# The Communication Burden of Payment Determination

Moshe Babaioff
Microsoft Research, Silicon Valley
moshe@microsoft.com

Liad Blumrosen
Department of Economics
and
The Center for the Study of Rationality
The Hebrew University
blumrosen@huji.ac.il

Michael Schapira
School of Computer Science and Engineering,
The Hebrew University
schapiram@huji.ac.il

**Abstract**

In the presence of self-interested parties, mechanism designers typically aim to implement some social-choice function in an equilibrium. This paper studies the cost of such equilibrium requirements in terms of communication. While a certain amount of information $x$ needs to be communicated just for computing the outcome of a certain social-choice function, an *additional* amount of communication may be required for computing the equilibrium-supporting payments (if exist).

Our main result shows that the total amount of information required for this task can be greater than $x$ by a factor linear in the number of players $n$, i.e., $n \cdot x$ (under a common normalization assumption). This is the first known lower bound for this problem. In fact, we show that this result holds even in single-parameter domains. On the positive side, we show that certain classic economic domains, namely, single-item auctions and public-good mechanisms, only entail a small overhead.

*keywords:* Implementation, Mechanism Design, Communication Complexity, Revelation Principle.

*JEL codes:* D82, D83.

# 1 Introduction

Consider the goal of designing mechanisms for environments with self-interested players. We seek mechanisms that admit the following two properties: first, *tractability* in the information-theoretic sense, i.e., a low amount of information needs to be communicated in order to realize the outcome of the mechanism. Second, *incentive compatibility*, i.e., the *existence* of some payments that supports the implementation of the outcome of the mechanism in equilibrium. In this work, we show that tractability and the existence of supporting payments are insufficient to establish that implementing the outcome of the mechanism in equilibrium will indeed be practical. This is due to the fact that a non-trivial amount of *additional* communication between the different parties may be required in order to compute equilibrium-supporting payments.

The question of how much overhead one incurs from the computation of incentive-compatible payments was recently raised by Fadel and Segal (2009) (henceforth, FS), who termed this overhead the "*communication cost of selfishness*". In their paper, they studied the communication overhead both for Bayesian equilibria and for ex-post equilibria. In this work we focus on ex-post equilibria - i.e., situations in which players would not want to change their behavior in retrospect, even if they were told (after the fact) everything about the other players. Our main result is that the communication overhead of computing equilibrium-supporting payments may be linear in the number of players. In our paper, we will use the term "Communication Cost of Incentive Compatibility" or when it is clear from the context we will just write *communication cost* (formal definitions will be given later in Section 3).

Our exploration of the communication cost of computing incentive-compatible payments takes place within the *communication-complexity* framework that was developed in the computer-science literature about three decades ago. The basic communication-complexity model is due to Yao (1979); in this 2-agent model, each agent $i \in \{1, 2\}$ privately holds a piece of information $v_i$ and the question is how many bits of information are required in order to determine the outcome of some function $f(v_1, v_2)$ in the worst case. One example, with a straightforward application in auction theory (given below), is where $f(v_1, v_2) = \max\{v_1, v_2\}$. To illustrate our framework and provide an intuition for why the computation of incentive-compatible payments might be costly in terms of information, consider the famous 2-bidder Vickrey auction: there are two bidders, 1 and 2, with private values $v_1$ and $v_2$, respectively, for a single item on sale. Assume that both $v_1$ and $v_2$ are in $\{0, ..., 2^k - 1\}$. The goal is to sell the item to the bidder with the highest value. It is well-known (Vickrey (1961)) that by allocating the item to that bidder, and charging him the second-highest value, one can induce truthful behavior of the participants.

How many bits of information are required to figure out who the bidder with the highest value (the winning bidder) is? The following simple procedure (or *protocol* in the communication-complexity jargon) does exactly that: Player 1 reveals his value to player 2. Player 2 compares the two values, determines the identity of the winning bidder, and informs 1. Observe that any value in $\{0, ..., 2^k - 1\}$ can be represented via $k$ bits (binary representation). Hence, this simple protocol requires $k + 1$ bits (after bidder 1 transmits $k$ bits, bidder 2 need only transmit one more bit indicating the identity of the bidder with the highest value).

However, to achieve incentive-compatibility, finding the bidder with the highest value does not suffice; one must *also* learn the value of the other bidder in order to determine the "second price". Now, the aforementioned simple protocol is no longer the solution, as it could be that the bidder revealing his value is not the one with the lower value. How many *additional* bits must the bidders exchange to learn not only who the winning bidder is, but also what he must pay? It turns out that in the single-item auction example, no much additional information needs to be communicated for determining the relevant payments; in this paper, we would like to explore whether there are environments where this overhead may be a real obstacle.

These kind of questions motivate our study of the communication cost of incentive compatibility. Our main result can roughly be stated as follows:

**Theorem: [Informal]** *There exist social-choice functions such that their outcome can be computed by communicating $x$ bits, but determining both the outcome* and *equilibrium-supporting payments requires about $n \cdot x$ bits of communication, where $n$ is the number of players.*

We prove that this result holds even for very simple domains where the private values of the players are one dimensional ("single-parameter" domains), where in each possible outcome every player either "wins" or "loses". The theorem is proven under the common normalization assumption, which, in the single-parameter domain that we consider, simply means that losing players pay zero. While this assumption does not seem very restrictive at first glance, we currently do not know how to relax it. (The normalization assumption is only required for proving the above impossibility result, and actually strengthens our positive results.) Whether a similar result can be proven without the normalization assumption is left as an intriguing open question, see further discussion in Section 6.

Our lower bound provides the first evidence that the communication cost due to the demand for incentive compatibility may be significant, and it matches the linear upper bound for single-parameter domains in FS. The fact that the communication requirements of mechanisms increases proportionally to the size of the market may have a tremendous effect on the scalability of these mechanisms. This is crucial when applying mechanism-design methods to large-scale Internet-based markets.

Informally, in order to prove our main result we need to construct a social-choice function (SCF) $f$ for which the following requirements hold:[1]

[1]Following a common tradition in computer science, in this paper we refer to a communication cost as "low" if it is within a constant factor of the communication requirement of the original (non-strategic) problem; By a "constant factor" we mean that the ratio is independent of the parameters of the problems

1. $f$ can be implemented in ex-post equilibrium (in single-parameter domains this means that $f$ should be monotone, see Section 3).

2. $f$ can be computed with "low" communication (about the same as the size of the private information of a *single* agent).

3. Computing equilibrium-supporting payments requires "high" communication (about the same as the combined size of the private information of *all* agents).

The difficulty in finding such a social-welfare function is demonstrated by the analysis of two classic economic problems: *public goods* and *single-item auctions*. For these problems, we show that the requirements are *not* met. This enables us to prove positive results for these two problems, by showing that the additional information required to compute the equilibrium-supporting payments is *low* (up to a small constant multiplicative factor). This claim is proven in an inherently different way for each one of these problems.

**Public goods:** Consider a social planner who wants to know whether a bridge should be built or not. A set of $n$ players have utilities $v_1, ..., v_n$ from using the bridge, where $v_i$ is private information of player $i$. To maximize the social welfare the bridge should be built if and only if $\sum_{i=1}^{n} v_i \geq C$ where $C$ is its construction cost. We prove that computing both the outcome and the payments merely requires about six times the communication of computing the outcome alone. Hence, the communication cost in this case is (relatively) small.

**Single-item auction:** In a single-item auction with $n$ buyers (players) each buyer $i$ has private value $v_i$ for the item on sale, and our goal is to sell the item to the player with the highest value. We prove that determining the right allocation and the appropriate payments requires at most three times the communication amount needed to determine the allocation alone.

As these two problems illustrate, coming up with a social-welfare function for which all of our requirements hold is a non-trivial task. We stress that achieving a better lower bound than the linear lower bound shown in this paper may be hard. The communication cost in known to be at most linear (in the number of players) for welfare-maximization objectives and in single-parameter domains (in FS). Thus a better lower bound seems likely to involve the construction of multi-parameter non-welfare maximizing social-choice functions, a class of functions that is relatively little understood (in terms of characterizing incentive-compatible payments).

## 1.1 Related Work

Fadel and Segal (2009) (FS) were the first to study the communication cost of incentive

(for example, the number of agents). A communication cost is considered "high" if it is non-constant, and increases as the original problem becomes more complex.

compatibility. They proved exponential upper bounds, both for Bayesian-Nash equilibria and ex-post equilibria. They presented a surprising matching exponential lower bound for the Bayesian case, but their only lower bound for the ex-post equilibrium case was 1 extra bit. The main open question posed in their paper remains unsolved: can the exponential upper bound for the communication cost in ex-post implementation be matched by a lower bound? FS also proved a linear (in the number of players) upper bound on the communication cost of incentive compatibility in single-parameter domains; Another positive result by FS is an upper bound on the communication cost of incentive compatibility for welfare maximization goals, using an approach by Reichelstein (1984) where agents pay the sum of the reported values of the other agents; Therefore, after realizing the efficient outcome supporting payments can be computed if each agent communicates its value for this outcome.

Another work which is very close to ours is the independent work of Lahaie and Parkes (2008) (henceforth, LP). They considered the amount of information overhead of computing Vickrey-Clarke-Groves (VCG) payments in socially-efficient mechanisms. LP showed settings with multi-dimensional types, in which the naïve VCG protocol that computes several efficient allocations in sequence (one execution with all $n$ agents, and $n$ executions with $n-1$ agents, where one of the agents is removed each time) is asymptotically optimal. This derives a result along the same line as our main impossibility result, that the communication cost of realizing equilibrium-supporting payments might be linear in the number of players. LP proved their result using a neat characterization of the communication requirement of the VCG outcome via the concept of *universal competitive equilibrium*.

Our results, and the result by LP are complementary yet incomparable, in the sense that our results do not imply the results of LP, and vice versa, as we now explain. The work of LP is different from ours in two important respects: the kind of environments considered and the kind of mechanisms considered. When proving negative results, the more restrictive the environment, and the class of mechanisms considered, the stronger the result. We point out that, both in terms of the environments considered, and of the class of mechanisms considered, our results do not apply to more restricted cases than those of LP, and vice versa. More specifically, LP study multi-parameter domains in which the objective is maximizing social welfare. We, in contrast, consider one-dimensional environments (single parameter domains), but do not restrict our attention to social-welfare maximization. Thus, our work and the work of LP consider two different restrictions on environments; while we remove from consideration complex (multi-dimensional) players' types, LP eliminate the possibility of social choice functions that do not optimize social welfare. In addition, the work of LP and ours consider two different restrictions on mechanisms; while LP focus on the popular VCG mechanisms, which charge every player the externality he imposes on the other players, in this work we focus on normalized incentive-compatible mechanisms. Indeed, LP do not claim that their results hold for all normalized incentive-compatible mechanisms. (In fact, this is not true; since for the class of valuations constructed in Section 6.2.2 of LP the efficient allocation is known before the bidding procedure starts, then a trivial mechanism that announces the efficient allocation and charges zero payments from all players is normalized and incentive compatible.) We do not know whether the lower bound by LP can be modified to hold for all incentive-compatible normalized mechanisms, and this remains an interesting open question (see Section 7).

Our linear lower bound for single-parameter domains holds for *deterministic* computation of payments that induce incentive computability. Babaioff et al. (2010) consider the case that one allows for randomized mechanisms that are only incentive compatible *in expectation*, and is also willing to allow two-directional payments (from the mechanism to the agents and vice versa). Babaioff et al. (2010) establishes that, for any social choice function, it is possible to implement a social choice function that outputs the same outcome with high probability (arbitrary close to 1), with no extra communication cost. Following the work of Babaioff et al. (2010), Segal[2] has observed that when considering domains where no information is revealed over time, as is the case with all the domains considered in this paper, one can take any single-parameter SCF and implement it (exactly) by a randomized, incentive compatible in expectation mechanism, with only a low communication cost. These results show that our negative results in this paper cannot be extended to randomized mechanisms that are only incentive compatible *in expectation*.

Both the work by FS and our paper belong to a more general line of research studying communication and information aspects of various economic environments, for example in auctions (Nisan and Segal (2006); Blumrosen and Nisan (2009, 2010); Segal (2010)) and in other economic domains (Blumrosen and Feldman (2006); Segal (2007); Hart and Mansour (2010)). A recent survey on this line of research in the context of combinatorial auctions is by Segal (2006). A recent paper by Papadimitriou et al. (2008) (see also Dobzinski and Nisan (2007)) presented settings where incentive-compatible mechanisms that are approximately efficient must entail exponential communication, while without the incentive-compatibility requirement such mechanisms only require polynomial communication.

The basic model for communication complexity was presented by Yao (1979), and a survey on communication complexity was given by Kushilevitz and Nisan (1997). Some early and influential work in the economics literature presented models in the same spirit as Yao's model, where the main difference is that these economic model allowed the communication of real numbers. This line of work goes back to the seminal work by Hurwicz (1960), and it was later shown that Walrasian mechanisms use message spaces of the lowest possible dimensions among the Pareto efficient mechanisms (see the work by Mount and Reiter (1974); Hurwicz (1977) and the survey in Hurwicz and Reiter (2006)).

Another related line of research studied the *computational* consequences of the need to determine equilibrium-supporting payments. The question of the computational burden of computing payments in social-welfare maximization environments was raised in the seminal paper by Nisan and Ronen (2001). This question has received attention since (Hershberger

[2]Private communication, 2010.

and Suri (2001); Emek et al. (2008); Bikhchandani et al. (2002)). Unlike our work, and that of FS and LP, these works considered the *computational* complexity, and not the *informational cost* of realizing the appropriate payments. Several works studied a different related question, of how to maximize the objective of the planner subject to being restricted to a fixed amount of communication that the mechanism can use. These papers studied either profit or efficiency maximization, see the recent survey by Mookherjee (2006), see also , e.g., Green and Laffont (1987); Melumad et al. (1992); Blumrosen et al. (2006).

### 1.2 Organization of the Paper

The rest of the paper is organized as follows: In Section 2 we give an informal discussion of ways for proving impossibility results regarding on the communication cost of incentive-compatibility, and demonstrate these ideas through a simple 2-player public good setting. We present our model and notations in Section 3. In Section 4, we prove constant upper bounds on the communication cost of incentive compatibility for the classic models of single-item auctions and public goods among $n$ players. Our main impossibility result is a linear lower bound which is presented in Section 5. Finally, in Section 6 we discuss the issue of extending our results to the non-normalized case, and in Section 7 we present open questions and directions for future research.

## 2 Informal Example: 2-Player Public-Good

Before making our framework and results precise, let us consider an informal example which would enable us to convey the essence of our techniques. Consider the classic "public good" problem, which deals with the construction of a public project. Each player $i$ in a set of players has a "benefit" of $v_i$ from using the public good, and the social planner aims to implement the efficient outcome and build it if and only if the sum of benefits is at least the fixed construction cost $C$. We restrict our attention to the 2-player case, where $v_1$ and $v_2$ are integers in $\{0, ..., 2^k - 1\}$ (for some integer $k > 0$), and $C = 2^k$. This SCF is described in Figure 1.[3]

**Input:** values $v_1, v_2 \in \{0, ..., 2^k - 1\}$.

**Output:** "Build" if $v_1 + v_2 \geq 2^k$, "Do Not Build" otherwise.

**Computing the outcome alone can be done with $k + 1$ bits.** How many bits are required to determine the outcome? Note that we are interested in discovering the number of bits that will be sufficient for this task for all possible inputs (i.e., pairs of player values).

[3]Note that we are interested in finding *some* payments that support an incentive-compatible mechanism. The question whether there exists such payments with good properties (e.g., that cover the construction cost) is secondary to our main goal.

As in our introductory example, the following simple protocol is sufficient: player 1 reveals his information ($v_1$) to player 2 (via $k$ bits). Player 2 computes $v_1 + v_2$, checks whether it is at least $C$, and informs player 1 (via one additional bit). Therefore, an exchange of $k+1$ bits is sufficient for both players to compute the outcome of the mechanism for every type profile.

**Computing the outcome and incentive-compatible payments requires at least $2k-1$ bits.** Observe that each player has $2^k$ possible values, and thus the two players have $2^k \cdot 2^k = 2^{2k}$ possible pairs of values $v_1, v_2$. Also observe that for more than half of these pairs, i.e., more than $2^{2k-1}$ pairs, the outcome is "Build". (These are the types profile drawn on and under the diagonal in Figure 1.) We shall denote the set of all pairs of values for which the outcome is "Build" by $\mathcal{B}$. An important fact about incentive-compatible payments is that, under the normalization (losers pay zero)assumption, then, for every $v_1, v_2$ in $\mathcal{B}$, the (only) incentive-compatible payments must be $C - v_2$ and $C - v_1$, respectively. This has a very intuitive interpretation: Each player is charged the minimal value he had to announce in order to reach a "Build" outcome.[4] This value is simply the cost of the public good ($C$) minus the value of the other player.

Therefore, for every two different pairs of values in $\mathcal{B}$ any incentive-compatible mechanism must output two different pairs of payments. This is shown in Figure 1, where the equilibrium payments are actually determined by a projection on the main diagonal (the line where the designer is indifferent between building the bridge or not). This means that any incentive-compatible mechanism must have at least $|\mathcal{B}|$ distinct outcomes. It is intuitive and well known (Kushilevitz and Nisan (1997)), that no communication protocol can output $|\mathcal{B}|$ distinct outcomes without requiring the communication of at least $\log |\mathcal{B}|$ bits. Hence, any incentive-compatible mechanism requires the transmission of at least $\log |\mathcal{B}|$ bits. As $|\mathcal{B}| \geq 2^{2k-1}$, we get that $\log |\mathcal{B}| \geq 2k - 1$.

**The communication cost of incentive compatibility.** We now know that, for any value of $k$, the outcome alone can be computed by transmitting at most $k+1$ bits, and, in contrast, computing the outcome and payments that support this outcome in equilibrium requires at least $2k-1$ bits. That is, the computation of the outcome and matching incentive-compatible payments is almost twice as costly as computing the outcome alone (observe that the ratio between computing the outcome and payments, and computing the outcome alone, is arbitrarily close to 2 as $k$ grows). Note that, in the 2-player case, one can always compute the outcome and payment via $2k$ bits by having each player transmit all of his private information ($k$ bits) to the other player. Hence, our result regarding the communication cost of incentive compatibility is essentially tight for the 2-player case.

[4] These equilibrium-supporting payments are unique under the normalization assumption, i.e., that the players pay 0 when losing. We will require this assumption also for our more general impossibility result. The solution in this public-good setting without the normalization assumption is discussed towards the end of this paper in Section 6.

Figure 1: (**2-player public good.**) The figure draws the 2-player public good social choice function. The row value represents the value of player A, and the column value represent the value of player B. Both values are between $[0, 1]$, and the public good should be built if the sum of these values is at least the construction cost of 1. The unique (normalized) payments that support an incentive compatible implementation of this function are determined by projection of the preference point onto the diagonal; for example, given the preference profile $x$, player $A$ should pay $p_A$ and player $B$ should pay $p_B$. It is easy that the payments are distinct in every valuation profile in which the public good is built.

**The $n$-player case.** Our goal in this paper is to show that, as in the 2-player case, the communication cost of computing payments might sometimes be linear in the number of players. That is, we wish to show that there exist social-choice functions for which computing the outcome and payments is roughly $n$ times more expensive, in terms of communication, than computing the outcome alone (where $n$ is the number of players). As we shall later see, the public good setting will no longer be of use to us in proving such a result and different social-choice functions will have to be considered.

# 3 Background and Model

## 3.1 Mechanism Design

The mechanism design setting considered in this paper is as follows: there is a set $N$ of $n$ players, and a set of outcomes $O$. In our paper, $O$ denotes all possible sets of winning players, that is, $O = 2^N$. Each player $i$ has a *valuation function*, or *type*, $v_i : O \to \mathbb{R}_{\geq 0}$, that belongs to a set of valuation functions $V$. A *social-choice function* (SCF) is a function that assigns every n-tuple of players' valuation functions $v = (v_1, ..., v_n) \in V^n$ ("*type-profile*") an outcome $o \in O$. Each $v_i$ is private and only known to $i$.

A *payment function* is a function $p : V^n \to \mathbb{R}^n$.

**Definition 1.** *A social-choice function $f$ is said to be* implementable *(in the ex-post Nash sense) if there is a payment function $p$ such that the following holds:*

$$\forall v = (v_1, ..., v_n) \in V^n,\ \forall i \in N,\ \forall v_i' \in V,$$

$$v_i(f(v)) - p(v) \geq v_i(f(v_i', v_{-i})) - p(v_i', v_{-i})$$

*(where $(v_i', v_{-i})$ is the type profile in which $i$ has type $v_i'$ and every player $j \neq i$ has type $v_j$)*

Informally, $f$ is implementable if it is possible to come up with a payment scheme that incentivizes players to report truthful information. For example, the social-choice function in single-item auctions (where the item should be sold to the player with the highest value) is $f(v_1, ..., v_n) \in \arg max_{i \in N} v_i$, where $v_i$ is the value of agent $i$ for the item. The payment scheme in a *second-price* (Vickrey) auction is known to implement this social-choice function in equilibrium.

In this paper, we provide several examples of single-parameter domains, where the type of a player can be represented by a single scalar. We consider specific single-parameter environments where every outcome defines whether each player "wins" or "loses" (i.e., $f(v) \subseteq N$ is the set of winners). A player gains a value of $v_i \geq 0$ if she wins, and she gains 0 when losing. We focus on normalized mechanisms in which losers pay 0.

**Definition 2.** ***(Normalization Assumption.)*** *In a single-parameter environment, a mechanism that implements the social-choice function $f$ satisfies the normalization (or losers-pay-zero) assumption if for every player $i$ and type profile $v$ such that $i \notin f(v)$ we have that $i$ pays zero, i.e., $p_i(v) = 0$ (where $f(v) \subseteq N$ is the set of chosen winners).*

The characterization of implementable single-parameter social-choice functions relies on the following property.

**Definition 3.** *A single-parameter SCF $f$ is* monotone *if for every player $i \in N$, all $v_{-i} \in V_{-i}$ and all $v_i' > v_i$ s.t. $v_i', v_i \in V$ it holds that if $i \in f(v_i, v_{-i})$ then $i \in f(v_i', v_{-i})$.*

The following characterization is well known (see, e.g., Mirrlees (1971), Myerson (1981), Mookherjee and Reichelstein (1992)).

**Observation 3.1.** *A single parameter SCF $f$ is implementable if and only if it is monotone. Under the normalization assumption, a winner has to pay the minimal bid she has to declare in order to win.*[5]

We note that in our proofs we use the Normalization assumption from Definition 2 only via the use of Observation 3.1.

[5]The payment is unique in the case that bids come from the continuous domain of non-negative real numbers. In our discrete bids setting, the payment can actually lie within an interval between two consecutive discrete bids which contains the minimal bid the winner has to declare in order to win. Yet, any such payment can be uniquely mapped to a payment taken from the domain of discrete bids, without any additional communication. Thus, our main impossibility result shows that computing any of these payments requires communicating many bits.

### 3.2 Communication cost of Incentive Compatibility

We consider the communication problem in which each $v_i$ is private and only known to $i$, and the players need to exchange information in order to compute the outcome of $f$.[6] We work in the broadcast model in which each sent bit is received by all players (and not addressed only to one player). Let $CC(f)$ denote the communication complexity of computing the outcome of $f$. Informally, the communication complexity of a function is the minimal number of bits that is required to compute the function (for any input). The communication complexity of a function is a worst-case measure, describing the minimal amount of information that can guarantee determining the outcome of $f$ on every input. This model was first suggested by Yao (1979), and is the standard communication model in computer-science theory which recently had several applications in economics (see, for example, Nisan and Segal (2006); Segal (2007); Hart and Mansour (2010); Fadel and Segal (2009)). For readers that are not familiar with the concept, we give the formal definition of a communication protocol and of communication complexity in Appendix A

How much additional communication burden is imposed by the necessity to compute payments that guarantee truthfulness? Let $CC_{IC}(f)$ denote the communication complexity of computing the outcome of $f$ and payments that guarantee incentive compatibility (that is, computing the outcome of both $f$ and *some* payment function that leads to the implementability of $f$.). Formally, given a social-choice function $f$ let $f_p : V^n \to O \times \mathbb{R}^n$ denote a function with two outputs, an outcome and a payment vector, such that the output selection in $f_p$ is identical to $f$. We say that $f_p$ is incentive compatible (IC) if $f$ is implementable via the payment function in $f_p$ as in Definition 1. We say that $f_p$ is normalized if it satisfies the normalization assumption. Then,

$$CC_{IC}(f) = \min_{f_p \text{ is IC and normalized}} \{CC(f_p)\} \tag{1}$$

We note that FS did not require the mechanisms to be normalized. We refer the reader to Section 6 for a discussion on the normalization assumption in our setting.

In order to formally define the communication cost of incentive compatibility we need to be concrete about the information each player holds: Let $f_k$ be a social-choice function with $n$ players such that each player's valuation is represented using $k$ bits of information, for some fixed $k \in \mathbb{N}$ (although $f_k$ depends on $n$, as the number of agents $n$ will always be clear from the context we simplify the notation and do not explicitly express that dependency in our notation). Formally, $f_k : \{0,1\}^{k \times n} \to O$, i.e., for every $(v_1, v_2, ..., v_n)$ with each $v_i \in \{0,1\}^k$, $f_k$ picks an outcome $f_k(v_1, v_2, ..., v_n)$.

**Definition 4.** *The* communication cost of incentive compatibility *of $f_k$ is defined to be* $\frac{CC_{IC}(f_k)}{CC(f_k)}$.

[6] We shall regard the exchange of information as taking place between the players themselves, without another entity that represents the social planner or the auctioneer. This strengthen our main result which is a lower bound. It is also easy to verify that our constant upper bounds for $n$ players single item auction and public good can be modified to constant upper bounds even with the additional entity.

Our main result shows that for some social-choice functions $f_k$ a significant communication overhead may be incurred, and we prove that this holds even in single-parameter domains. Formally, in such domains the valuation of a player is given by a number in $[0, 1]$ represented by $k$ bits ($k$ is the precision in the representation of $v_i$). That is, for every player $i$ there is some $t_i \in \{0, ..., 2^k - 1\}$, such that $v_i = t_i \cdot 2^{-k}$.

One alternative to Def. 4 is to define the communication cost of selfishness as $\lim_{k\to\infty} \frac{CC_{IC}(f_k)}{CC(f_k)}$. This would make sense as most of our results hold given any small $\epsilon$ and for large enough $k$'s, and this alternative definition would eliminate the need for the $\epsilon$ qualification. Definition 4, however, enables us to show the tradeoff between $k$ and $\epsilon$ in a more explicit way, while the results in the limit are can be immediately deduced.

### 3.3 Communication Complexity: Background and Basic Observations

This section presents some basic background of some of the tools we use from the theory of communication complexity. For a comprehensive survey on the subject we refer the reader to the book by Kushilevitz and Nisan (1997).

We will start by presenting two easy observations that essentially show that when one needs to distinguish between many different outcomes, the communication complexity cannot be very small.

**Observation 3.2.** *If the range of function $f$ is of size $m$ ($|f(V)| = m$), then any communication protocol for $f$ requires at least* $\log(m)$ *bits. ($f(V)$ is the set of different outcomes in the range of $f$.)*

**Observation 3.3.** *For any implementable function $f_k$ with $n$ players each holding $k$ bits it holds that $CC(f_k) \leq k(n-1) + \lceil \log(|f_k(V)|) \rceil$ and $CC_{IC}(f_k) \leq kn$ .*

*Proof.* This is shown by considering two trivial protocols: To compute $f_k$, each of the players but one transmits all his information, and the last player computes the outcome and transmits this outcome. As there are $|f_k(V)|$ possible relevant outcomes, $\lceil \log(|f_k(V)|) \rceil$ bits are clearly sufficient to encode all these outcomes. To compute $CC_{IC}(f_k)$ simply let all players transmit all of their private information. □

Our proofs use a common communication-complexity technique called "*fooling sets*". Intuitively, a fooling set is a large set of possible inputs such that any communication protocol must have a different execution (exchange of bits) for every two of them. Fooling set arguments are based on a well known property of communication protocols (Kushilevitz and Nisan (1997)): if a protocol execution is exactly the same on two inputs, it must output the exact same outcome for any possible "combination" of these inputs. That is, suppose that some communication protocol that computes a social function $f$ for a two-player setting executes exactly the same for $(v_1, v_2)$ and $(v'_1, v'_2)$ (thus outputs the same outcome on both), then, that protocol would execute exactly the same for the inputs $(v'_1, v_2)$ and $(v_1, v'_2)$, and thus must also output the same outcome for these inputs.

This suggests a way for proving lower bounds on the communication complexity of social-choice functions: Find a subset $X$ of the inputs, where parts of each input are held by the players as their types; Then, prove that even though the social choice function assigns the same outcome for every input in this subset, a combination (in the above sense) of *every* two members is assigned a different outcome. This will imply that *any* communication protocol *must* distinguish between *every* two members of the subset of inputs. This, in turn, would mean that $\log |X|$ is a lower bound on the number of bits that need to be transmitted in order to compute $f$.

Formally, let $f$ be a social-choice function. Let $v, v'$ be two valuation functions. Let

$$V_{v,v'} = \{v'' = (v''_1, ..., v''_n) | \forall i \in N \; v''_i = v_i \text{ or } v''_i = v'_i\}$$

A well known fact in communication complexity (see Kushilevitz and Nisan (1997)) is the following:

**Theorem 3.4.** **[Fooling Set Argument]** *Let $f : V = V_1 \times \ldots \times V_n \to O$ be a social-choice function. For every $V' \subseteq V$ such that:*

- *there is an outcome $o^* \in O$ such that for every $v' \in V'$ $f(v') = o^*$.*
- *for every $v, v' \in V'$ $\exists v'' \in V_{v,v'}$ such that $f(v'') \neq o^*$*

*it holds that $CC(f) \geq log(|V'|)$.*

# 4 Environments with a Small Overhead

We start by discussing the communication cost required for computing the equilibrium-supporting payments in two classic environments: single-item auctions and public goods. It turns out that in both environments this additional information is low – at most a constant factor of 3 for the first and at most 6 for the latter. That is, even with numerous bidders, determining the relevant payments requires at most 6 times the information that is needed to determine the outcome alone in these examples.

## 4.1 Single-Item Auctions

In a socially-efficient single-item auction, a seller aims to sell an item to the bidder who values it the most. We now formally define the corresponding social-choice function Single-Item-Auction.

**Definition 5** (Single-Item-Auction)**.**
*Input: valuations $v_1, ..., v_n \in \mathbb{N}$*
*Output: a bidder with the highest value, i.e.,* $\arg max_{i \in N} v_i$ *(breaking ties lexicographically).*

If bidder $i$ wins, his value for the outcome is $v_i$; otherwise, his value is 0. From SINGLE-ITEM-AUCTION we can derive the function SINGLE-ITEM-AUCTION$_k$ for the case where the values of the players are numbers between $[0, 1]$ represented by $k$ bits of information, that is, $v_i = t_i \cdot 2^{-k}$ for some integer $t_i \in \{0, ..., 2^k - 1\}$. It is well known that this function is implementable: if winners pay the second-highest bid then the auction is (dominant-strategy) truthful.

Next, we show that for SINGLE-ITEM-AUCTION the communication cost of incentive compatibility is at most a small constant (i.e., 3).

**Proposition 4.1.** *The communication cost of incentive compatibility for the social-choice function* SINGLE-ITEM-AUCTION *is at most 3.*

*Proof.* To prove the claim we show that for every $k$, it holds for the social-choice function $f_k$ =SINGLE-ITEM-AUCTION$_k$ that

$$CC_{IC}(f_k) \leq 2 \cdot CC(f_k) + k \leq 3 \cdot CC(f_k) \tag{2}$$

To show that $CC_{IC}(f_k) \leq 2CC(f_k) + k$ we present a simple protocol for $CC_{IC}(f_k)$: let $\mathcal{P}$ be the optimal protocol for computing $f_k$; by definition, $\mathcal{P}$ communicates at most $CC(f_k)$ bits when running on any input. Now, we first run $\mathcal{P}$ to find the player with highest value. We then remove the highest bidder and run the protocol $\mathcal{P}$ for $f_k$ again (or more formally, we replace this player with a bidder with a zero valuation so we can use the same protocol as a "black box"). Now the players know who is the player with the second highest value. Finally, this player transmits his value (the payment), which requires $k$ more bits. Overall, the new protocol runs two executions of $\mathcal{P}$ plus at most $k$ bits, a total of $2CC(f_k) + k$.

The last weak inequality in Equation 2 is a result of the following claim. (We prove this claim for completeness. Similar claims were proven before, e.g., by Nisan and Segal (2006).)

**Claim 1.** *For every $n \geq 2$, $CC$(*SINGLE-ITEM-AUCTION$_k$*) $\geq k$*

*Proof.* We shall prove that $CC(f_k)$ is large by constructing a "fooling set" of size $2^k$. By Theorem 3.4 this shows that $CC(f_k) \geq k$. Consider all pairs $(v_1, v_2)$ such that $v_1 = v_2$. No two such pairs can have exactly the same execution of the communication protocol: assume in contradiction that $(v_1, v_2)$ and $(v_1', v_2')$ are two distinct type-profiles that have the same execution. Observe that for all these type-profiles bidder 1 wins and bidder 2 loses. W.l.o.g., let $v_1 < v_1'$. Then, $(v_1, v_2')$ should also have the same execution as $(v_1, v_2)$ and $(v_1', v_2')$. However, this leads to a contradiction because in this case player 2 should win and not player 1. Hence, there are at least $2^k$ pairs such that each two must have a different execution by any communication protocol, and so any protocol that computes $f$ must transmit at least $k$ bits. □

This concludes the proof of Proposition 4.1 □

## 4.2 The Public-Good Setting

We now consider another classic economic setting - the construction of a public project ("public good"). Each player in a set of players has a "benefit" of $v_i$ from using the public good, and the social planner aims to build it only if the sum of benefits is at least the construction cost $C$. The function is defined given the parameter $C \geq 0$.

**Definition 6** ($C$-Public-Good)**.**
*Input: valuations $v_1, ..., v_n \in \mathbb{N}$.*
*Output: "Build" if $\sum_{i=1}^{n} v_i \geq C$; "Do not build", otherwise.*

It is easy to observe that the payments that implement this SCF in a normalized mechanism are $p_i = C - \sum_{j \neq i} v_j$ in the case of "Build" (and all players win). Again, we consider the derived SCF $C$-Public-Good$_k$ for the case where the types of the players are in numbers in $[0, 1]$ represented by $k$-bit strings, i.e., $v_i = t_i \cdot 2^{-k}$ for some integer $t_i \in \{0, ..., 2^k - 1\}$.

In Section 2, we showed that for the public good problem with 2 players, the communication cost of incentive compatibility is essentially 2. Yet, it turns out that when moving to an $n$-player setting the overhead remains constant and does not grow with $n$. Proving this upper bound requires more work than the previous auction example. The main intuitive argument here is that solving the allocation problem alone (whether to build or not) already requires all the players to reveal almost their entire private information, so the overhead incurred by realizing the desired payments must be limited.

**Theorem 4.2.** *Fix $\epsilon > 0$. For any $C$, the communication cost of incentive compatibility of the social-choice function $C$-Public-Good$_k$ with $n \geq 3$ agents and $k$ that is large enough is at most $2 \cdot \frac{n}{n-2} + \epsilon \leq 6 + \epsilon$.*

*Proof.* The players values $v_1, ..., v_n$ are all in $[0, 1]$ and we are interested in figuring out whether $\Sigma_i v_i \geq C$. Obviously, when $C = 0$ or $C \geq n$ then the answer is trivial even without any communication. So, we can consider the case $C \in (0, n)$. For convenience, we will multiply all inputs (player values $v_i$ and cost $C$) by $2^k$, and this clearly does not change the complexity of the problem. That is, we will consider agents with values $\tilde{v_i} = v_i \cdot 2^k \in \{0, ..., 2^k - 1\}$ and the cost $\tilde{C} = C \cdot 2^k \in (0, n2^k)$. We denote by $k_c$ the integer such that $\tilde{C} \in (2^{k_c - 1}, 2^{k_c}]$. Note that $k_c$ can be made arbitrarily large by increasing $k$. In the rest of the proof we will abuse notations and refer to these normalized values by $C$ and $v_i$'s.

**Lemma 4.3.** $CC_{IC}(f_k) \leq nk_c + n$

*Proof.* The following protocol elicits enough information for computing the outcome of $f_k$ and the desired payments with $nk_c + n$ bits: Ask each player if his value is at least $C$ (this requires $n$ bits). Ask all the players whose values are lower than $C$ to transmit their $v_i$'s (this requires $k_c$ bits per player, i.e., at most $n \times k_c$ bits). An easy observation is that this simple protocol provides us with sufficient information to calculate the payments for all players. $\square$

We shall now prove a lower bound on $CC(f_k)$ that will imply the theorem. We prove the following lemmas:

**Lemma 4.4.** *If $C \leq \frac{n}{2}$, then for large enough $k_c$ the communication complexity of determining whether $\Sigma_i v_i \geq C$ or whether $\Sigma_i v_i \leq C$ is at least $(\frac{n}{2} - 1) \cdot k_c - \gamma_n$, where $\gamma_n$ is a function of $n$ which is independent of $k_c, k$.*

*Proof.* We shall construct a large fooling set and invoke Theorem 3.4. Consider all the possible type-profiles $(v_1, \ldots, v_n)$ such that $\Sigma_i v_i = C$. Obviously for all these type-profiles $\Sigma_i v_i \geq C$. However, any protocol that tries to determine whether $\Sigma_i v_i \geq C$ cannot avoid making a distinction between any two such type-profiles for the following reason: Let $v = (v_1, \ldots, v_n)$ and $v' = (v'_1, \ldots, v'_n)$ be two different such type-profiles. Let $j$ be a coordinate such that $v_j \neq v'_j$. W.l.o.g, assume that $v_j < v'_j$. Then, if $v$ and $v'$ are undistinguished then so is the type profile $v'' = (v''_1, \ldots, v''_n)$ in which $v''_i = v'_i$ for every $i \neq j$ and $v''_j = v_j$. However, this leads to a contradiction because $\Sigma_i v''_i < C$ (since the outcome of the protocol cannot be the same for $v$ and $v''$).

Hence, by finding a lower bound $L$ on the number of possible type-profiles $(v_1, \ldots, v_n)$ such that $\Sigma_i v_i = C$, we also find a lower bound on the number of bits transmitted by any protocol that computes $C$-Public-Good$_k$. Specifically, $\log L$ is a lower bound on the number of bits transmitted by any such protocol (see Observation 3.2). In Lemma B.1 in the appendix, we show that $\log L \geq (\frac{n}{2} - 1)k_c - n \cdot \log n$ (here we use the assumption that $C \leq \frac{n}{2}$) which completes the proof.

The complexity of determining whether $\Sigma_i v_i \leq C$ is proved analogously. $\square$

Lemma 4.4 allows us to bound $CC(f_k)$ from below when $C \leq \frac{n}{2}$. That is, we have that when $C \leq \frac{n}{2}$, $CC(f_k) \geq (\frac{n}{2} - 1) \cdot k_c - \gamma_n$ for some $\gamma_n$ independent of $k_c$. Now consider the case that $C > \frac{n}{2}$. Clearly, finding whether $\Sigma_i v_i \geq C$ is equivalent to figuring out whether $n - \Sigma_i v_i \leq n - C$. Thus, our problem now is equivalent to the problem in which every player has the value $1 - v_i$ and the players are trying to figure out whether the sum of their values is at most $n - C$ (note that $n - C \leq \frac{n}{2}$ as in this case $C > \frac{n}{2}$). This problem is just as hard as the original problem, as shown by the second part of Lemma 4.4.

Finally, the theorem follows by plugging in the lower bound on $CC(f_k)$ and the upper bound on $CC_{IC}(f_k)$ (from Lemma 4.3), showing that their ratio is at most

$$\frac{nk_c + n}{(\frac{n}{2} - 1) \cdot k_c - \gamma_n} \tag{3}$$

which converges (when $n \geq 3$) to $2\frac{n}{n-2}$ as $k_c$ goes to infinity. $\square$

# 5 Main Result: A Linear Lower Bound

In order to prove a lower bound of $n$ we must identify a social-choice function $f$ such that $CC(f)$ is smaller than $CC_{IC}(f)$ by a factor of $n$. In our proof, we will construct an ad-hoc $n$-player social-choice function in which the communication cost of incentive compatibility is large. We are not aware of a natural economic interpretation for this social-choice function, and its main goal is for proving the existence of functions with large overhead. We leave the question whether there are "natural" problems with large information overhead open.

As a warm-up, we start by presenting our construction for 2 players.

## 5.1 Another Lower Bound for 2 Players

Recall that we have seen that for the public good setting with 2 players the communication cost of incentive compatibility is essentially 2, but the lower bound used in that proof does not extend to a linear lower bound for the $n$ players case. In fact, the $n$-player public-good problem is such that the communication cost of incentive compatibility is never greater than $6 + \epsilon$.

We will now present a construction that does extend to $n$ players, and thus obtains a linear lower bound. We will start by presenting the construction and the proof for 2 players, and then we will describe the general construction and proof.

Consider the 2-player social-choice function depicted in Figure 2. As before, we have 2 players now denoted by $A, B$. Each player holds a value between $[0, 1]$ represented by a $k$-bit string, i.e., $v_i = t_i \cdot 2^{-k}$ for an integer $t_i \in \{0, ..., 2^k - 1\}$. Player $i$'s utility from winning is $v_i$, and his utility from losing is 0. Player $A$ wins if and only if $v_B \geq 1/2$ and $v_A \geq v_B - 1/2$. Similarly, player $B$ wins if and only if $v_A \geq 1/2$ and $v_B \geq v_A - 1/2$. We shall refer to $f_k$ as the "NOT-TOO-FAR$_k$" social-choice function. It is easy to check that NOT-TOO-FAR$_k$ is monotone and thus it is implementable.

**Proposition 5.1.** *Assume $n = 2$. For any $\epsilon > 0$ and for $k$ large enough, the communication cost of incentive compatibility for the social-choice function* NOT-TOO-FAR$_k$ *is at least* $2 - \epsilon$.

*Proof.* Let $f_k =$NOT-TOO-FAR$_k$. By Observation 3.3 it holds that $CC(f_k) \leq k+2$ as there are 4 outcomes in the range (any subset of the player can win).

We show below that $CC_{IC}(f_k) \geq 2k - 2$. From this we derive that $\frac{CC_{IC}(f_k)}{CC(f_k)} \geq \frac{2k-2}{k+2} = 2 - \frac{6}{k+2}$. Clearly, as this is a monotone function of $k$ that converges to 2, for any $\epsilon > 0$ there is a $k$ such that it is larger than $2 - \epsilon$. We next derive the promised lower bound on $CC_{IC}(f_k)$.

**Claim 2.** *For $n = 2$, $CC_{IC}$(NOT-TOO-FAR$_k$) $\geq 2k - 2$.*

Figure 2: The description of the 2-players NOT-TOO-FAR$_k$ social-choice function. For every profile of values for the players, the figure shows whether A wins, B wins or both. In all other profiles both lose. The hardness of this example is due to the fact that every two profiles of values for which the two players win (e.g., $x$ and $y$ in the picture) is associated with different payments. Note that the payments are determined by a projection on the diagonals defined by the social-choice function.

*Proof.* Consider the type-profiles $(v_1, v_2)$ such that both $v_1$ and $v_2$ are at least $1/2$. Observe, that there are exactly $2^{2k-2}$ such type-profiles, and that both players win for each such type profile.

We shall prove that for every two such type-profiles $v = (v_1, v_2)$ and $v' = (v'_1, v'_2)$ it must hold that any communication protocol that computes incentive-compatible payments outputs $p(v_1, v_2) \neq p(v'_1, v'_2)$. Therefore, any such protocol has at least $2^{2k-2}$ outcomes in its range. By Observation 3.2 this implies that the minimal number of bits that must be transmitted by any such protocol is at least $\log(2^{2k-2}) = 2k - 2$.

W.l.o.g., assume that $v_1 \neq v'_1$. From Observation 3.1 (see also Figure 2) one can deduce that in the event that both players win the payment of each is the other's player's value minus $1/2$. We shall show that $p_2(v_1, v_2) \neq p_2(v'_1, v'_2)$. Clearly $p_2(v_1, v_2) = v_1 - 1/2$ (the minimal value for which 2 would win). Similarly, $p_2(v'_1, v'_2) = v'_1 - 1/2$. Since $v_1 \neq v'_1$ we conclude that $p_2(v_1, v_2) \neq p_2(v'_1, v'_2)$. □

The theorem follows. □

### 5.2 A Linear Lower Bound for $n$ Players

We are now ready to prove the main theorem of this paper. The social-choice function for which we shall prove a lower bound of about $n$ is an extension of NOT-TOO-FAR$_k$ presented above for 2 players, to the $n$-players setting.

**Definition 7** ($\textsc{Not-Too-Far}_k$)**.**
*<u>Input:</u> valuations $v_0, ..., v_{n-1} \in \mathbb{N}$ represented by strings of $k$ bits.*
*<u>Output:</u> A set of winning players from $N$. Player $i$ wins if one of the following happen:*

1. $v_j \geq 1/2$ *for every* $j \in N$.
2. $v_j \geq 1/2$ *for every* $j \in N \setminus \{i\}$ and $v_i \geq v_{i+1 \ (mod\ n)} - 1/2$.

We first observe that the function $\textsc{Not-Too-Far}_k$ is implementable, and therefore the communication cost of incentive compatibility is well defined. Indeed, it is easy to see that if player $i$ wins in $\textsc{Not-Too-Far}_k$ and increases its bid, $i$ will still win.

**Observation 5.2.** *The social-choice function* $\textsc{Not-Too-Far}_k$ *is monotone.*

We are now ready to present the main result, a linear lower bound on the communication cost of incentive compatibility for a single-parameter SCF. This is done by showing that computing both the function and the payments ($CC_{IC}(f)$) requires lots of communication since there are many different payment-vectors the mechanism should distinguish between; also, computing the function alone ($CC(f)$) requires a low amount of communication the players announce whether their values is at least $1/2$ using 1 bit each, and the additional relevant information is held by up to two players ($v_i$ and $v_{i+1}$ for some $i$).

**Theorem 5.3.** *For any $\epsilon > 0$ and for $k$ large enough, the communication cost of incentive compatibility for the social-choice function* $\textsc{Not-Too-Far}_k$ *is at least $n - \epsilon$. In particular,* $CC_{IC}(\textsc{Not-Too-Far}_k) \geq n(k-1)$ *and* $CC(\textsc{Not-Too-Far}_k) \leq n + k + 1$.

*Proof.* Let $f_k = \textsc{Not-Too-Far}_k$. We show below that $CC_{IC}(f_k) \geq n(k-1)$ (Claim 3) and that $CC(f_k) \leq n + k + 1$ (Claim 4). From these two facts we derive that $\frac{CC_{IC}(f_k)}{CC(f_k)} \geq \frac{n(k-1)}{k+n+1} = n - \frac{n(n+2)}{n+k+1}$. Clearly for any $\epsilon > 0$ there is a $k$ such that this is larger than $n - \epsilon$. We next derive the promised bounds on $CC_{IC}(f_k)$ and $CC(f_k)$.

**Claim 3.** $CC_{IC}(\textsc{Not-Too-Far}_k) \geq n(k-1)$.

*Proof.* There are $2^{n(k-1)}$ type-profiles $(v_0, \ldots, v_{n-1})$ such that each $v_i$ is at least $1/2$. For all these type-profiles all players win. Let $(v_0, \ldots, v_{n-1})$ and $(v'_0, \ldots, v'_{n-1})$ be two different such type-profiles. Let $p(v_0, \ldots, v_{n-1})$ and $p(v'_0, \ldots, v'_{n-1})$ be the incentive-compatible payments outputted by a communication protocol for these two type-profiles. We shall show that these two payment vectors must be different. This is derived from the fact that $p_i(v_0, \ldots, v_{n-1}) = v_{i+1 \ (mod\ n)} - 1/2$, and similarly,
$p_i(v'_0, \ldots, v'_{n-1}) = v'_{i+1 \ (mod\ n)} - 1/2$. Hence, if any coordinate $j \in \{0, 1, ..., n-1\}$ is such that $v_i \neq v'_i$ this implies that $p_{j-1}(v_0, v_2, ..., v_{n-1}) \neq p_{j-1}(v'_0, v'_2, ..., v'_{n-1})$. Therefore, any protocol that computes incentive-compatible payments has at least $2^{n(k-1)}$ outcomes in the range of $f_k$ and thus requires at least $n(k-1)$ bits (by Observation 3.2). We conclude that $CC_{IC}(f_k) \geq n(k-1)$. □

**Claim 4.** $CC(\text{Not-Too-Far}_k) \leq n + k + 1$.

*Proof.* We show that $CC(f_k) \leq k + n + 1$ by exhibiting a communication protocol that computes $f_k$ and only requires $k + n + 1$ bits: First, each player $i$ transmits a single bit $b_i$ that indicates whether his value is at least $1/2$ ($i$ transmits 1 if $v_i \geq 1/2$). If $b_i = 1$ for all $i$ then all players win. If for two or more players $b_i = 0$ then all players lose. If there is a player $j$ such that $b_j = 0$ and for all other players it holds that $b_i = 1$ then all other players (but $j$) lose. In this case, in order to determine whether player $j$ wins, player $j + 1$ $(mod\ n)$ transmits all of his bits. Player $j$ (who now knows $v_{j+1\ (mod\ n)}$) checks whether $v_j \geq v_{j+1\ (mod\ n)} - 1/2$ (in which case player $j$ wins). He now broadcasts an additional bit informing the others of the result (1 indicating "I win" and 0 indicating "I lose"). Observe, that overall $k+n+1$ bits were transmitted, and that the protocol does indeed compute $\text{Not-Too-Far}_k$. So, $CC(f_k) \leq k + n + 1$. □

The theorem follows. □

# 6 Discussion: The General Case

We have shown that, subject to the natural normalization assumption (Definition 2), the communication cost of computing payments that support an outcome in equilibrium can be significant. Specifically, we have shown that this communication overhead cost can be linear in the number of players. An obvious open question is that of extending our lower bound to the case in which the normalization assumption is removed (or prove that no such lower bound is attainable in that case).

Recall the 2-player public good setting discussed in Section 2. This setting played an important part in advancing our understanding of the communication cost of incentive compatibility. It was through this setting that we were able to prove a tight result of 2 for 2-player settings. For this simple setting, we now prove a surprising result: if the normalization assumption is removed then the communication cost of incentive-compatibility in the 2-player public-good setting is negligible. In fact, similar constructions to the one presented in the next subsection imply that social-choice functions like the one used to prove our linear lower bound for $n$-player settings (Theorem 5.3) will not be helpful in extending our lower bound to the general case. At this point, it is unclear to us whether this is an evidence for the existence of a sub-linear upper bound for the general case, or merely implies that inherently different constructions are required to obtain a linear lower bound for the general case.

## 6.1 Removing the normalization assumption.

Let us revisit the 2-player public-good setting, where the values of the two players are once again in $\{0, \ldots, 2^{k-1}\}$ and $C = 2^k$. As explained in Section 4.2, for each pair of values

$(v_1, v_2)$ for which both players "win" (meaning $v_1 + v_2 \geq 0$), the players are required to pay $C - v_2$, and $C - v_1$, respectively. If the two players "lose", then they are both charged 0. If we no longer insist on the normalization assumption, this implies that for every value of $v_1$, we can force player 2 to pay, in addition to what he was required to pay so far, and regardless of whether he wins or loses (and of his value $v_2$), some extra amount $x_{v_1}$ that can only depend on the value of the *other* player. By monotonicity, as long as $x_{v_1}$ depends only on $v_1$ the resulting mechanism is incentive compatible.

Consider the following incentive-compatible payment scheme for every pair of values $(v_1, v_2)$ for which both players win, the players are required to pay $C - v_2$, and $C - v_1 + x_{v_1}$, respectively, where $x_{v_1} = v_1$. Observe that if both players win, player 1's payment is $C - v_2$, while player 2's payment is always $C$. In contrast, if both players lose, then player 1 is charged 0, but player 2 is charged $x_{v_1} = v_1$.

**A protocol.** The above suggests the following simple protocol:

1. Find out the value of $\min\{v_1, C - v_2\}$. Denote this value by $\alpha$.

2. If $\alpha \neq v_1$, then $\alpha = C - v_2$, and so $v_1 \geq C - v_2$ (which means that $v_1 + v_2 \geq C$). Then, output "Build", charge player 1 an amount of $C - v_2$, and charge player 2 an amount of $C$.

3. Otherwise, it must be that $\alpha = v_1$, and $v_1 < C - v_2$ (which means that $v_1 + v_2 < C$). Then, output "Do Not Build", do not charge player 1, and charge player 2 an amount of $v_1$.

Naor shows that computing the minimum of two values in $\{0, \ldots, 2^k - 1\}$ (i.e., who has the highest value and the actual value of the second highest) can be done via transmission of $k + O(\log k)$ bits (see proof in Babaioff et al. (2008)). Observe that as both $v_1$ and $C - v_2$ are such numbers, this implies that the protocol above only requires the transmission of $k + O(\log k)$ bits (once we learn $\alpha = \min\{v_1, C - v_2\}$ we have all the information we need to determine the outcome and compute payments).

**Insignificant communication cost of incentive-compatibility.** Recall that we showed a lower bound of roughly $2k$ for computing payments if the normalization assumption holds. Hence, the removal of this assumption can be allowed for mechanisms that are significantly more frugal in terms of communication. Specifically, determining the winner and computing supporting payments can be done with $k + O(\log k)$ bits, giving an upper bound on the communication cost of $1 + O(\frac{\log k}{k})$.

## 7 Open Questions and Directions for Future Research

We leave the following open questions:

1. The obvious question left open is proving lower bounds on the communication cost of incentive compatibility for the non-normalized case (see discussion in Section 6), or exhibiting a sub-linear upper bound for this case.

2. A big open question posed by Fadel and Segal (2009) is determining the communication cost of incentive compatibility in *multi-parameter* domains. Recall that for proving our main impossibility result, we constructed an environment with one-dimensional types where the communication cost is linear; with multi-dimensional types, one may hope to prove even stronger lower bounds. (Note that FS proved a linear upper bound for single parameter environments but only an exponential upper bound for the general case.) This would be a challenging task, as the characterization of implementable social-choice functions in multi-dimensional domains is less well understood.

3. It was shown by Lahaie and Parkes (2008) (LP) that in some multi-parameter welfare-maximization settings, the naïve VCG protocol that computes the efficient outcome $n + 1$ times in sequence (removing one player at a time) is asymptotically optimal. However, in such settings there could be *other* incentive-compatible normalized mechanisms that can perform significantly better. It would be interesting to strengthen the result of LP by exhibiting a welfare-maximizing social-choice function for which not only does a linear lower bound on VCG mechanisms exists, but VCG mechanisms are the *only* incentive-compatible normalized mechanisms (in the spirit of the work by Green and Laffont (1977)).

**Acknowledgements:**
The authors are grateful to Moni Naor, Ilya Segal and two anonymous referees for valuable discussions and comments. Liad Blumrosen was supported by the Israeli Science Foundation grant number 230/10, and his work was done in part while being a post-doc researcher at Microsoft Research, Silicon Valley. Michael Schapira was supported by NSF grant 0331548 and by grants from the Israel Science Foundation and the USA-Israel Bi-national Science Foundation, and his Work done in part while interning at Microsoft Research, Silicon Valley, as a graduate student at the Hebrew University of Jerusalem.

# References

Babaioff, M., Blumrosen, L., Naor, M., Schapira, M., 2008. Informational overhead of incentive compatibility. In: Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08). pp. 88–97.

Babaioff, M., Kleinberg, R. D., Slivkins, A., 2010. Truthful mechanisms with implicit payment computation. In: Proceedings of the 11-th ACM Conference on Electronic Commerce (EC'10). pp. 43–52.

Bikhchandani, S., Vohra, R., de Vries, S., Schummer, J., 2002. Linear programming and Vickrey auctions. Mathematics of the Internet: E-Auction and Markets, 75–115.

Blumrosen, L., Feldman, M., 2006. Implementation with a bounded action space. In: Proceedings of the 7th ACM Conference on Electronic Commerce (EC'06). pp. 62–71.

Blumrosen, L., Nisan, N., 2009. On the computational power of demand queries. SIAM Journal on Computing 39 (4), 1372–1391, an extended abstract in EC'05 contained preliminary results ,.

Blumrosen, L., Nisan, N., 2010. Informational limitations of ascending combinatorial auctions. Journal of Economic Theory 145 (3), 1203–1223, an extended abstract in EC'05 contained preliminary results.

Blumrosen, L., Nisan, N., Segal, I., 2006. Auctions with severely bounded communications. Journal of Artificial Intelligence Research 28, 233–266.

Dobzinski, S., Nisan, N., 2007. Limitations of VCG-based mechanisms. In: The 39th Annual ACM Symposium on Theory of Computing. pp. 338–344.

Emek, Y., Peleg, D., Roditty, L., 2008. A near-linear time algorithm for computing replacement paths in planar directed graph. In: SODA'08.

Fadel, R., Segal, I., 2009. The communication cost of selfishness. Journal of Economic Theory 144 (5), 1895–1920.

Green, J., Laffont, J., 1977. Characterization of satisfactory mechanism for the revelation of preferences for public goods. Econometrica, 427–438.

Green, J. R., Laffont, J.-J., 1987. Limited communication and incentive-compatibility. Information, Incentives and Economic Mechanisms: Essays in Honor of Leonid Hurwicz. T. Groves, R. Radner, and S. Reiter, (eds.),.

Hart, S., Mansour, Y., 2010. How long to equilibrium? the communication complexity of uncoupled equilibrium procedures. Games and Economic Behavior 69 (1), 107–126.

Hershberger, J., Suri, S., 2001. Vickrey prices and shortest paths: What is an edge worth? In: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science. p. 252.

Hurwicz, L., 1960. Optimality and informational efficiency in resource allocation. Arrow, Karlin and Suppes, editors, Mathematical Methods in Social Sciences., 413–424.

Hurwicz, L., 1977. On the dimensional requirements of informationally decentralized pareto-satisfactory processes. K.J. Arrow and L. Hurwicz, Editors, Studies in Resource Allocation, 413–424.

Hurwicz, L., Reiter, S., 2006. Designing Economic Mechanisms. Cambridge University Press.

Kushilevitz, E., Nisan, N., 1997. Communication Complexity. Cambridge University Press.

Lahaie, S., Parkes, D. C., 2008. On the communication requirements of verifying the VCG outcome. In: Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08). pp. 78–87.

Melumad, N., Mookherjee, D., Reichelstein, S., 1992. A theory of responsibility centers. Journal of Accounting and Economics 15 (4), 445–484.

Mirrlees, J., 1971. An exploration in the theory of optimum income taxation. Review of Economic Studies 38, 175–208.

Mookherjee, D., 2006. Decentralization, hierarchies, and incentives: A mechanism design perspective. Journal of Economic Literature 44 (2), 367–390.

Mookherjee, D., Reichelstein, S., 1992. Dominant strategy implementation of Bayesian incentive compatible allocation rules. Journal of Economic Theory 56 (2), 378–399.

Mount, K., Reiter, S., 1974. The informational size of message spaces. Journal of Economic Theory 8 (2), 161–192.

Myerson, R. B., 1981. Optimal auction design. Mathematics of Operations Research 6 (1), 58–73.

Nisan, N., Ronen, A., 2001. Algorithmic mechanism design. Games and Economic Behaviour 35, 166 – 196, a preliminary version appeared in STOC 1999.

Nisan, N., Segal, I., 2006. The communication requirements of efficient allocations and supporting prices. Journal of Economic Theory 129 (1), 192 – 224.

Papadimitriou, C. H., Schapira, M., Singer, Y., 2008. On the hardness of being truthful. In: 49th Annual IEEE Symposium on Foundations of Computer Science. pp. 250–259.

Reichelstein, S., 1984. Incentive compatibility and informational requirements. Journal of Economic Theory 34 (1), 32–51.

Segal, I., 2006. In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 11. The Communication Requirements of Combinatorial Allocation Problems. MIT Press.

Segal, I., 2007. The communication requirements of social choice rules and supporting budget sets. Journal of Economic Theory 136, 341–378.

Segal, I., 2010. Nash implementation with little communication. Theoretical Economics 5 (1), 51–71.

Vickrey, W., 1961. Counterspeculation, auctions and competitive sealed tenders. Journal of Finance, 8–37.

Yao, A. C.-C., 1979. Some complexity questions related to distributive computing. In: ACM Symposium on Theory of Computing. pp. 209–213.

# A Communication Complexity - Formal Definition

Consider environments with $n$ players, each player $i$ privately holds a piece of information $v_i$ (his type), and the goal is compute a function $f(v_1, ..., v_n)$. Informally speaking, the communication complexity of $f$ is the minimal number of bits required (in the worst case) to compute this function. For a formal definition, we start by defining communication protocols using binary trees, where at each node some player needs to decide what bit to communicate.

**Definition 8.** *A communication protocol $\mathcal{P}$ with a set of players $N$, type space $V = V_1 \times ... \times V_n$, and an outcome space $O$ is a binary tree with a set of nodes $U$ and a set of leaves $L \subset U$, where:*

- *The set of internal nodes $U \setminus L$ is partitioned to $n$ subset $U_1, ..., U_n$, one per each player. Each set $U_i$ represents the nodes where decisions are made by player $i$.*
- *Each leaf $l \in L$ is labeled with an outcome $o(l) \in O$*
- *Each player $i$ has a strategy function that maps his type to a decision in each one of his decision nodes, $\sigma_i : V \to \{0,1\}^{U_i}$.*

*For every decision profile $s = (s_1, ..., s_n) \in \prod_{i \in N} \{0,1\}^{U_i}$, let $p(s)$ denote the leaf that is reached when each player $i$ uses the decisions specified in $s_i$.*

*Finally, we say that a protocol $\mathcal{P}$* computes *the function $f : V \to O$ when for every $v \in V$ we have $f(v) = o(p(\sigma(v)))$.*

The *cost* of a protocol $\mathcal{P}$ is the height of its binary tree, that is, the maximum number of edges between the root node and a leaf. The height of the tree represents the longest worst-case execution of the protocol. The *communication complexity* of a function is the minimal cost of a protocol that computes this function:

**Definition 9.** *The* communication complexity *$CC(f)$ of a function $f : V \to O$ is defined as the minimal cost of a protocol $\mathcal{P}$ that computes $f$, over all protocols that compute $f$.*

# B Missing proofs

**Lemma B.1.** *In the setting described in the proof for Lemma 4.4 in Theorem 4.2, the number of valuation profiles L where $\Sigma_i v_i = C$ holds that:*

$$\log L \geq (\frac{n}{2} - 1)k_c - n \cdot \log n \tag{4}$$

*Proof.* We first consider the case where $C \leq 2^k - 1$. We consider the following family of type-profiles: $v_1 = C - 2^{k_c - 1}, v_2 = ... = v_{\frac{n}{2}-1} = 0$, and $\Sigma_{i=\frac{n}{2}}^n v_i = 2^{k_c-1}$. Observe, that any type-profile in this family is such that $\Sigma_{i=1}^n v_i = C$. How many such type-profiles are there? There are $\binom{2^{k_c-1}+\frac{n}{2}}{\frac{n}{2}}$ ways to distribute $2^{k_c-1}$ between $\frac{n}{2}+1$ players. (Recall that we study a discrete environment.) This is bounded from below by $\frac{2^{(k_c-1)\frac{n}{2}}}{(\frac{n}{2})^{\frac{n}{2}}}$ (as $\binom{a+b}{b} \geq \frac{a^b}{b^b}$). So, any protocol that determines whether $\Sigma_i v_i \geq C$ needs to transmit at least $\log(\frac{2^{(k_c-1)\frac{n}{2}}}{(\frac{n}{2})^{\frac{n}{2}}}) = \frac{n}{2}k_c - \frac{n}{2}(\log(n))$ bits.

What if $C \geq 2^k - 1$? We will now describe another set of valuations as our fooling set. We will distribute $2^k - 1$ to players $\frac{n}{2}+1, ..., n$, and distribute the rest of the cost $C - (2^k - 1)$ to players $1, ..., \frac{n}{2}$ in some arbitrary way. We first argue that we can indeed distribute $C-(2^k-1)$ to players $1, ..., \frac{n}{2}$, namely, this amount does not exceed the maximal total value they can have. Indeed, the total possible value of $\frac{n}{2}$ players is $(2^k-1)\cdot\frac{n}{2}$, and for large enough $k$ it holds that $(2^k - 1)\cdot\frac{n}{2} \geq 2^k \cdot \frac{n}{2} - (2^k - 1) \geq C - (2^k - 1)$.

So now we are left with a cost of $C' = 2^k - 1$ to distribute between agents $\frac{n}{2}+1, ..., n$. We can now achieve $L$ by looking at the different ways to distribute $C'$ between players $\frac{n}{2}+1, ..., n$, i.e., such that $\Sigma_{i=\frac{n}{2}+1}^n v_i = 2^k - 1$. How many such type-profiles are there? There are $\binom{(2^k-1)+\frac{n}{2}-1}{\frac{n}{2}-1}$ ways to distribute $C'$ between $\frac{n}{2}$ players. Now

$$\binom{(2^k-1)+(\frac{n}{2}-1)}{\frac{n}{2}-1} \geq \frac{(2^k-1)^{(\frac{n}{2}-1)}}{(\frac{n}{2}-1)^{(\frac{n}{2}-1)}} \tag{5}$$

So, any protocol that determines whether $\Sigma_i v_i \geq C$ needs to transmit at least a logarithmic factor of the right hand side of Eq.(5), which is $(\frac{n}{2}-1)\log(2^k-1) - (\frac{n}{2}-1)\log(\frac{n}{2}-1)$. By the conditions of Lemma 4.4 it holds that $\frac{n}{2}\cdot 2^k \geq C$. Now $\frac{n}{2}\cdot 2^k \geq C \geq 2^{k_c-1}$ thus $2^k \geq \frac{2^{k_c}}{n}$. We conclude that $\log\ L \geq (\frac{n}{2}-1)\log(\frac{2^{k_c}}{n}-1) - (\frac{n}{2}-1)(\log(\frac{n}{2}-1)) \geq (\frac{n}{2}-1)k_c - (\frac{n}{2}-1)\cdot(2\log(n)) \geq (\frac{n}{2}-1)k_c - n\cdot\log(n)$ □