

The Communication Cost of Selfishness: Ex Post Implementation*

Ronald Fadel[†] Ilya Segal[‡]

May 9, 2005

Abstract

We consider the communication complexity of implementing a given decision rule when the protocol must also calculate payments to motivate the agents to be honest in an ex post equilibrium (agents' payoffs are assumed to be quasi-linear in such payments). We find that the communication cost of selfishness when measured with the average-case communication complexity may be arbitrarily large. For the worst-case communication complexity measure, we provide an exponential upper bound on the communication cost of selfishness. Whether this bound is ever achieved remains an open question. We examine several special cases in which the communication cost of selfishness proves to be very low. These include cases where we want to implement efficiency or where we have only two agents, and the precision of agents' utilities is fixed.

1 Introduction

This paper straddles two literatures on allocation mechanisms. One literature, known as “mechanism design,” examines the agents' incentives, and uses the “Revelation Principle” to focus on mechanisms in which agents fully reveal their preferences (e.g., [9, Chapter 23]). However, full revelation of private information would be prohibitively costly in most practical settings. The other literature examines how much communication, measured with the number of bits or real variables, is required in order to achieve the social goals, assuming that agents communicate truthfully (e.g., [5],[10], and references therein). However, in most practical settings we should expect agents to communicate strategically to maximize their private benefit.

This paper considers how much communication is required in order to implement a given decision function when agents are selfish.¹ Thus, we consider communication mechanisms that reveal enough information to calculate not only the allocation but also payments to the agents that motivate them to send honest reports in equilibrium (agents' payoffs are assumed to be quasi-linear in such payments). In this extended abstract we focus on the equilibrium concept of Ex Post (Nash) Incentive-Compatibility (EPIC for short) (The full paper [?] also presents results on the weaker concept of Bayesian-Nash Incentive-Compatibility.) Our results shed some light on the “communication cost of selfishness,” i.e., the additional communication complexity needed to calculate such

*The authors acknowledge the support of the National Science Foundation (grants SES 0318447, ITR 0427770).

[†]Department of Computer Science, Stanford University, Stanford, CA 94305, email rfadel@stanford.edu.

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

¹For simplicity here we restrict attention to decision *functions*, but see [2] for the problem of implementing a decision *correspondence* (relation), specifying a set of decisions any of which could be implemented in a given state.

payments, which we measure as the number of bits exchanged in either the average case or the worst case.

First we observe that a protocol computing a decision function that is EPIC-implementable need not reveal enough information to construct the transfers making it an EPIC mechanism, and so additional communication may be needed to ensure EPIC. In fact, the communication cost of selfishness may be quite large: we show that the average-case communication complexity of allocating one object efficiently between two agents whose valuations are independently drawn from a uniform distributions is at most 4 bits, but the average-case communication complexity of any EPIC mechanism implementing such allocation is infinite. We proceed to examine the worst-case communication complexity measure, and show that for this measure the cost of selfishness is at most exponential in the communication complexity. Whether this upper bound is ever achieved remains an open question. We proceed to examine several special cases in which the communication cost of selfishness proves to be very low. These include the cases when we want to implement efficiency or when we have only two agents, and the precision of agents' utilities is fixed (or we are content with approximate incentive compatibility).

A number of papers have examined incentive-compatible indirect communication mechanisms in various special settings. The first paper we know of is Reichelstein [11], who considered incentive compatibility in nondeterministic real-valued mechanisms, and showed that the communication cost of selfishness in achieving efficiency is low (we obtain a parallel result for the communication cost in bits). Lahaie and Parkes [7] characterized the communication problem of finding Vickrey-Groves-Clarke (VCG) transfers as that of finding a “universal price equilibrium,” but did not examine the communication complexity of finding such transfers, or the possibility of implementing efficiency using other transfers. Neither paper examined the communication complexity of allocation rules other than surplus maximization. Also, see Roughgarden and Tardos [12] for results on the latency cost of “selfish routing” in networks, and see Feigenbaum et al. [3] for some analysis of the communication requirements of distributed incentive-compatible multicast cost sharing mechanisms.

2 Background: Communication Complexity

The concept of *communication complexity*, introduced by Yao [13] and surveyed in [5], describes how much communication is needed for several agents (from a set I) to compute the value of a function $f : \prod_{i \in I} \Theta_i \rightarrow X$ when each agent $i \in I$ knows privately a part $\theta_i \in \Theta_i$ of the “state” $(\theta_1, \dots, \theta_I)$ (we will refer to θ_i as agent i 's “type”). Communication is modeled using the notion of a protocol. In the language of game theory, a protocol is described with an extensive-form game form and the agents' strategies in it. We restrict attention to games of perfect information (i.e., each decision is broadcast to all agents), which is without loss of generality for our purposes. Also, we restrict attention to protocols in which each agent has two possible moves at a decision node (interpreted as sending one bit), since any message from a finite alphabet can be coded with a fixed number of bits. Formally,

Definition 1 *A protocol \mathcal{P} with agents $I = \{1, \dots, I\}$ over state space $\Theta = \prod_{i \in I} \Theta_i$ and decision space X is a binary tree, with a set of nodes N and a set of leaves $L \subseteq N$, where:*

- *The set $N \setminus L$ of non-leaf nodes is partitioned into I subsets N_1, \dots, N_I , with N_i representing the set of decision nodes of agent $i \in I$.*
- *Each leaf $l \in L$ of the tree is labeled with a decision $o(l) \in X$*

- Each agent $i \in N$ has a strategy plan $\sigma_i : \Theta_i \rightarrow S_i$, where $S_i = \{0, 1\}^{N_i}$ is the set of the agent's possible strategies in the game (moves made at his decision nodes).²

For each strategy profile $s = (s_1, \dots, s_I) \in \prod_{i \in I} S_i$, let $\lambda(s) \in L$ denote the leaf reached when the agents follow the respective strategies. The decision function $f : \Theta \rightarrow X$ computed by protocol \mathcal{P} , which denoted by $\text{Fun}(\mathcal{P})$, is defined by $f(\theta) = o(\lambda(\sigma(\theta)))$ for all $\theta \in \Theta$.

The communication complexity of a protocol is defined as the number of bits sent in it, which could be measured in the worst case (as the maximum depth of the tree), or in the average case given a probability distribution over the states:

Definition 2 For $\theta \in \Theta$, let $d_{\mathcal{P}}(\theta)$ be the depth of leaf $\lambda(\theta)$ in protocol \mathcal{P} , i.e., the number of edges between the root and $\lambda(\theta)$.

- The worst-case communication cost of a protocol \mathcal{P} is defined as the maximum depth of the protocol i.e., $d_{\mathcal{P}} = \max_{\theta \in \Theta} d_{\mathcal{P}}(\theta)$. The (worst-case) communication complexity of decision function $f : \Theta \rightarrow X$ is defined as $\min_{\mathcal{P}: \text{Fun}(\mathcal{P})=f} d_{\mathcal{P}}$.
- The average-case communication cost of a protocol \mathcal{P} given a probability distribution p on Θ is defined as $\text{ACC}_p(\mathcal{P}) = E_p[d_{\mathcal{P}}(\theta)]$. The average-case communication complexity of decision function $f : \Theta \rightarrow X$ given distribution p is defined as $\min_{\mathcal{P}: \text{Fun}(\mathcal{P})=f} \text{ACC}_p(\mathcal{P})$.

3 Binary Dynamic Mechanisms

3.1 The Formalism

A protocol induces an extensive-form game with perfect information, and prescribes a strategy in this game for each agent of any type. When agents are selfish, we need to consider their incentives to deviate from the prescribed strategies. We will say that an agent i of type θ_i is *honest* if he follows the prescribed strategy $\sigma_i(\theta_i)$. An agent's incentive to be honest are affected by monetary payments assigned to the agents, which the protocol can compute along with the decision:

Definition 3 A binary dynamic mechanism (BDM) is a pair $\langle \mathcal{P}, \tau \rangle$ such that:

- \mathcal{P} is a protocol with set of leaves L .
- $\tau : L \rightarrow \mathbb{R}^I$ describes a profile of payments given to the agents in all leaves..

The function $t : \Theta_i \rightarrow \mathbb{R}^I$ defined by $t(\theta) = \tau(\lambda(\sigma(\theta)))$ is called the transfer function computed by the BDM.

We describe agents' preferences by a Utility Function Profile (UFP) $U = (u_1, \dots, u_I)$, where for all $i \in I$, the utility function $u_i : X \times \Theta_i \rightarrow \mathbb{R}$ gives the utility of agent i of each type for each decision. We assume the utilities are quasi-linear in the payments, i.e., the total payoff of agent i of type θ_i from decision x and transfer t_i is $u_i(x, \theta_i) + t_i$.

²It more customary in game theory to call the "strategy" of agent i the whole function σ_i , contingent of the agent's type $\theta_i \in \Theta_i$, which is interpreted as a "move of nature." For our purposes, however, it is convenient to reserve the term "strategy" to denote the agent's behavior $s_i \in S_i$ in the protocol.

3.2 Ex Post Incentive Compatibility

Definition 4 *BDM $\langle \mathcal{P}, \tau \rangle$ is Ex Post Incentive Compatible (EPIC) with respect to utility function profile $U = (u_1, \dots, u_I)$ if in any state $\theta \in \Theta$, the strategy profile $s = (\sigma_1(\theta_1), \dots, \sigma_I(\theta_I)) \in \prod_{i \in I} S_i$ is an ex post Nash Equilibrium of the induced game, i.e.,*

$$\forall i \in I, \forall s'_i \in S_i : u_i(o(\lambda(s)), \theta_i) + \tau_i(\lambda(s)) \geq u_i(o(\lambda(s'_i, s_{-i})), \theta_i) + \tau_i(\lambda(s'_i, s_{-i})).$$

In words, for any state $\theta \in \Theta$, each agent's prescribed strategy is optimal for him as long as long he expects other agents to follow their prescribed strategies.³ In this case, we will say that BDM $\langle \mathcal{P}, \tau \rangle$ *implements* $f = \text{Fun}(\mathcal{P})$ in EPIC with respect to U .

Note that:

- By the Revelation Principle, if decision function f is implementable in an EPIC BDM that computes a transfer function $t : \Theta \rightarrow \mathbb{R}^I$, the following Dominant-strategy Incentive-Compatibility (DIC) constraints must be satisfied:

$$u_i(f(\theta_i, \theta_{-i}), \theta_i) + t_i(\theta_i, \theta_{-i}) \geq u_i(f(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i}) \quad \forall \theta_i, \theta'_i \in \Theta_i \quad \forall \theta_{-i} \in \Theta_{-i}. \quad (1)$$

In particular, for $\theta_i, \theta'_i \in \Theta_i$ such that $f(\theta_i, \theta_{-i}) = f(\theta'_i, \theta_{-i})$, these inequalities imply that $t_i(\theta_i, \theta_{-i}) = t_i(\theta'_i, \theta_{-i})$, and therefore the transfer to each agent i can be written in the form

$$t_i(\theta) = T_i(f(\theta), \theta_{-i}) \quad \text{for some } T_i : X \times \Theta_{-i} \rightarrow \mathbb{R}. \quad (2)$$

- Our restriction to mechanisms of perfect information is without loss of generality for EPIC implementation, since ex post Nash equilibrium requires honesty to be optimal even for an agent who knows the other agents' types and therefore all their moves.
- This setting can be extended to the case where types are revealed to agents in real time as the mechanism is executed, as long as each agent has enough information at each node to compute the prescribed move.

3.3 Incentivability of Protocols

In standard mechanism design, according to the Revelation Principle, a decision function is DIC-implementable if and only if the direct revelation protocol for this rule can be incentivized with some transfers. Now we define incentivability for general protocols:

Definition 5 *A protocol \mathcal{P} with set I of agents and set of leaves L is EPIC-incentivable (or just incentivable) with respect to UFP $U = (u_1, \dots, u_I)$ if there is a transfer function $\tau : L \rightarrow \mathbb{R}^I$ such that $\langle \mathcal{P}, \tau \rangle$ is an EPIC BDM with respect to U .*

There exist simple examples of incentivable protocols other than direct revelation protocols. For example, if decision function f is DIC implementable with transfer functions t that depend only on

³In our setting of *private values*, in which agent i 's utility does not depend on others' types θ_{-i} , this is equivalent to saying that agent i 's strategy be optimal for him assuming that each agent $j \neq i$ follows a strategy prescribed for *some* type θ_j . Note that we do not require the stronger requirement of Dominant-strategy Incentive-Compatibility (DIC), which would allow agent i to expect agents $j \neq i$ to use contingent strategies $s_j \in S_j$ that are inconsistent with any type θ_j , and which would be violated in even the simplest dynamic mechanisms.

the decision (i.e., take the form $t_i(\theta) = T_i(f(\theta))$), then any protocol computing f is incentivable.⁴ However, for other DIC implementable decision functions, and in particular for efficient decision functions, it is easy to find examples of protocols that are not incentivable:

Example 1 *There are two agents and one indivisible object, which can be allocated to either agent. The two agents' valuations (utilities from receiving the object) lie in type spaces $\Theta_1 = \{1, 2, 3, 4\}$ and $\Theta_2 = [0, 5]$ respectively (their utilities for not receiving the object are normalized to zero). An efficient decision function (which allocates the object to the agent with the higher valuation) can be computed with the following protocol \mathcal{P}_0 : Agent 1 sends his type (using $\log_2 4 = 2$ bits), and then agent 2 outputs an efficient allocation $x \in \{1, 2\}$ (using 1 bit). The computed decision function is DIC implementable. Now, suppose in negation that the protocol computes a transfer function t that make it into an EPIC BDM. Using (2) and the fact that the protocol does not reveal anything about agent 2's type θ_2 beyond the decision $f(\theta_1, \theta_2)$, the transfer to agent 1 must take the form $t_1(\theta_1, \theta_2) = T_1(f(\theta_1, \theta_2))$. But then agent 1 would always want to take the object when $\theta_1 > T_1(1) - T_1(2)$, and would always choose not to take it when $\theta_1 < T_1(1) - T_1(2)$, so it cannot compute an efficient decision function. Hence, Protocol \mathcal{P}_0 is not incentivable.*

3.4 Incentive Communication Complexity

Since the cheapest protocol computing a decision function may not be incentivable, the agents' selfishness may raise the communication cost:

Definition 6 $ICC^U(f)$, the incentive communication complexity of a decision function f with respect to the UFP U , is the depth $d_{\mathcal{P}}$ of the shallowest protocol \mathcal{P} that computes f and that is incentivable with respect to U .

$AICC_p^U(f)$, the average incentive communication complexity of a decision function f with respect to the UFP U and probability distribution p , is the minimal average communication cost $ACC_p(\mathcal{P})$ over all protocols \mathcal{P} that compute f and that are incentivable with respect to U .

We can now talk of the communication cost of selfishness as the difference between $CC(f)$ and $ICC^U(f)$ when we consider the worst-case costs, or the difference between $ACC_p(f)$ and $AICC_p^U(f)$ when we consider the average-case costs. Note that the inequalities $ACC_p(f) \leq AICC_p^U(f) \leq ICC^U(f)$ and $ACC_p(f) \leq CC(f) \leq ICC^U(f)$ must always hold.

4 Infinite Average-Case Cost of Selfishness

We show that the average-case communication cost of selfishness can be unbounded, even for the problem of computing an efficient allocation with two agents:

Proposition 1 *For any $\alpha > 0$ there exists a problem with two agents, a state space $\Theta = \Theta_1 \times \Theta_2$, a UFP $U = (u_1, u_2)$, and an efficient decision function f such that, given the uniform probability distribution p over Θ , $ACC_p(f) < 4$ but $AICC_p^U(f) > \alpha$.*

⁴For example, this property is satisfied by decision functions that depend only on a single agent's type (this is known as the "Taxation Principle"). To have another example satisfying this property, suppose there is one object and two agents, and consider the decision function that gives the object goes to agent 1 if his valuation is more than a , otherwise, it gives the object to agent 2 if his valuation is more than a , and does not give it to either agent if neither condition is satisfied.

Proof Sketch: Consider the problem of allocating an indivisible object to one of the two agents, as in Example 1 above, but with the agents' valuations drawn independently from the uniform distribution over $\Theta_1 = \Theta_2 = \{k2^{-\gamma} | k = 0, \dots, 2^\gamma - 1\}$. Let f be the efficient decision function that allocates the object to the agent with the higher valuation, and gives it to agent 1 in the case of a tie. f can be computed with the following bisection protocol suggested in [4]: At each round $m = 1, \dots, \gamma$, each agent i reports the m th bit in the binary expansion of his valuation θ_i . The protocol stops as soon as the two agents report different bits, and then gives the object to the agent who just reported 1 (he is proven to have the higher valuation). If the agents have not disagreed in the γ steps, the object is given to agent 1 (in this case the two agents are shown to have the same valuations). The probability that the protocol stops at any given round conditional on arriving there is $1/2$. Therefore, the expected number of rounds is at most 2, and so the average-case communication complexity is at most 4, regardless of γ .

Now, consider an EPIC BDM implementing decision function f . (In fact, the argument below applies to any efficient decision function). By (2), the transfer to agent 2 can be written as $T_2(f(\theta_1, \theta_2), \theta_1)$. Furthermore, the DIC inequalities (1) imply that to implement an efficient decision rule f , we must have

$$|T_2(2, \theta_1) - T_2(1, \theta_1) - \theta_1| \leq 2^{-\gamma} \text{ for any } \theta_1 \in (0, 1 - 2^{-\gamma}),$$

for otherwise agent 2 would prefer either to overstate his valuation when $\theta_1 = \theta_2 - 2^{-\gamma}$ or to understate it when $\theta_1 = \theta_2 + 2^{-\gamma}$.

Suppose that $\gamma \geq 4$. Let us now run the EPIC BDM twice, drawing θ_2 independently for the two runs but using the same θ_1 for the two runs. In the event where $\theta_2 \geq 3/4$ in the first run, $\theta_2 < 1/4$ in the second run, and $\theta_1 \in \tilde{\Theta}_1 \equiv \{k \cdot 2^{2-\gamma} : k = 1/4 \cdot 2^{\gamma-2}, \dots, 3/4 \cdot 2^{\gamma-2} - 1\}$, the object goes to agent 2 in the first run and to agent 1 in the second run, and the difference between agent 1's transfers computed in the two runs pins down the realization of $\theta_1 \in \tilde{\Theta}_1$ (it must be within $2^{-\gamma}$ of the difference). Thus, in this event, whose probability is $1/16$, the two runs of the protocol should generate $|\tilde{\Theta}_1| = 2^{\gamma-3}$ distinct equiprobable leaves. This implies that at least half of these leaves must have depth at least $\gamma - 3$. Since the leaves are equiprobable, with probability $1/2 \cdot 1/16 = 1/32$ the two-run protocol sends at least $\gamma - 3$ bits, hence the expected number of bits transmitted in the two-run protocol is at least $(\gamma - 3)/32$. The average-case communication complexity of a single run of the EPIC BDM is then at least half this number, i.e., $(\gamma - 3)/64$.

Although infinite protocols are not formally defined here, we can use a bisection protocol [4] to allocate an object efficiently between two agents whose valuations are uniformly distributed on $[0, 1]$ using only 4 bits on average. However, no protocol with a finite average-case communication cost can be incentivized in this case.

5 Exponential Worst-Case Upper Bound

Now we turn to the communication cost of selfishness for the worst-case complexity measure. We show that this cost can be bounded above by an exponential function of the original communication complexity of the decision function.

By the Revelation Principle, given a DIC-implementable decision function, a full revelation of the agents' types, would allow us to compute a transfer function that makes the DIC inequalities (1) hold. But is full revelation necessary to compute such "acceptable" transfers? For example, if some of the agents' type spaces are infinite but our decision function can be computed with a finite

protocol, is it possible that we need infinite communication to compute acceptable transfers? We show that the answer is “no:” In fact, given a protocol \mathcal{P} that computes a DIC-implementable decision function, we can compute some acceptable transfers using only enough information about the types to know what each agent would do at any node of \mathcal{P} .

Formally, we define the *Normal-Form Expansion* of protocol \mathcal{P} , denoted by $\text{NFE}(\mathcal{P})$, as a protocol in which each agent i describes his strategy $\sigma_i(\theta_i)$ in \mathcal{P} , i.e., the decisions $\sigma_i(\theta_i)(n)$ prescribed by the strategy at each of the agents’ decision nodes $n \in N_i$. (The order in which the agents describe their strategies in $\text{NFE}(\mathcal{P})$ is irrelevant.) Note that the depth of $\text{NFE}(\mathcal{P})$ equals to the number of decision nodes in \mathcal{P} , which is bounded above by $2^{d_{\mathcal{P}}} - 1$. We show that if protocol \mathcal{P} computes a DIC implementable decision function, then $\text{NFE}(\mathcal{P})$ is EPIC incentivable, which yields the following upper bound on Incentive Communication Complexity:

Proposition 2 *For any UFP $U = (u_1, \dots, u_I)$ and any DIC-implementable decision function f ,*

$$ICC^U(f) \leq 2^{CC(f)} - 1$$

Proof Sketch: Suppose that protocol \mathcal{P} computes a decision function f that is DIC implementable with some transfer rule $t : \Theta \rightarrow \mathbb{R}^I$. For each agent i , consider a reduced state space $\Theta'_i \subseteq \Theta$ such that any strategy s_i in \mathcal{P} is used by at most one type in Θ'_i . Note that if $\sigma_i(\theta'_i) = \sigma_i(\theta_i)$ for $\theta'_i, \theta_i \in \Theta_{-i}$, then the decision function computed by the protocol must have $f(\theta'_i, \theta_{-i}) = f(\theta_i, \theta_{-i})$ for all $\theta_{-i} \in \Theta_{-i}$. To each leaf $s = (s_1, \dots, s_I)$ of protocol $\text{NFE}(\mathcal{P})$, assign the transfers $t(\theta') \in \mathbb{R}^I$ corresponding to the state $\theta' \in \Theta'$ for which $\sigma(\theta') = s$. (Announcement of a strategy s_i that is not used by any type in \mathcal{P} should be banned or, equivalently, punished with a large enough negative transfer so that no type would want to do it.) DIC inequalities (1) imply that such transfers incentivize $\text{NFE}(\mathcal{P})$.

This upper bound shows that the communication cost of selfishness is not unbounded, and is at most exponential. In particular, it shows that a decision function f can be implemented in an EPIC BDM if and only if f is a DIC-implementable decision function that can be computed with finite communication.

The upper bound of Proposition 2 can be improved by eliminating from $\text{NFE}(\mathcal{P})$ those strategies in \mathcal{P} that are not used by any type:

Example 2 *Consider the setting in Ex. 1. Protocol \mathcal{P}_0 has depth 3, so by Prop. 2, there exists an incentivable protocol of depth $2^3 - 1 = 7$. But we can go further: agent 1 uses only 4 strategies in \mathcal{P}_0 (one for each of his types) out of the $2^3 = 8$ possible strategies, and agent 2 uses only 5 strategies out of $2^4 = 16$, each of the 5 strategies being described by a threshold of agent 1’s announcement below which agent 2 takes the good. Since full description of such strategies in \mathcal{P}_0 by the two agents takes $\lceil \log_2 4 \rceil + \lceil \log_2 5 \rceil = 5$ bits, we have a protocol of depth 5 that is incentivable.*

It is unknown, however, whether there exists a UFP U and a DIC-implementable decision function f such that $ICC^U(f)$ is even close to the obtained upper bound of $2^{CC(f)} - 1$. In particular, it remains an open problem to determine the highest attainable upper bound, or to find any “canonically hard” instances, in the spirit of [3].

Open Problem 1 *Are there “canonically hard” DIC-implementable decision functions f combined with utility function profiles U , for which the incentive communication complexity, $ICC^U(f)$, is much higher than the communication complexity of f , $CC(f)$? How high can be the (worst-case) communication cost of selfishness?*

6 Cases with Low Cost of Selfishness

While we do not expect the communication cost of selfishness to be low in general, in this section we identify some cases where it proves to be reasonable. Some of these cases involve decision functions that are f efficient, i.e. choose decisions that maximize the sum of utilities. Since the communication complexity of computing exact efficiency may in itself be infinite or prohibitive, we also allow approximately efficient functions:

Definition 7 *Decision function $f : \Theta \rightarrow X$ is δ -efficient for UFP $U = (u_1, \dots, u_I)$ (for some $\delta \geq 0$) if $\forall \theta \in \Theta$,*

$$\sum_{i \in I} u_i(f(\theta), \theta_i) \geq \max_{x \in X} \sum_{i \in I} u_i(x, \theta_i) - \delta \quad \forall \theta \in \Theta.$$

The applicability of our results goes beyond efficient decision functions: Under some natural conditions [8], any DIC-implementable decision function maximizes a non-negatively weighted affine combination of the agents' utilities, which can be interpreted as efficiency upon rescaling of the utilities and adding a fictitious agent with a known utility function. We consider efficient and approximately efficient decision functions in general, and also for the special case of "single-parameter" agents, who have the same nonnegative utility for a known subset of decisions and a zero utility for the other decisions. In the general case, the communication cost of selfishness is bounded when agents' utilities are given with a finite precision (equivalently, we are content with approximate incentive compatibility), while for single-parameter agents we obtain a bound without any constraints on the utility range. Finally, we show that the communication cost of selfishness is very low for *any* DIC decision function if there are only two agents whose utilities are given with a finite precision (or if we are content with approximate incentive compatibility).

6.1 Efficiency with Discrete Utility Range

It is well known that any efficient decision function is DIC implementable. In particular, it can be implemented in dominant strategies by giving each agent a transfer equal to the sum of other agents' utilities from the computed decision (as in the VCG mechanism). Therefore, starting with any protocol computing an efficient decision function f , we can satisfy EPIC by having the agents report their utility values from the decision computed by the protocol, and then transfer to each agent the amount equal to the sum of the reported utilities of the other agents. (Note that an agent would have no incentive to misreport his utility value since this would have no effect on his own transfer.) This idea dates back to Reichelstein [11] and was recently used in [1]. To obtain a BDM, we need the agents to have finite utility ranges. For example, the following condition would suffice:

Definition 8 *UFP $U = (u_1, \dots, u_I)$ has discrete range with precision γ if $u_i(x, \theta_i) \in \{k2^{-\gamma} : k = 0, \dots, 2^\gamma - 1\}$ for all $\theta_i \in \Theta_i, x \in X$.*

In this case, each agent can report his utility value using γ bits, and so we obtain

Proposition 3 *For an UFP U with discrete precision- γ range, and an efficient decision function f ,*

$$\begin{aligned} ICC^U(f) &\leq CC(f) + I\gamma, \text{ and} \\ AICC_p^U(f) &\leq ACC_p(f) + I\gamma \text{ for any probability distribution } p \end{aligned}$$

Thus, the communication cost of selfishness is at most γI bits, both for average-case and worst-case communication.

Furthermore, even if utility range is continuous but bounded (without loss of generality lies in $[0, 1]$, up to a rescaling of utilities), we can use the same approach to achieve ε -EPIC (i.e., an agent will be honest unless he can get more than ε by deviating):

Definition 9 *BDM $\langle \mathcal{P}, \tau \rangle$ is ε -EPIC for some $\varepsilon \geq 0$ if in any state $\theta \in \Theta$, the prescribed strategy profile in this state is an ε -Nash equilibrium of the game, i.e., the inequalities in Def. 4 above are violated by at most ε .*

If each agent outputs his utility rounded down to a multiple of $\varepsilon/(I-1)$, we can construct a transfer to each agent that is within $1/2 \cdot \varepsilon/(I-1)$ of the sum of others' true utilities, and so ensure ε -EPIC. This will take at most $\lceil \log_2 (I-1)/\varepsilon \rceil$ bits by each agent, which bounds above the communication cost of selfishness. The argument also extends to δ -efficient decision functions with $0 \leq \delta < \varepsilon$: It is sufficient for each agent to output his utility rounded down to a multiple of $(\varepsilon - \delta)/(I-1)$, which takes at most $\lceil \log_2 (I-1)/(\varepsilon - \delta) \rceil$ bits. To summarize:

Proposition 4 *For any UFP U with utility range contained in $[0, 1]$, and decision function f is δ -efficient, and the solution concept is ε -EPIC with $0 \leq \delta < \varepsilon$ then the communication cost of selfishness is at most $I \lceil \log_2 (I-1)/(\varepsilon - \delta) \rceil$, for either the worst-case or average-case communication measure with any probability distribution.*

Thus, for efficient and approximately efficient decision functions, the communication cost of selfishness is bounded above by a number that does not depend on the communication complexity of the decision function, as long as the utility range is discrete or approximate EPIC is allowed.

6.2 Efficiency with Single-Parameters Agents

Here we restrict attention to agents who have the same nonnegative utility for a known subset of decisions and a zero utility for the other decisions:

Definition 10 *Agent i is a Single-Parameter (SP) agent if his type space is $\Theta_i \subset \mathbb{R}_+$ and there exists a set $X_i \subseteq X$ of decisions such that $u_i(x, \theta_i) = \theta_i$ if $x \in X_i$ and $u_i(x, \theta_i) = 0$ otherwise.*

For the problem of achieving efficiency or approximate efficiency with SP agents, we can bound above the communication cost of selfishness regardless of the utility range, by a factor of at most the number of agents:

Proposition 5 *Consider UFP U such that all I agents are single-parameter agents, and let f be a δ -efficient decision functions for some $\delta \geq 0$. Then*

$$ICC^U(f') \leq I \times (CC(f) + 1)$$

for some δ -efficient decision function f' .⁵

⁵For exact efficiency (i.e., $\delta = 0$), the proposition easily extends to agents whose utility functions take the form $u_i(x, \theta_i) = \theta_i a_i(x) + b_i(x)$, for some functions $a_i, b_i : X \rightarrow \mathbb{R}$.

Proof Sketch: Consider a protocol \mathcal{P} computing f . It is well known that the set of states in which a given leaf $l \in L$ of the protocol is reached is a product set $\prod_{i \in I} \Theta_i(l)$ [5, Chapter 1]. For each agent i , let $\underline{\theta}_i(l) = \inf \Theta_i(l)$ and $\bar{\theta}_i(l) = \sup \Theta_i(l)$. Given the utilities of SP agents, the decision $o(l)$ assigned to leaf l must be δ -efficient on $\prod_{i \in I} [\underline{\theta}_i(l), \bar{\theta}_i(l)]$. The threshold points $\underline{\theta}_i(l), \bar{\theta}_i(l)$ for $l \in L$ partition agent i 's type space Θ_i into at most $2|L|$ intervals. Consider now a protocol \mathcal{P}' in which each agent i reports the interval G_i in which his type lies. There is at least one decision known to be δ -efficient on the event $G = \prod_{i \in I} G_i$ (namely, the decision assigned by protocol \mathcal{P} to any leaf $l \in L$ for which $G_i \subset [\underline{\theta}_i(l), \bar{\theta}_i(l)]$ for all $i \in I$). Now, let protocol \mathcal{P}' assign a decision to leaf $G = \prod_{i \in I} G_i$ by trying, within all the decisions that are δ -efficient on G , lexicographically, to achieve $x_1 \in X_1$, then $x_2 \in X_2, \dots$, then $x_I \in X_I$. Let $f' = \text{fun}(\mathcal{P}')$.

Note that in protocol \mathcal{P}' , each agent sends at most $\log_2(2|L|) \leq 1 + d_{\mathcal{P}}$ bits, and it computes a δ -efficient decision function f' . It remains to show that \mathcal{P}' is incentivable. For this purpose, first observe that f' is DIC implementable. Indeed, by construction, f' has the property that $\forall \theta \in \Theta \forall \theta'_i \in \Theta_i$ such that $\theta'_i > \theta_i$, $f'(\theta) \in X_i \Rightarrow f'(\theta'_i, \theta_{-i}) \in X_i$. Therefore, f' is DIC implementable using transfers of the form (2) with $T_i(x, \theta_{-i}) = \inf \{\theta_i \in \Theta_i : f(\theta_i, \theta_{-i}) \in X_i\}$. Furthermore, each agent's strategy in protocol \mathcal{P}' is not contingent on others' messages, and so by the argument of Section 5, \mathcal{P}' is incentivable.

In particular, this proposition applies to the problem of allocating one indivisible object among I agents without externalities, in which case the agents are single-parameter agents.⁶ Note that the result does not extend to the average-case communication cost, which, as shown in Prop. 1, can be unbounded in this case.

6.3 Any DIC Decision Function with Two Agents

Recall from (2) that when decision function f is DIC implementable with transfers t , the transfer to agent i can be written as $t_i(\theta) = T_i(f(\theta), \theta_{-i})$. Furthermore, if the utilities have discrete range with precision γ , then we can restrict attention to discrete transfers with the same precision. With two agents, agent $-i$ can output the transfer $T_i(f(\theta), \theta_{-i})$ at the end of any protocol computing $f(\theta)$, thus yielding an EPIC BDM implementing f . This argument yields

Proposition 6 *Suppose that $I = 2$ and UFP U has discrete range with precision γ . Then for any DIC decision function f ,*

$$\begin{aligned} ICC^U(f) &\leq CC(f) + 2(\gamma + 1), \text{ and} \\ AICC_p^U(f) &\leq ACC_p(f) + 2(\gamma + 1) \text{ for any probability distribution } p \end{aligned}$$

Proof Sketch: We can fix $T_i(x_0, \theta_{-i}) = 0 \forall \theta_{-i}$ for an arbitrary fixed decision $x_0 \in X$, and then DIC inequalities (1) imply that $|T_i(x, \theta_{-i})| \leq 1 - 2^{-\gamma}$. Furthermore, since utilities have a discrete range with precision γ , we can round down all transfers to multiples of $2^{-\gamma}$ while preserving DIC. Reporting such a transfer takes $\gamma + 1$ bits.

⁶A related theorem for this case is stated in [6]: It shows that any sequential communication in this setting can be replaced with simultaneous communication achieving the same approximation of efficiency while increasing the communication complexity by at most a factor of I .

We can obtain a parallel result for a continuous but bounded utility range if ε -EPIC allowed, since we can round down the transfers to multiples of ε^{-1} :

Proposition 7 *Let $I = 2$ with UFP U with utility range contained in $[0, 1]$, decision function f be DIC implementable, and the solution concept be ε -EPIC with $\varepsilon > 0$, then the communication cost of selfishness is at most $2\lceil \log_2 \varepsilon^{-1} + 1 \rceil$, for either the worst-case or average-case communication measure with any probability distribution.*

7 Conclusion

This note has shown that the communication cost of selfishness when defined by ex post incentive-compatibility may be quite high—even infinite when measured with the average-case complexity. In the full paper, we also consider the weaker equilibrium concept of Bayesian Incentive-Compatibility (BIC). For this concept, we show that for DIC-implementable decision function, the communication cost of selfishness is zero. In fact, we show that any protocol computing a DIC decision function is BIC incentivable, i.e., it yields enough information to compute payments to agents that make it a BIC mechanism (and the payments can even be made to add up to zero, thus ensuring a balanced budget). The situation is different for decision functions that are not DIC-implementable but still BIC-implementable. To implement such decision functions, it proves important not to broadcast all the moves in the game to all the agents, i.e., to create some non-trivial information sets. Intuitively, if an agent knows too much, this may destroy his incentive to be honest. However, hiding previous moves may increase the communication complexity (e.g., single-round communication complexity is known to be sometimes exponentially higher than multi-round complexity [5]). Thus, we conjecture that BIC implementable decision functions may have a large communication cost of selfishness. For now, however, the only general formal result we have in this regard is an exponential upper bound for the worst-case communication cost of selfishness (proven along the same lines as Prop. 2).

References

- [1] S. Athey and I. Segal. An efficient dynamic mechanism, 2005. Forthcoming.
- [2] R. Fadel. Contributions to canonical hardness, 2005. Workshop on Agent-Mediated Electronic Commerce VII.
- [3] J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Hardness results for multicast cost sharing, 2002. Yale University Technical Report YALEU/DCS/TR1232, June 2002.
- [4] E. Grigorieva, P. J.-J. Herings, R. Muller, and D. Vermeulen. The private value single item bisection auction. Research Memoranda 051, Maastricht : METEOR, Maastricht Research School of Economics of Technology and Organization, 2002.
- [5] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [6] N. Nisan L. Blumrosen and I. Segal. Multi-player and multi-round auctions with severely bounded communication. ESA 2005.
- [7] S. Lahaie and D. Parkes. Applying learning algorithms to preference elicitation, 2004.

- [8] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.
- [9] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [10] N. Nisan and I. Segal. The communication requirements of efficient allocations and supporting prices, 2004. Forthcoming in the *Journal of Economic Theory*.
- [11] S. Reichelstein. Incentive compatibility and informational requirements. *Journal of Economic Theory*, 1984.
- [12] T. Roughgarden and E. Tardos. How bad is selfish routing? In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.
- [13] A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213. ACM Press, 1979.