tag{P}$$

$$\text{s.t.} \sum_{i \in N, S \ni l} x_{iS} \leq 1 \text{ for all objects } l \in L,$$

$$\sum_{S \subset L} x_{iS} \leq 1 \text{ for all agents } i \in N.$$

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

$$\min_{p \in \mathbb{R}_+^L, w \in \mathbb{R}_+^N} \sum_{l \in L} p_j + \sum_{i \in N} w_i \tag{D}$$

$$\text{s.t. } w_i - \left[v_i\left(S\right) - \sum_{S \ni l} p_l \right] \geq 0 \text{ for all } i \in N,\ S \subset L,$$

where p_l and w_i denote the Lagrange multipliers with the constraints in (P) associated with object l and agent i 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 (see, e.g., Karloff (1991)). 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.[28]

---

[28] 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

9   Average-case Analysis

Suppose that we are given a probability distribution over the states of the world. Then 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, which allows a savings from coding more frequent messages with fewer bits (as in Shannon's (1948) information theory). 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 outcome is known and can be implemented with no communication. Thus, it is only interesting to consider distributions that are sufficiently diffuse so that no outcome has a high a priori probability of being efficient. It turns out that for *some* such distributions, the communication complexity of approximating efficient combinatorial allocations still grows exponentially with the number of objects:

Proposition 13 *In the combinatorial allocation problem, there exists a sequence of joint probability distributions over valuation profiles $\left(v^1, \ldots, v^N\right)$ 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 $c \exp\left\{L/\left(2N^2\right) - 5\ln N\right\}$, 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

---

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.

complexity and distributional complexity (which follows from the Minimax Theorem—see Kushilevitz and Nisan (1997, 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. ∎

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, Frankl, and Simon (1986):

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 to the "disjointness problem" from communication complexity theory (See Kushilevitz and Nisan 1997). In this problem, two agents are each given a subset of a set M, and the objective is to decide whether the subsets are disjoint. Babai, Frankl, and Simon (1986) prove a lower bound on the distributional complexity of disjointness for product distributions:

Theorem 1 (Babai, Frankl, and Simon (1986)) *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 according to D_1, D_2 (to be defined below) can be used to obtain a protocol for disjointness 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.

Here is the definition of the distributions D_1, D_2 on valuations: Let M be the collection of subsets of L of size exactly $L/2$, hence $|M| = m$. The valuation v is chosen by first choosing a random subset X of M according the distribution D of Babai, Frankl, and Simon (1986). In both D_1 and D_2, we define $v(S) = 0$ for $|S| < L/2$; $v(S) = 1$ for $|S| > L/2$. In D_2 we define for $|S| = L/2$, $v(S) = 1$ if $S \in X$ and $v(S) = 0$ otherwise. In D_2 we define for $|S| = L/2$, $v(S) = 1$ if $N \backslash S \in X$ and $v(S) = 0$ 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 and then solve the combinatorial allocation problem. Finding an allocation with surplus 2 means finding a partition of L into two sets $(S, N \backslash S)$ of size $L/2$ each such that $S \in X_1$ and $S \in X_2$, thus proving X_1 and X_2 are not disjoint. Any inefficient allocation has at most surplus 1. Now, if the allocation

protocol loses at most 0.5% of expected surplus, then the probability that it produces an inefficient allocation in a state with maximum surplus 2 is at most 1%. Thus, if we declare that X_1 and X_2 are disjoint whenever the obtained allocation has value 1, we err with probability of at most 1%. ∎

The proof is done for $c = 99.5\%$, which is derived from a constant quoted in Babai, Frankl, and Simon (1986). No optimization of the constant was attempted and it seems likely that a substantial improvement is possible.

10   Comparison with Computational Complexity

The communication problem examined here is 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. Even in such simple cases from the viewpoint of communication, the efficient combinatorial allocation problem has been shown to be NP-complete.

Nevertheless, we believe that the communication bottleneck is more severe in practice than the computational one. Recall that NP-completeness only indicates that the problem may be 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 (Vohra and de Vries 2002, Sandholm et al. 2001) and thousands of items (with tens of thousand of bids) near-optimally (Zurel and Nisan 2001). 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 objects, 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 $L = 50$ items, any improvement would require the bidders to send at least $\binom{50}{25} \simeq 1.3 \times 10^{14}$ bits $\simeq 500$ Gigabytes of data, which roughly corresponds to 250 million typewritten pages.[29]

A model of computational complexity that is designed to handle large inputs is "black box complexity." In the general model, the algorithm can ask a "black box" (oracle) an arbitrary query about an agent's valuation v^i. (More restricted models define which queries the black box for v^i will answer—e.g., in the "valuation oracle" model only allows the requests to report of the agents' valuations v^i_k for particular outcomes k.) It is easy to see that the number of queries in the general black box model is bounded below by the deterministic communication complexity of the problem.

11 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 the communication game. In this section, we show how in the quasilinear case, these constraints may be satisfied using monetary transfers.

Suppose that after running the protocol, we ask each agent to report his payoff u_i at the resulting allocation, and pay each agent i a transfer $t_i = \sum_{j \neq i} u_i$. This transfer scheme (first proposed by Reichelstein (1984, pp.45-46)) ensures that each agent's total payoff equals the total surplus, and so converts the communication game into one of common interest. (In the terminology of Marschak and Radner (1972), 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

---

[29]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 package (Banks et al. 1989, Nisan and Ronen 2000), all the communication in a combinatorial auction passes through the auctioneer.

an agent can increase his payoff in any state.[30]

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 the agents have a common prior over the states of the world, and that they play a Bayesian-Nash Equilibrium (BNE) of the game defined by the protocol. Since the game is one of common interest under the proposed transfer scheme, the strategy profile maximizing the expected surplus constitutes a BNE of the game. In particular, this BNE cannot have a lower expected surplus than that achieved by the original protocol. In other words, if the agents coordinate on this BNE, it will achieve the best average-case approximation consistent with the game.[31] Of course, this relies heavily on the agents' rationality—both individual (being able to calculate an optimal strategy profile) 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.

A possible criticism of the proposed transfer scheme is that it is very costly. The cost can be covered with lump-sum participation fees, but these fees may be restricted by the agents' participation constraints. The largest fee that still guarantees agent i's participation is the surplus that could be achieved in his absence. (This makes an agent indifferent about participating when he does not contribute anything to the surplus, and strictly prefer to participate otherwise.) Together with the transfer scheme, this gives each agent his Vickrey-Groves-Clarke (VCG) payoff. As noted by Reichelstein (1984) and Ausubel (2002), calculation of the participation fees requires running the protocol with successively removing each agent, which multiplies the communication burden by $N + 1$. However, some efficient protocols (such as those in Feigenbaum et al. (2000) and

---

[30]In general, obedience will not be a dominant strategy, since an agent i may gain from a deviation if he expects another agent j to use a strategy that is not consistent with any type u_j.

[31]To take this observation to its extreme, under our transfer scheme, rational agents need not be offered a protocol at all! Namely, given an explicit expression for communication costs, the agents could be made to internalize these costs (e.g., by paying for sending bits). Then if the agents play some "free form" game in which they can individually send messages and implement an allocation, the protocol that maximizes the expected surplus net of communication costs will constitute a BNE of the game.

Bikhchandani et al. (2001)) yield the VCG transfers as a side product, which provides a substantial communication savings when N is large.

12 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 non-price mechanisms perform better than price mechanisms in nonconvex economies. The designers of combinatorial auctions have proposed numerous designs that purport to find efficient allocations without finding all the prices supporting them.

The present paper has shown that in fact, *any* efficient mechanism is "essentially" a price mechanism, in the sense that it must reveal supporting Lindahl prices (in the general social choice problem, budget sets) along with the efficient allocation itself. Thus, the indispensable role of prices for implementing efficient allocations is now made clear. This result holds regardless of the agents' incentives, even if the agents report truthfully.

In the special case of the combinatorial allocation problem, we demonstrate that efficient communication must name (at least) one Lindahl price for each of the $2^L - 1$ possible bundles, where L is the number of objects. Even if we only require a better approximation of efficiency than that guaranteed by auctioning off all objects as a bundle, participants in a two-budder auction still must send at least $\binom{L}{L/2}$ bits. With $L = 50$ (a realistic number), this amount of inormation roughly corresponds to more than 250 million typewritten pages, and which is not feasible for any reasonable view of communication costs. We also show that even if approximation is required only on expectation, for *some* probability distribution over valuations, it still requires exponential communication in L. These results imply that for realistic values of L, any combinatorial auction design or other "preference elicitation" scheme suggested in the literature would either run for

a prohibitively long time, or, if stopped after some reasonable time, would not produce an efficient or even 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 institution 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, by showing that there does not exist a practical auction design that works well for all possible combinatorial preferences over many objects, we hope to motivate auction designers to focus on specific classes of preferences or probability distributions over them. However, the burden should be on the proposer of a particular design to characterize the environments on which it works well. The tools developed in the present paper will be useful for this purpose, as we have demonstrated by examining the communication burden for 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 reflects the difficulty of approximating efficiency with a discretized mechanism. We found 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 may sometimes be achieved with much less communication. In such cases, the economic measure seriously overstates the "hardness" of the communication problem. A dramatic example of this is offered by Calsamiglia's (1977) model of allocating a homogeneous divisible good, in which exact efficiency requires infinite-dimensional communication, but a fully polynomial approximation exists.

Appendix: A Set Packing Lower Bound

Our lower bound for the set packing problem will use a reduction from the following "approximate-disjoitness" problem which was studied in Alon et al. (1999):

The Approximate Disjointness Problem: N players are each holding an n-bit string. The string for player i specifies a subset $A_i \subseteq \{1...n\}$. They are required to distinguish between the following two extreme cases:

- Negative: $\cap_{i \in N} A_i \neq \emptyset$,

- Positive: for every $i \neq j$, $A_i \cap A_j = \emptyset$.

A lower bound on the required communication of cn/N^4 for some fixed constant $c > 0$ was given in Alon et al. (1999) for randomized protocols (with two sided error). The lower bound was improved by Jaikumar Radhakrishnan and Venkatesh Srinivasan to $\ln(2) \cdot n/N$ for deterministic and nondeterministic protocols.

Let us now define the set packing problem. We have N players, each holding a collection $A_i \subseteq 2^L$. The objective is to approximate the *packing number*—the number of subsets in the union of the collections that are packed together, i.e., are pairwise disjoint. We will prove a lower bound on the communication complexity needed in order to distinguish between the case where the packing number is 1 and the case that it is N. I.e., to distinguish the case where there exists N disjoint sets $S_i \in A_i$ (a N-packing), and the case where any two sets $S_i \in A_i$ and $S_j \in A_j$ intersect (packing number is 1). We will reduce this problem from the approximate-intersection problem on vectors of size $t = e^{L/(2N^2)}/N$. The reduction uses a set of partitions $F = \{P^s | s = 1...t\}$, where each P^s is a partition $(P_1^s, ..., P_N^s)$ of set L into N subsets. This set of partitions will have the following property:

Definition 5 *A set of partitions $F = \{P^s | s = 1...t\}$ has the pair wise-intersection property if for every choice of $1 \leq i \neq j \leq k$ and every $1 \leq s_i \neq s_j \leq t$ we have that $P_i^{s_i} \cap P_j^{s_j} \neq \emptyset$. I.e. that any 2 parts from different partitions intersect.*

Lemma 2 *There exists a set F of partitions with the pair wise-intersection property of size $|F| = t = e^{L/(2N^2)}/N$.*

Proof. We will use the probabilistic method where each partition in the set will be chosen independently at random, and for each partition, each element is placed independently at random in one of the parts of the partition. Now fix $1 \le i \ne j \le k$ and two indices of partitions $1 \le s_i \ne s_j \le t$. The probability that they do not intersect can be calculated as:

$$Pr[P_i^{s_i} \cap P_j^{s_j} = \emptyset] = (1 - 1/k^2)^n \le e^{-n/k^2}$$

Since there are at most $k^2 t^2$ such choices of indices, we get that as long as $k^2 t^2 < e^{n/k^2}$, the required set of partitions exists. ∎

We can now specify the reduction of approximate disjointness on vectors of size t to the approximate set packing problem. Player i who gets as input the set $B_i \subset \{1...t\}$ will construct the collection $A_i = \{P_i^s | s \in B_i\}$. Now, if there exists $s \in \cap_i B_i$ then a N-packing exists: $P_1^s \in A_1...P_N^s \in A_N$. If, one the other hand, for all $i \ne j$, $B_i \cap B_j = \emptyset$, then for any two sets $P_i^{s_i} \in A_i$ and $P_j^{s_j} \in A_j$ we have $s_i \ne s_j$ and thus $P_i^{s_i} \cap P_j^{s_j} \ne \emptyset$.

From the lower bounds described above for the approximate disjointness problem, letting $n = |F|$, we obtain the following lower bounds for the set packing problem:

Theorem 2 *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] Abelson, H. (1980). "Lower Bounds on Information Transfer in Distributed Computations," *Journal of the Association for Computer Machinery* 27(2): 384-392.

[2] Alon, N., Y. Matias, and M. Szegedy (1999). "The Space Complexity of Approximating the Frequency Moments," *Journal of Computer and System Sciences*, 58(1); 137-137.

[3] Ausubel, L. (2000). "An efficient dynamic auction for heterogeneous commodities," mimeo, University of Maryland.

[4] Ausubel, L., and P. Milgrom (2002). "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics* 1(1).

[5] Babai, L., P. Frankl, and J. Simon. "Complexity Classes in Communication Complexity Theory," *Proceedings of FOCS*, 337–347, 1986.

[6] Banks, J., J. Ledyard, and D. Porter (1989). "Allocating Uncertain and Unresponsive Resources: An Experimental Approach," *Rand Journal of Economics*, 20:1 – 25.

[7] Bikhchandani, S., and J.W. Mamer (1997). "Competitive equilibrium in an exchange economy with indivisibilities," *Journal of Economic Theory*, 74:385 – 413.

[8] Bikhchandani, S., R. Vohra, S. de Vries, and J. Schummer (2001). "Linear programming and Vickrey auctions," mimeo, 2001. Availailabe from http://www.kellogg.nwu.edu/faculty/vohra/htm/res.htm.

[9] Calsamiglia, X. (1977). "Decentralized Resource Allocation and Increasing Returns," *Journal of Economic Theory* 14: 262-283.

[10] Edgar, G.E. (1990). *Measure, Topology, and Fractal Geometry*. New York: Springer-Verlag.

[11] Feigenbaum, J. , C. H. Papadimitriou, and S. Shenker (2000). "Sharing the cost of multicast transactions," STOC 2000.

[12] Gul, F., and E. Stacchetti (1999). "Walrasian equilibrium with gross substitutes," *Journal of Economic Theory*, 87:95 – 124.

[13] Gul, F., and E. Stacchetti (2000). "The English Auction with Differentiated Commodities," *Journal of Economic Theory.*

[14] Hayek, F. (1945). "The Use of Knowledge in Society," *American Economic Review* 35: 519-30.

[15] Holzman, R., N. Kfir-Dahav, D. Monderer, and M. Tennenholtz (2001), "On Bundling Equilibrium in Combinatorial Auctions" mimeo, Technion University, http://iew3.technion.ac.il/~moshet/rndm11.ps

[16] Hurwicz, L. (1960), "Optimality and Informational Efficiency in Resource Allocation Processes," in K.J. Arrow, S. Karlin, and P. Suppes, eds., *Mathematical Methods in the Social Sciences*, 27-46, Stanford: Stanford University Press.

[17] Hurwicz, L. (1977). "On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes," in K.J. Arrow and L. Hurwicz, eds., *Studies in Resource Allocation Processes*, 413-424, New York: Cambridge University Press.

[18] Hurwicz, L., and T. Marschak (2003a), "Finite Allocation Mechanisms: Approximate Walrasian versus Approximate Direct Revelation," *Economic Theory* 21, 545-572.

[19] Hurwicz, L., and T. Marschak (2003b), "Comparing Finite Mechanisms," *Economic Theory* 21, 783-841.

[20] Karloff, H. (1991). *Linear Programming.* Birkhäuser Verlag.

[21] Kelso, A.S. Jr., and V.P. Crawford (1982). "Job matching, coalition formation, and gross substitutes." *Econometrica*, 50:1483-1504.

[22] Korshunov, A.D. (1981). "O Chisle Monotonnykh Bulevykh Funktsi*i", Problemy Kibernetiki 38, 5-108.

[23] Kushilevitz, E., and N. Nisan (1997). *Communication Complexity*. Cambridge University Press.

[24] Kwasnica, A., J. Ledyard, D. Porter, and C. DeMartini (2002). "A new and improved design for multi-object iterative auctions," available from http://www.ices-gmu.net/pdf/materials/372.pdf

[25] Lang, S. (1987). *Linear Algebra*, New-York: Springer-Verlag, 3rd edition.

[26] Lehmann, D., L. Ita O'Callaghan, and Y. Shoham (1999). "Truth revelation in rapid, approximately efficient combinatorial auctions." *1st ACM conference on electronic commerce.*

[27] Lehmann, B., D. Lehmann, and N. Nisan (2001). "Combinatorial Auctions with Decreasing Marginal Utilities." In *ACM conference on electronic commerce.*

[28] Luo, Z.-Q., and J.N. Tsitsiklis (1993). "On the Communication Complexity of Distributed Algebraic Computation." *Journal of the Association for Computer Machinery* 40(5): 1019-1047.

[29] Marschak, J., and R. Radner (1972), *Economic Theory of Teams*, New Haven: Yale University Press.

[30] Marschak, T. (1996) "On Economies of Scope in Communication," *Economic Design* 2(1): 1-30.

[31] Milgrom, P. (2000). "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy.*

[32] Mount, K., and S. Reiter (1974). "The Information Size of Message Spaces," *Journal of Economic Theory* 28:1-18.

[33] Mount, K., and S. Reiter (1977). "Economic Environments for which there are Pareto Satisfactory Mechanisms," *Econometrica* 45:821-842.

[34] Nisan, N. (2000). "Bidding and Allocation in Combinatorial Auctions." In *ACM Conference on Electronic Commerce.*

[35] Nisan, N., and A. Ronen (2000). "Computationally Feasible VCG-Based Mechanisms." In *ACM Conference on Electronic Commerce.*

[36] Parkes, D.C. (1999). "ibundle: An Efficient Ascending Price Bundle Auction," mimeo, Harvard University.

[37] Parkes, D.C. (2002). "Price-Based Information Certificates for Minimal-Revelation Combinatorial Auctions," in *Agent-Mediated Electronic Commerce IV*, Padget et al. (eds), LNAI 2531 pp. 103-122, Springer-Verlag.

[38] Parkes, D.C., and Lyle H. Ungar (2000). Iterative combinatorial auctions: Theory and practice. In *AAAI/IAAI*, pages 74–81.

[39] Reichelstein, S. (1984). "Incentive Compatibility and Informational Requirements," *Journal of Economic Theory* 34(1), 32-51.

[40] Rothkhof, M.H., A. Pekeč, and R.M. Harstad (1998). "Computationally manageable combinatorial auctions." *Management Science*, 44(8):1131–1147.

[41] Sandholm, T., S. Suri, A. Gilpin, and D. Levine "Cabob: A fast optimal algorithm for combinatorial auctions." In *IJCAI*, 2001.

[42] Shannon, C.E. (1948). "A Mathematical Theory of Communication," *Bell System Technical Journal* 27, 379-423, 623-656.

[43] Van Zandt, T. (1999), "Decentralized Information Processing in the Theory of Organizations," in Murat Sertel, ed., *Contemporary Economic Issues, Vol. 4: Economic Design and Behavior*. London: MacMillan Press Ltd., pp. 125-160.

[44] Vazirani, V.V. (2001), *Approximation Algorithms*, Berlin: Springer-Verlag.

[45] Vohra, R., and S. de Vries (2002). "Combinatorial auctions: A survey," *INFORMS Journal of Computing*, forthcoming.

[46] Walker, M. (1977). "On the Informational Size of Message Spaces," *Journal of Economic Theory* 15, 366-375.

[47] Yao, A. C.-C. (1979), "Some complexity questions related to distributive computing," In *ACM Symposium on Theory of Computing*, pp. 209–213.

[48] Zinkevich, M., Blum, A., T. and Sandholm (2003), "Preference Elicitation and Allocation with Read-Once Preferences.," to appear in *2003 ACM conference on Electronic Commerce*.

[49] Zurel, E., and N. Nisan (2001). "An efficient approximate allocation algorithm for combinatorial auctions," *ACM conference on 2001 Electronic Commerce*.
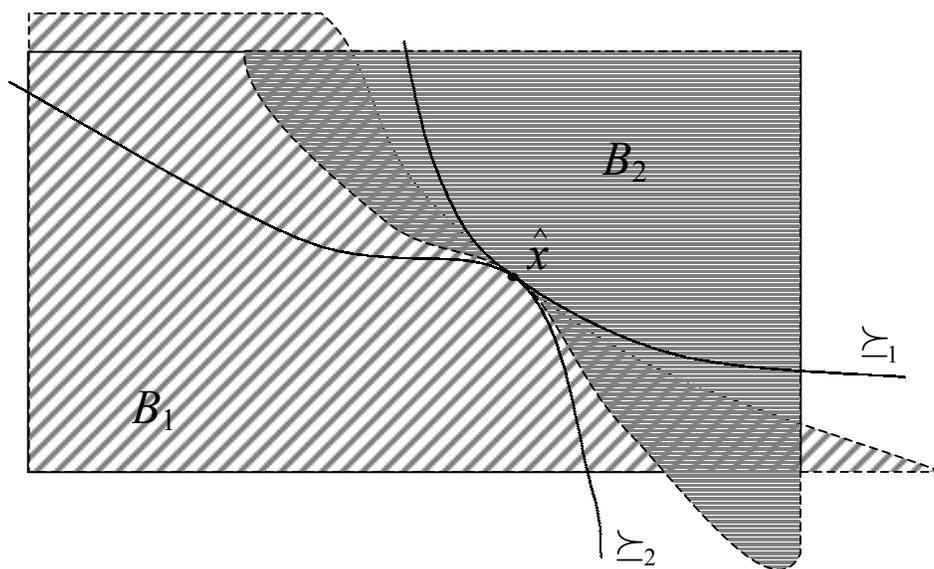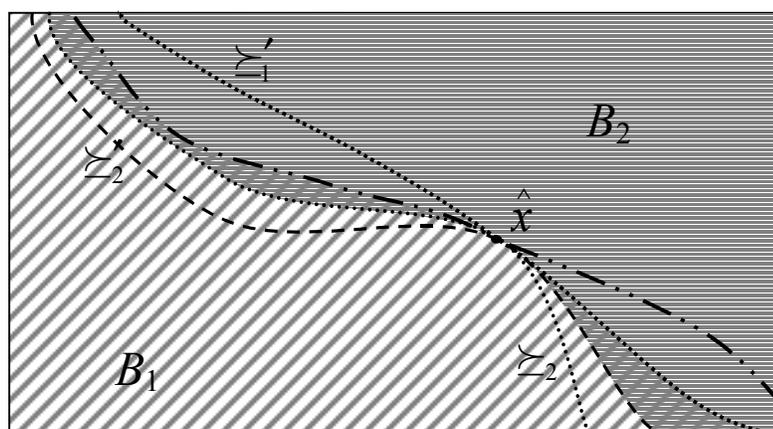
Figure 1: Lindahl Equilibrium

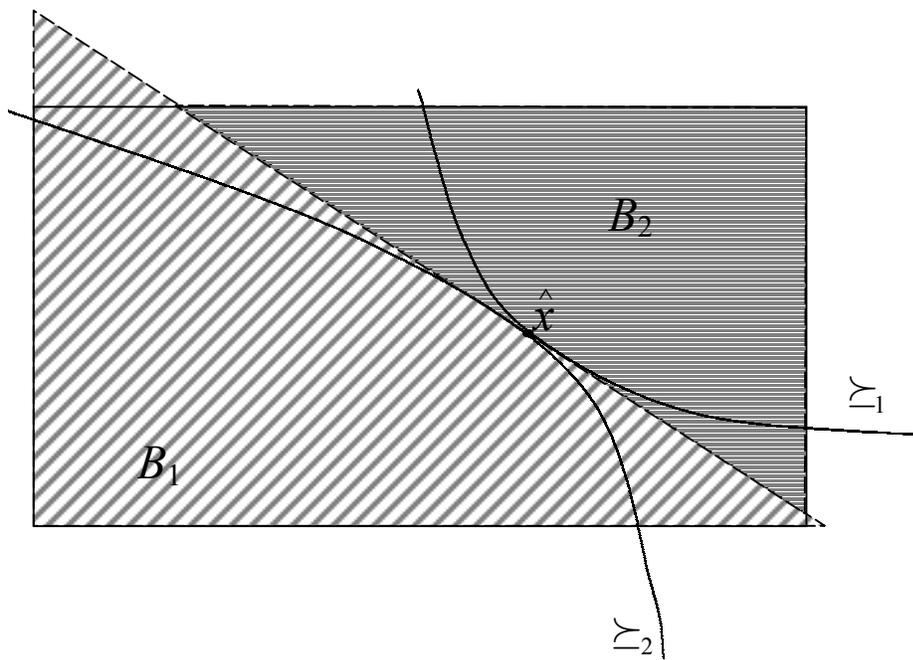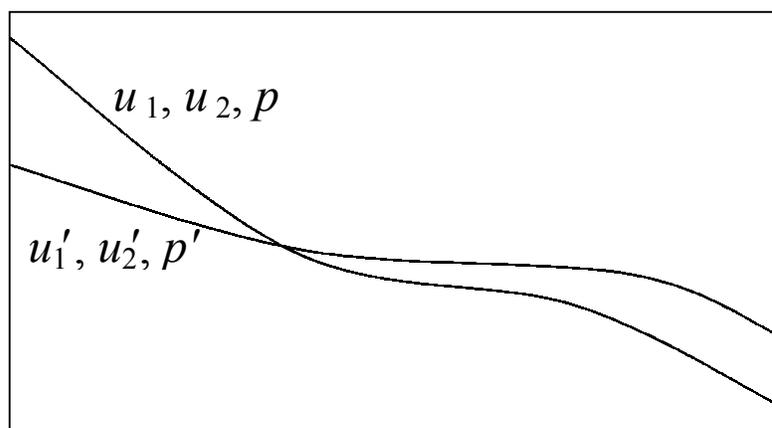Figure 2: Construction of Lindahl Equilibrium from communication

Figure 3: Walrasian Equilibrium

Figure 4: Lindahl Equilibrium with dual preferences