Informational Overhead of Incentive Compatibility

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

Liad Blumrosen
Microsoft Research
Silicon Valley, California
liadbl@microsoft.com

Moni Naor
Dept. of Computer Science
and Applied Mathematics,
Weizmann Institute of
Science, Rehovot, Israel
moni.naor@weizmann.ac.il

Michael Schapira*
School of Engineering and
Computer Science,
Hebrew University of
Jerusalem, Israel
mikesch@cs.huji.ac.il

ABSTRACT

In the presence of self-interested parties, mechanism designers typically aim to achieve their goals (or social-choice functions) in an equilibrium. In this paper, we study the cost of such equilibrium requirements in terms of communication, a problem that was recently raised by Fadel and Segal [15]. 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 prices (even if such prices are known to exist).

Our main result shows that the total communication needed for this task can be greater than x by a factor linear in the number of players n, i.e., $n \cdot x$. This is the first known lower bound for this problem. In fact, we show that this result holds even in single-parameter domains (under the common assumption that losing players pay zero). On the positive side, we show that certain classic economic objectives, namely, single-item auctions and public-good mechanisms, only entail a small overhead. Finally, we explore the communication overhead in welfare-maximization domains, and initiate the study of the overhead of computing payments that lie in the core of coalitional games.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of algorithms and problem complexity; J.4 [**Social and behavioral sciences**]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Mechanism Design, Communication Complexity, Ex-post Implementation

1. INTRODUCTION

Consider the goal of implementing algorithms in environments with self-interested players. We seek algorithms that admit the following two preliminary 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 algorithm. Second, *incentive compatibility*, i.e., the *existence* of some payment scheme that supports the implementation of the algorithm in equilibrium. In this work, we show that tractability and the existence of supporting payments are insufficient to establish that implementing the algorithm in equilibrium will indeed be simple. This is due to the fact a non-trivial amount of *additional* communication between the different parties may be required in order to compute the equilibrium-supporting prices.

The question of how much overhead one incurs from the computation of incentive-compatible payments was recently introduced by Fadel and Segal [15], 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 only 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 shows that the amount of communication that is required to compute equilibrium-supporting prices may increase the communication complexity of the algorithm by a factor that is linear in the number of players.

Theorem: [**Informal**] *There are social-choice functions such that their outcome can be computed by communicating x bits, but determining both the outcome* and *equilibrium-supporting prices may require about $n \cdot x$ bits of communica-*

*Supported by grants from the Israel Science Foundation.

EC'08, July 8–12, 2008, Chicago, Illinois, USA.

tion, where n is the number of players.

We prove that this result holds even for very simple 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 lower bound and actually strengthen our upper bounds.) Whether a similar result can be proven without the normalization assumption is left as an open question. Our lower bound is the first evidence that the communication overhead due to the demand for incentive compatibility may be significant. This linear factor in the number of players may become substantial in large-scale electronic-commerce systems.

Informally, in order to prove our main result we need to construct a social-choice function f for which the following requirements hold:

1. f can be implemented in ex-post equilibrium (in single-parameter domains this means that f should be monotone, see Section 2).

2. f can be computed with low communication complexity.

3. Computing the equilibrium-supporting prices requires high communication complexity.

The difficulty in finding such a social-welfare function is demonstrated by contrasting such desirable functions with 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 upper bounds for these two problems, by showing that the additional information required to compute the equilibrium-supporting prices 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 players have privately known utilities from using the bridge $v_1, ..., v_n$ and the bridge should be built only if $\sum_{i=1}^{n} v_i \geq C$ where C is its construction cost. We prove that computing the outcome plus the payments merely requires about three times the communication requirements of computing the outcome alone. Hence, the overhead in this case is small.

Single-item auction: In a single-item auction, players have private values $v_1, ..., v_n$ 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 needed to determine the allocation alone. For the special case of $n = 2$, we prove an even better upper bound. The 2-player problem turns out to be equivalent to the following interesting communication complexity problem: there are 2 players, each holding a number represented by k bits. What is the communication complexity of computing the minimum of these two numbers (such that both players will know the result)? While proving an upper bound of $k + O(\sqrt{k})$ is fairly easy, we prove even a better bound of $k + O(\log k)$.

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. This is due to the fact that it is likely to involve the construction of multi-parameter non-welfare maximizing social-choice functions, a class of functions that is little understood.

Welfare Maximizing Social-Choice Functions. Finally, we turn our attention to welfare-maximizing environments. This is the prominent example of an objective function that can always be implemented in equilibrium (using the family of VCG mechanisms). We discuss two well studied pricing schemes: (1) prices that support an ex-post Nash equilibrium. (2) price levels in the core.

Fadel and Segal [15] showed that maximizing social welfare in ex-post equilibrium incurs a low communication overhead. However, their simple solution, that belongs to the VCG family (each player is paid the sum of the values of the others), is impractical in many settings since it involves paying the players. Under the assumption of *No-Positive-Transfers*, we argue that the solution to this problem is not straightforward anymore. We show a simple linear upper bound on the informational overhead of computing the equilibrium-supporting prices, for a special case of this problem, and pose the following two open questions: (1) can a similar linear upper bound can be proven for general valuations? (2) does a matching lower-bound exist, even for our special case?

Finally, we initiate the study of the complexity of computing payment schemes that lie in the core of the respective coalitional game. A payment scheme is in the core if there is no subset of players that can deviate (with the "seller") and reach an outcome that betters their payoffs. We present a simple argument showing that computing core outcomes incurs, in general, a low informational overhead. Unfortunately, this positive result necessitates the computation of points in the core that are unreasonable in practice. A more reasonable core outcome is the one defined by Ausubel and Milgrom [1]. We show that the communication overhead of computing such core points is unbounded.

1.1 Related Work

Fadel and Segal [15] were the first to study the communication overhead of incentive compatibility. They proved exponential upper bounds, both for Bayesian-Nash equilibria and an 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 overhead in ex-post implementation be matched by a lower bound?[1] [15] also proves a linear (in

[1]This question appears to be hard to solve. The overhead in known to be at most linear (in the number of players) for welfare-maximization objectives and in single-parameter domains [15]. In other (multi-dimensional, arbitrary objectives) domains, the space of implementable social-choice function is not well understood. Therefore, constructing such a negative example is hard.

the number of players) upper bound on the communication overhead of incentive compatibility in single-parameter domains. Both [15] and our paper belong to a more general line of research studying communication and information aspects of various economic environments, for example in auctions [12, 5, 6] and in more general economic domains [4, 14]. A recent survey on this line of research in the context of combinatorial auctions is found in [13]. The basic model for communication complexity was presented by Yao [16]. A survey on communication complexity can be found in [9].

The question of the computational burden of computing payments in social-welfare maximizing environments has received some attention in the past: one of the open questions raised in the seminal paper by Nisan and Ronen [11] is whether VCG payments could be discovered with a smaller computational cost than that of the in naïve solution that solves $n+1$ separate problems (once for determining the outcome, and n additional solutions where players are excluded in turns). Indeed, [8] proved that when auctioning shortest paths, VCG prices can be determined by solving a single shortest-path problem (see also the recent work by [7]). [3] showed, via linear-programming duality, that calculating VCG prices can be done in certain domains by solving two optimization problems. The above results consider *computational* complexity, while in this paper we measure the additional *communication* complexity of computing the appropriate payments.

1.2 Organization of the Paper

The rest of the paper is organized as follows: we present our model and notations in Section 2. We prove a constant upper bound on the informational overhead of incentive compatibility for the classic model of single-item auctions in Section 3. In Section 4, we prove our results for public goods problems. We present a construction that proves a linear lower bound in Section 5. Finally, in Section 6, we explore welfare-maximizing social-choice functions.

2. BACKGROUND AND THE MODEL

2.1 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. 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_i. A *social-choice function* (SCF) is a function that assigns every n-tuple of players' valuation functions $v = (v_1, ..., v_n) \in V_1 \times \ldots \times V_n$ ("*type-profile*") an outcome $o \in O$. Each v_i is private and only known to i.

A *payment function* is a function $p : V_1 \times \ldots \times V_n \to \mathbb{R}^n$.

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:*

$$\forall v = (v_1, ..., v_n) \in V_1 \times \ldots \times V_n,\ \forall i \in N,\ \forall v_i' \in V_i,$$

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

(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 on 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" ($f(v) \subseteq N$ is the set of winners). The 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. The following is a well known characterization of implementable single-parameter social-choice functions.

DEFINITION 2. *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. $v_i', v_i \in V_i$ it holds that if $i \in f(v_i, v_{-i})$ then $i \in f(v_i', v_{-i})$.*

The following observation is well known (see, e.g., [10] and the reference therein).

OBSERVATION 2.1. *A single parameter SCF f is implementable if and only if it is monotonic. In the case of normalized mechanisms a winner has to pay the minimal bid she has to declare in order to win.*

2.2 Communication Overhead 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. 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).

How much additional communication burden is imposed by the necessity to compute payments that guarantee truthfulness? Let $CC(f, p)$ 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.).

In order to formally define the informational overhead 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 $k \in \mathbb{N}$. 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 3. *The* informational overhead of incentive compatibility *of f_k is defined to be $\frac{CC(f_k, p)}{CC(f_k)}$.*

Our main result shows that for some social-choice functions f_k a significant informational 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 $v_i = t_i \cdot 2^{-k}$.

2.3 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 [9].

OBSERVATION 2.2. *If the range of function f is of size m ($|Range(f)| = m$), then any communication protocol for f requires at least* $\log(m)$ *bits.*

OBSERVATION 2.3. *For any implementable function f_k with n players each holding k bits it holds that $CC(f_k) \leq k(n-1) + \lceil \log(|Range(f_k)|) \rceil$ and $CC(f_k, p) \leq kn$ ($Range(f_k)$ is the set of different outcomes in the range of f_k).*

PROOF. This is shown by considering two trivial protocols: To compute f_k, each of the players but the last one transmits all his information, and the last player computes the outcome and transmits the outcome. As there are $|Range(f_k)|$ possible relevant outcomes, $\lceil \log(|Range(f_k)|) \rceil$ bits are clearly sufficient to encode all these outcomes. To compute $CC(f_k, p)$ 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 be able to distinguish between every two of them. Fooling sets arguments are based on following well known property of communication protocols. If a protocol executes exactly the same on two inputs, it must do so on any possible "combinations" of these inputs, thus must output the same outcome. For completeness, we give a formal definition of the fooling-set technique in Appendix A.1.

3. SINGLE-ITEM AUCTIONS

A well known economic setting is the single-item auction, where a seller aims to sell an item to the bidder who values it the most.

DEFINITION 4 (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 for the outcome is 0. From SINGLE-ITEM-AUCTION we can derive the function SINGLE-ITEM-AUCTION$_k$ for the case that $v_i = t_i \cdot 2^{-k}$ for some integer $t_i \in \{0,...,2^k-1\}$ and is represented by a k-bit string. It is well known that if the winner pays the second highest price then the auction is truthful.

Next we show that for SINGLE-ITEM-AUCTION the informational overhead of incentive compatibility is at most a small constant (3).

PROPOSITION 3.1. *The informational overhead 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 $CC(f_k, p) \leq 2 \cdot CC(f_k) + k \leq 3 \cdot CC(f_k)$. The last weak inequality is a result of the following claim:[2]

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 A.1 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 be mapped by a protocol that computes f to the same monochromatic rectangle (exactly the same execution of the protocol which leads to the same outcome). Let (v_1, v_2) and (v_1', v_2') be two such type-profiles that are mapped to the same rectangle. 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 be mapped to the same rectangle. However, this leads to a contradiction because in this case player 2 should win. Hence, there are at least 2^k rectangles, and so any protocol that computes f must transmit at least k bits. □

To show that $CC(f_k, p) \leq 2CC(f_k) + k$ we present a simple protocol for $CC(f_k, p)$: run the protocol for f_k to find the player with highest value. Remove the highest bidder and run the protocol for f_k again. Now the players know who is the player with the second highest value. Finally, this player transmits his value, which requires k more bits.

3.1 An improved analysis for $n = 2$

Auctioning a single item is probably the most fundamental problem in mechanism design. In this section we present a better analysis of the communication burden in second-price auctions with two players. We present an iterative mechanism where players broadcast their information in turns, and an alphabet-changing trick that allows us to save in information (a similar protocol without changing alphabets proves a $1 + O(\sqrt{k})$ bound).

THEOREM 3.2. *For the social-choice function* SINGLE-ITEM-AUCTION$_k$ *and $n = 2$, the information overhead of incentive compatibility is at most* $1 + O(\frac{\log(k)}{k})$.

The proof appears in Appendix B. We note that the proof of this result implies that the informational overhead of incentive compatibility for SINGLE-ITEM-AUCTION with $n = 2$ is extremely small and tend to 1 very fast with k.

4. 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 exceeds the construction cost C. The function is defined given the parameter $C \geq 0$.

DEFINITION 5 (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.

[2]We prove this claim for completeness. Similar claims were proven before, e.g., in [12].

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 that $v_i = t_i \cdot 2^{-k}$ for some integer $t_i \in \{0, ..., 2^k - 1\}$ and is represented by a k-bit string.

In the next section we consider the problem for the case of 2 agents. We show that the informational overhead of incentive compatibility for that case is almost 2. Yet, in the following section we show that when moving to an n-player setting the overhead remains constant and does not grow with n.

4.1 A Lower Bound for 2-Player Settings

We start by considering the case of only 2 players and show that there is lower bound that is almost 2 on the informational overhead of incentive compatibility.

PROPOSITION 4.1. *Assume $n = 2$. For any $\epsilon > 0$, for any large enough k there exists a cost C such that the informational overhead of incentive compatibility for the social-choice function C-PUBLIC-GOOD$_k$ is at least $2 - \epsilon$.*

PROOF. To prove the claim we consider the SCF $f_k = C$-PUBLIC-GOOD$_k$ for the case that the cost C of the public good is $1 - 2^{-k}$.

By Observation 2.3 it holds that $CC(f_k) \leq k+1$ (as there are only 2 possible outcomes $|Range(f_k)| = 2$). In order to prove the proposition we show below that $CC(f_k, p) \geq 2k - 1$. This implies that the informational overhead of incentive compatibility of C-PUBLIC-GOOD$_k$ is at least $\frac{2k-1}{k+1} = 2 - \frac{3}{k+1}$. Clearly for any $\epsilon > 0$ there is a k such that this is larger than $2 - \epsilon$. To conclude the proof of the theorem we next show that $CC(f_k, p) \geq 2k - 1$.

CLAIM 2. *When $n = 2$ and $C = 1 - 2^{-k}$, $CC(C$-PUBLIC-GOOD$_k, p) \geq 2k - 1$.*

PROOF. Figure 1 describes the function f_k. We shall call any pair of possible values (v_1, v_2) a *type-profile*. Observe, that there are 2^{2k} possible type-profiles, out of which $2^{2k-1} + 2^{k-1} > 2^{2k-1}$ are type profiles in which $v_1 + v_2 \geq C$. We shall prove that for every two type-profiles $v = (v_1, v_2)$ and $v' = (v'_1, v'_2)$, such that $v_1 + v_2 > C$, $v'_1 + v'_2 > C$, and $v \neq v'$, it must hold that any protocol that computes incentive-compatible payments outputs $p(v_1, v_2) = (p_1(v_1, v_2), p_2(v_1, v_2)) \neq p(v'_1, v'_2) = (p'_1(v_1, v_2), p'_2(v_1, v_2))$. This would imply that f_k has at least 2^{2k-1} different outcomes in its range, which means that at least $log(2^{2k-1}) = 2k - 1$ bits must be transmitted (by Observation 2.2).

Let $v = (v_1, v_2)$ and $v' = (v'_1, v'_2)$ be two type-profiles such that $v_1 + v_2 > C$, $v'_1 + v'_2 > C$, and $v \neq v'$ (different type-profiles in which both players win). W.l.o.g. assume that $v_1 \neq v_2$. Then, we shall show that $p_2(v_1, v_2) \neq p_2(v'_1, v'_2)$. As shown in Figure 1, the payment of a winning player must be the minimal value it needed in order to win. That is, the payment of player 1 is the horizontal projection on the diagonal line, and the payment of player 2 is the vertical projection on the diagonal line. Hence, if the value of player 1 is v_1 then the payment of player 2, $p_2(v_1, v_2)$, is exactly $C - v_1$ (that is, the minimal value of player 2 for which they both win). Similarly, $p_2(v'_1, v'_2) = C - v'_1$. 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.

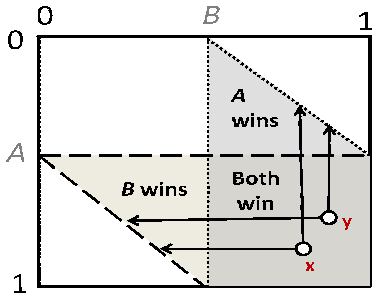

Figure 1: The description of the 2-NOT-TOO-FAR social-choice function. For every profile of values for the players, the figure shows whether A wins, B wins or both. In all other profiles both lose. The hardness of this example is due to the fact that every two profiles of values for which the two players win (e.g., x and y in the picture) is associated with different prices. Note that the prices are determined by a projection on the diagonals defined by the social-choice function.

4.2 A Constant Upper Bound for n-Player Settings

As we have seen, for $n = 2$ the informational overhead of the public good function is essentially 2.[3] One might hope that a generalization of this function to an n-player setting leads to a lower bound of n. Yet we show that this is not the case.

THEOREM 4.2. *Fix $\epsilon > 0$. For any C, the informational overhead 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 $2 \cdot \frac{n}{n-1} + \epsilon < 3 + \epsilon$.*

We refer the reader to Appendix C for the proof of the theorem.

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 substantially (essentially factor n) smaller that $CC(f, p)$.

5.1 Another Lower Bound for 2 Players

The reader may expect a straightforward generalization of the 2-player public-good problem to enable us to get an $\Omega(n)$ lower bound for general single-parameter domains. However, as shown in Section 4.2 this is not the case. In fact, the n-player public-good problem is such that the communication cost of selfishness is never greater than $3 + \epsilon$.

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 1. As before, we have 2 players $1, 2$, each holding a

[3]It is relatively easy to derive a matching upper bound of about 2 as it is clear that $2k$ bits are always sufficient to compute the payments and it is relatively easy to show that k bits are necessary to compute the outcome.

value $v_i = t_i \cdot 2^{-k}$ for an integer $t \in \{0, ..., 2^k - 1\}$, represented by a k-bit string. Player i's utility from winning is v_i, and his utility from losing is 0. Player 1 wins if and only if $v_2 \geq 1/2$ and $v_1 \geq v_2 - 1/2$. Similarly, player 2 wins if and only if $v_1 \geq 1/2$ and $v_2 \geq v_1 - 1/2$. 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 its informational overhead of incentive compatibility is finite.

PROPOSITION 5.1. *For any $\epsilon > 0$, for any k large enough the informational overhead 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 2.3 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 $CC(f_k, p) \geq 2k - 2$. From this we derive that $\frac{CC(f_k,p)}{CC(f_k)} \geq \frac{2k-2}{k+2} = 2 - \frac{6}{k+2}$. Clearly as this a monotonic function of k that converge to 2, for any $\epsilon > 0$ there is a k such that this is larger than $2 - \epsilon$. We next derive the promised lower bound on $CC(f_k, p)$.

CLAIM 3. *For $n = 2$, CC(NOT-TOO-FAR$_k, p) \geq 2k - 2$.*

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 2.2 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 Figure 1 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 2-NOT-TOO-FAR$_k$ to n-player settings.

DEFINITION 6 (NOT-TOO-FAR$_k$).
Input: valuations $v_0, ..., v_{n-1} \in \mathbb{N}$ represented by strings of k bits.
Output: 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 $j \in N, j \neq i$, and $v_i \geq v_{i+1 \ (mod\ n)} - 1/2$.*

We first observe that the function NOT-TOO-FAR$_k$ is implementable, and therefore the informational overhead of incentive compatibility is well defined. Indeed, it is easy to see that if player i wins in NOT-TOO-FAR$_k$ and increases its bid, i will still win.

OBSERVATION 5.2. *The social-choice function* NOT-TOO-FAR$_k$ *is monotone.*

We are now ready to present the main result, a linear lower bound on the informational overhead of incentive compatibility for a single-parameter SCF. This is done by showing that computing both the function and the payments (CC(f,p)) requires lots of communication since there are many different price-vectors the mechanism should distinguish between; also, computing the function alone (CC(f)) requires a low amount of communication since after all players declared if their value exceeds 1/2 using 1 bit each, the only 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 informational overhead of incentive compatibility for the social-choice function* NOT-TOO-FAR$_k$ *is at least $n - \epsilon$.*

PROOF. Let $f_k =$NOT-TOO-FAR$_k$. We show below that $CC(f_k, p) \geq n(k-1)$ (Claim 4) and that $CC(f_k) \leq n+k+1$ (Claim 5). From these two facts we derive that $\frac{CC(f_k,p)}{CC(f_k)} \geq \frac{n(k-1)}{k+n+1} = n - \frac{n(n+2)}{n+k+1}$. Clearly for any $\epsilon > 0$ there is a k such that this is larger than $n - \epsilon$. We next derive the promised bounds on $CC(f_k, p)$ and $CC(f_k)$.

CLAIM 4. CC(NOT-TOO-FAR$_k, p) \geq n(k-1)$.

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, as shown in Proposition 4.1, 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 2.2). We conclude that $CC(f_k, p) \geq n(k-1)$. □

CLAIM 5. CC(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 NOT-TOO-FAR$_k$. So, $CC(f_k) \leq k + n + 1$. □

The theorem follows.

6. WELFARE MAXIMIZATION: VCG AND THE CORE

In this section we study social-choice functions of a particular type – social choice functions that maximize the social welfare. We consider two payment scheme for such environments. In Subsection 6.1 we study prices supporting ex-post Nash equilibria (which are VCG payments under some conditions). In Subsection 6.2, we introduce the problem of measuring the communication overhead of computing points in the core.

Recall that we consider mechanism-design settings with a social-choice function f, outcome set O and valuation spaces $V_1, ..., V_n$. Note that in this section we allow multi-parameter valuations.

DEFINITION 7. *A social choice function f is* welfare maximizing *if for every* $\mathbf{v}$, $f(v) \in argmax_{o \in O} \sum_{i=1}^{n} v_i(o)$.

6.1 Ex-post equilibrium: VCG prices

6.1.1 VCG prices: Overview

Recall that with VCG prices the welfare maximizing outcome o is chosen, and each player pays $h_i(v_{-i}) - \sum_{j \neq i} v_j(o)$ (where the $h_i(\cdot)$ is a function that does not depend on v_i). It is well known that such payments, for any choice of the h_i's, supports a dominant strategy equilibrium (in our case, the mechanisms are iterative and the ex-post equilibrium concept is used).

Probably the most common choice for the h_i terms for VCG mechanisms is the Clarke pivot rule, where the payment of each player becomes $max_{a \in O} \sum_{j \neq i} v_i(a) - \sum_{j \neq i} v(o)$. This choice of the h_i's implies that the mechanism has the following important property:[4]

DEFINITION 8. *A payment scheme has* No Positive Transfers (NPT) *if for every valuation profile v and every player i, $p_i(\mathbf{v}) \geq 0$.*

6.1.2 Communication Overhead

Fadel and Segal [15] gave an elegant solution to the communication overhead of computing some equilibrium prices. They show that the overhead is very small: it is an additive term that is not related to the complexity of the social-choice function, but only depends on the representation of values by the players. Recall that the value each player has for an alternative is represented by at most k bits.

PROPOSITION 6.1 ([15]). *For any welfare maximizing function, $CC(f,p) \leq CC(f) + n \cdot k$.*

However, Proposition 6.1 computes VCG prices with positive transfers to the players. When this is not allowed, the bound on the communication level is no longer trivial. Under the NPT assumptions, we were able to bound the communication overhead for settings where every bidder has a *zero valuation*. That is, a player with such a valuation will never affect the outcome. For example, in a combinatorial-auction model, if the player may have a valuation that gives him a value of 0 for every possible bundle of items he gets, this is a zero valuation. Note that settings may exist where the valuation space has no such valuation (e.g., in combinatorial auction a player may always have a value of 10 for either item a or item b).

A player that reports a valuation function where the value for every bundle is zero, will actually ensure that the outcome is computed regardless of his preferences.

DEFINITION 9. *We say that a player i has a* zero valuation *if there exists $v_i^0 \in V_i$ such that for every v_{-i}, $f(v_i^0, v_{-i}) \in argmax_{o \in O} \sum_{j \neq i} v_i(o)$.*

Let $CC_{NPT}(f,p)$ denote the communication complexity of computing the social-choice function f and some payments with no positive transfers. Note that this definition allows us to find payments that are not necessarily VCG with the Clarke pivot rule; this makes the construction of lower bounds for this measure actually harder.

PROPOSITION 6.2. *Assume that every player has a zero valuation. Then, for every welfare maximizing function f, $CC_{NPT}(f,p) \leq (n+1)CC(f) + n \cdot k$.*

PROOF. We are given a "black-box" that computes the function $f(v)$ using $CC(f)$ bits for every valuation profile v. We will use this black box once for computing $f(v)$, and n more times for (v_i^0, v_{-i}), where $i = 1, ..., n$. For computing the VCG prices, we should also ask the players for $v_i(o)$ (o is the chosen alternative), which may take up to $\epsilon = n \cdot k$ bits (again, k is the precision or number of bits of the value of a player for a single alternative) and this term is independent of $CC(f)$ and is therefore negligible. □

We direct the attention of the reader to the following open questions. Any non-linear upper or lower bounds in this context will be considered surprising.

Open questions: For welfare maximizing social-choice functions, can equilibrium supporting prices that admit NPT can be computed:

1. with a linear communication overhead, even without assuming the zero-valuation property?
2. with o(n) communication overhead, even with the zero-valuation property?

6.2 Core Outcomes

Consider a coalitional game with a set of $n+1$ players $N^+ = \{0, 1, ..., n\}$ (player 0 is interpreted as the "seller") – where the coalitional value is defined as follows:

- $w(S) = max_{o \in O} \sum_{i \in S} v_i(o)$, if $0 \in S$.
- $w(S) = 0$, if $0 \notin S$.

A vector $\pi \in \mathbb{R}^{n+1}$ is in the *core* of this coalitional game, if $w(N^+) = \sum_{i \in N^+} \pi_i$ and for every $S \subset N^+$ we have $w(S) \leq \sum_{i \in S} \pi_i$. Intuitively, π should be thought of as a vector of payoffs (utilities), and this payoff vector is in the core if no subset of players can have a better deal with the social planner (the "seller"). It is known ([1]) that the core of this game is non-empty, and that every core allocation is welfare maximizing.

[4] It also guarantees another close property of *individual rationality*, that is, the utility of the players is never negative.

DEFINITION 10. *A social-choice function f and a payment scheme p are in the core, if for every profile of valuations $\mathbf{v}$, the vector of payoffs*

$$(v_1(o) - p_1(\mathbf{v}), v_2(o) - p_2(\mathbf{v}), ..., v_n(o) - p_n(\mathbf{v}))$$

is in the core of the respective coalitional game (where $f(\mathbf{v}) = o$).

We denote the communication complexity of determining the function f and some set of payments p in the core as $CC_{core}(f, p)$.

It turns out that computing a core outcome can be done with a small communication overhead that is independent of the communication complexity of f. We can thus prove a result that is similar in spirit to Proposition 6.1.

THEOREM 6.3. $CC_{core}(f, p) \leq CC(f) + n \cdot k$

PROOF. The payoff vector (0,...,0) is always in the core (i.e., every player pays his exact value for the outcome: $p_i(\mathbf{v}) = v_i(o)$). It is easy to see that no coalition can then deviate. Thus, we can first compute the welfare-maximizing outcome using $CC(f)$ bits, and then ask each bidder to report his value from this outcome using k bits. □

As in the VCG case, the above positive result is unrealistic, since eliciting the exact valuation of rational players seems to be impossible. A probably more reasonable point in the core is the one described by Ausubel and Milgrom [1]. Their payment method can be computed using an iterative ascending-price bidding process, and also have some desired incentive properties (see more details in [1, 2]). These prices are computed by an ascending-price auctions that use personalized bundle prices. A formal definition of this pricing scheme can be found in [1], together with a detailed comparison to VCG payments. We denote the communication complexity of computing a welfare-maximizing allocation and the Ausubel-Milgrom core payments by $CC^{AM}_{core}(f, p)$. We show by a trivial example that the communication overhead of computing the Ausubel-Milgrom outcome over computing the welfare-maximizing outcome is unbounded.

PROPOSITION 6.4. $\frac{CC^{AM}_{core}(f,p)}{CC(f)}$ *cannot be bounded from above by any number α (even when α is allowed to be a function of the input).*

PROOF. It is known that for auctions with (gross) substitutes valuations, the Ausubel-Milgrom core prices coincide with VCG prices. We will show a trivial example where the optimal allocation can be computed with out any communication (i.e., known in advance), but computing VCG prices requires positive amount of communication. Note that this does not prove the same claim for finding equilibrium-supporting prices, since if the outcome is known in advance, any set of prices is incentive compatible. The example: consider an auction for one item among two bidders 1,2. $V_1 = \{10\}$, $V_2 = \{0, 1\}$. Bidder 1 clearly wins, but the VCG prices need to disclose v_2. □

7. REFERENCES

[1] L. M. Ausubel and P. R. Milgrom. Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1:1–42, 2002.

[2] Lawrence M. Ausubel and Paul Milgrom. *In P. Cramton and Y. Shoham and R. Steinberg (Editors), Combinatorial Auctions. Chapter 1. The Lovely but Lonely Vickrey Auction.* MIT Press., 2006.

[3] Sushil Bikhchandani, Rakesh Vohra, Sven de Vries, and James Schummer. Linear programming and Vickrey auctions. *Mathematics of the Internet: E-Auction and Markets*, pages 75–115, 2002.

[4] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 62–71, 2006.

[5] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions I: demand queries. Discussion paper no. 381, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC'05 contained preliminary results.

[6] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions II: ascending auctions. Discussion paper no. 382, The Center for the Study of Rationality, The Hebrew University., 2005. An extended abstract in EC'05 contained preliminary results.

[7] Yuval Emek, David Peleg, and Liam Roditty. A near-linear time algorithm for computing replacement paths in planar directed graph. In *SODA'08*, 2008.

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

[9] Eyal Kushilevitz and Noam Nisan. *Communication Complexity.* Cambridge University Press, 1997.

[10] Noam Nisan. Introduction to Mechanism Design (for Computer Scientists) In Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors), Algorithmic Game Theory. Chapter 9.

[11] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in STOC 1999.

[12] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192 – 224, 2006.

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

[14] Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136:341–378, 2007.

[15] Ilya Segal and Ronald Fadel. The communication cost of selfishness, 2006. Journal of Ecnomic Theory, to appear.

[16] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *ACM Symposium on Theory of Computing*, pages 209–213, 1979.

APPENDIX

A. BACKGROUND

A.1 The fooling sets techniques

Suppose that some communication protocol that computes a social function f for a two-player setting cannot distinguish between (v_1, v_2) and (v_1', v_2'), then, that protocol cannot also tell the difference between these inputs and (v_1', v_2), (v_1, v_2'). This suggests a way for proving lower bounds on the communication complexity of social-choice functions: Find a subset of the inputs and prove that every member in this subset is assigned the same outcome, but a combination 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 the logarithmic value of the cardinality of this subset 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 [9]) is the following:

THEOREM A.1. **[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 $CC(f) \geq log(|V'|)$.

B. AN IMPROVED UPPER BOUND FOR SINGLE ITEM AUCTIONS WITH TWO PLAYERS: THE PROOF

In this section we prove Theorem 3.2.

THEOREM B.1. *For the social-choice function* SINGLE-ITEM-AUCTION$_k$ *and $n = 2$, the information overhead of incentive compatibility is at most $1 + O(\frac{\log(k)}{k})$.*

PROOF. By Observation 1 proved in the proof for Proposition 3.1, $CC(f_k) \geq k$. The theorem is a direct result of the following communication-complexity lemma. In the terms of our model, it shows that for $n = 2$, $CC(\text{SINGLE-ITEM-AUCTION}_k, p) \leq k + O(\log(k))$ (recall that in single-item auctions, an equilibrium-supporting price is well known to be the second-highest price).

LEMMA 1. *Consider two players 1 and 2, each holds a number represented by k bits v_1, v_2 (resp.). Both players can realize $\min\{v_1, v_2\}$ with communication of at most $k + O(\log(k))$ bits.*

PROOF. We present a protocol that computes f_k and incentive-compatible payments, and requires the transmission of at most $k + O(\log(k))$ bits. It is well known that if the winning bidder is charged the value of the losing bidder then incentive compatibility is guaranteed (this is the celebrated "second-price auction"). So, computing f_k and incentive-compatible payments can be done by finding the identity and value of the bidder with the lowest value.

We consider the following communication protocol: Let $m = \lceil \log(k) \rceil$. The two parties (bidders) encode their inputs (values) using an alphabet of size $2^m - 1$. The two parties now take turns sending the symbols ("blocks") of their input values (from most significant to least significant). The alphabet symbols are encoded using m bits and are represented by the strings 0^m through $1^{m-1}0$. The string 1^m is assigned a reserved meaning: "your previous string was not equal to mine". If one party sends the 1^m string, then it also sends an additional bit indicating who has the larger value, and that party then sends all of its remaining symbols (encoded normally).

What is the maximal cost of this protocol in bits? We are wasting at most two blocks and one more bit when one party discovers the differing blocks (the most wasteful possibility is if the non-sending party realizes that its input is smaller, in which case he needs to send the string with the reserved meaning, an additional bit indicating that he has the smaller value, and repeat the last symbol). The cost is therefore at most $m \times R + 2m + 1$, where R is the length of the k-bit input encoded in the alphabet of size $2^m - 1$. That is, the number of symbols (blocks) sent times the number of bits needed to represent each symbol, plus the number of bits in two additional blocks, and another bit. The size of R must be such that $(2^m - 1)^R \geq 2^k$, that is, we can set R to be $\lceil k / \log(2^m - 1) \rceil$. Hence, the total cost in bits (setting $m = \lceil \log(k) \rceil$) is at most $\lceil \log(k) \rceil \times \lceil k / \log(k-1) \rceil + 2\lceil \log(k) \rceil + O(1)$, which is $k + O(\log(k))$. □

The theorem follows.

C. A CONSTANT UPPER BOUND FOR PUBLIC GOODS: THE PROOF

In this section we prove Theorem 4.2.

THEOREM C.1. *Fix $\epsilon > 0$. For any C, the informational overhead 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 $2 \cdot \frac{n}{n-1} + \epsilon \leq 3 + \epsilon$.*

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 in $C = 0$ or $C \geq n$ then the answer is trivial. So, we can consider the case $C \in (0, n)$. For any integer k' that is large enough, we can pick k such that $C \cdot 2^k \in [2^{k'-1}, 2^{k'}]$. We now consider the function f_k and normalize the input by 2^k, that is, we think about the input as if each agent i has integer value $v_i \in \{0, ..., 2^k - 1\}$ and the cost C is $\lceil C \cdot 2^k \rceil \in [2^{k'-1}, 2^{k'}]$. (note that as the values are now integers we can round up $C \cdot 2^k$ while being left with the same decision problem). The communication complexity of the original problem is clearly equivalent to the communication complexity of the new problem after normalization.

It is easy to see that $CC(f_k, p)$ is at most $nk' + n$, because of the following protocol: 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' bits per player, i.e., at most $n \times k'$ bits). An easy observation is that this simple protocol provides us with sufficient information to calculate the prices for all players. We shall now prove a lower bound on $CC(f_k)$ that will imply the theorem.

We prove the following lemmas:

LEMMA C.2. *If $C \leq \frac{n}{2}$, then for large enough k' the communication complexity of determining whether $\Sigma_i v_i \geq C$ is at least $(\frac{n}{2}-1)\cdot k' - f(n)$ for some function $f(\cdot)$.*

PROOF. We shall construct a large fooling set and invoke Theorem A.1. 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 place any two such type-profiles in the same monochromatic rectangle (see [9]) 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 in the same rectangle 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 monochromatic rectangles of any protocol that computes C-PUBLIC-GOOD$_k$. This implies that $\log L$ is a lower bound on the number of bits transmitted by any such protocol. We reach L as follows: First, consider the case that $C \leq 2^k - 1$. We consider the following family of type-profiles: $v_1 = C - 2^{k'-1}, v_2 = \ldots = v_{\frac{n}{2}-1} = 0$, and $\Sigma_{i=\frac{n}{2}}^{n} v_i = 2^{k'-1}$. 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 $\binom{2^{k'-1}+\frac{n}{2}}{\frac{n}{2}}$ ways to distribute $2^{k'-1}$ between $\frac{n}{2}+1$ players. This is bounded from below by $\frac{2^{(k'-1)\frac{n}{2}}}{(\frac{n}{2})^{\frac{n}{2}}}$. So, any protocol that determines whether $\Sigma_i v_i \geq C$ needs to transmit at least $\log(\frac{2^{(k'-1)\frac{n}{2}}}{(\frac{n}{2})^{\frac{n}{2}}}) = \frac{n}{2}k' - \frac{n}{2}(\log(n))$ bits.

What if $C \geq 2^k - 1$? Recall that $C \leq \frac{n}{2}\cdot 2^k$ (after the normalization) so in this case observe that for large enough k we can distribute $C - (2^k - 1)$ to players $1,\ldots,\frac{n}{2}$ (in some arbitrary way). This is so as these $\frac{n}{2}$ agent can split up to $(2^k-1)\cdot\frac{n}{2}$, 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 $\frac{n}{2}+1,\ldots,n$. We can now achieve L by looking at the different ways to distribute C' between players $\frac{n}{2}+1,\ldots,n$, i.e., such that $\Sigma_{i=\frac{n}{2}+1}^{n} v_i = C'$. How many such type-profiles are there? There are $\binom{C'+\frac{n}{2}-1}{\frac{n}{2}-1}$ ways to distribute C' between $\frac{n}{2}$ players. Now $\binom{C'+\frac{n}{2}-1}{\frac{n}{2}-1} = \binom{(2^k-1)+(\frac{n}{2}-1)}{\frac{n}{2}-1} \geq \frac{(2^k-1)^{(\frac{n}{2}-1)}}{(\frac{n}{2}-1)^{(\frac{n}{2}-1)}}$

So, any protocol that determines whether $\Sigma_i v_i \geq C$ needs to transmit at least a logarithmic factor of this number, which is $(\frac{n}{2}-1)\,log(2^k-1) - (\frac{n}{2}-1)\log(\frac{n}{2}-1)$. Note that $\frac{n}{2}\cdot 2^k \geq C \geq 2^{k'-1}$ thus $2^k \geq \frac{2^{k'}}{n}$. We conclude that this number is at least $(\frac{n}{2}-1)\log(\frac{2^{k'}}{n}-1) - (\frac{n}{2}-1)(\log(n)-2) \geq (\frac{n}{2}-1)k' - (\frac{n}{2}-1)(2\log(n)-1)$

□

LEMMA C.3. *If $C \leq \frac{n}{2}$, then for large enough k' the communication complexity of determining whether $\Sigma_i v_i \leq C$ is at least $(\frac{n}{2}-1)\cdot k' - f(n)$ for some function $f(\cdot)$.*

PROOF. The proof of this lemma is almost identical to that of the previous one (it involves the construction of the very same fooling set). □

The theorem will now follow because if $C \geq \frac{n}{2}$ in the original formulation then Lemma C.2 implies that $CC(f,p)$ is at least $(\frac{n}{2}-1)\cdot k' - f(n)$ for some function $f(\cdot)$. If $C > \frac{n}{2}$ then the problem is equivalent to figuring out whether $n - \Sigma_i v_i \geq n - C$. That is, it 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 $n - C$. This problem is just as hard as the original problem, as shown by Lemma C.3. By plugging in the values of $CC(f_k,p)$ and $CC(f,p)$ we get an overhead of at most $\frac{nk'+n}{(\frac{n}{2}-1)\cdot k' - f(n)}$, which converges to $2\frac{n}{n-1}$ as k' goes to infinity.