Cost-sharing Models in Participatory Sensing

Georgios Birmpas[1], Costas Courcoubetis[1,2],
Ioannis Giotis[1], and Evangelos Markakis[1]

[1] Department of Informatics, Athens University of Economics and Business, Greece.
[2] Singapore University of Technology and Design, Singapore.

Abstract. In Smart City and Participatory Sensing initiatives the key concept is for user communities to contribute sensor information and form a body of knowledge that can be exploited by innovative applications and data analytics services. A key aspect in all such platforms is that sensor information is not free but comes at a cost. As a result, these platforms may suffer due to insufficient sensor information made publicly available if applications do not share efficiently the cost of the sensor information they consume.
We explore the design of specialized market mechanisms that match demand to supply while taking into account important positive demand externalities: sensors are digital goods and their cost can be shared by applications. We focus on the buyer side and define different demand models according to the flexibility in choosing sensor data for satisfying application needs. We then investigate the properties of various cost-sharing mechanisms with respect to efficiency and budget balance. In doing so, we also propose and study a new mechanism, which although lacks strategyproofness, it exhibits important efficiency improvement along with certain fairness properties.

1 Introduction

A remarkable incorporation of sensors has occurred in the last few years in a wide range of devices. Starting from the inclusion of GPS receivers, accelerometers and barometers in smartphones, lately we are seeing a wave of health related sensors being used in the form of fitness bands and smartwatches. Aside from personal devices, home automation and power management devices are distinctively on the rise and include a different variety of sensor data.

Such sensor information can potentially be collected in more precise detail and volume, opening up possibilities for research on unprecedented scales. Towards this, participatory sensing initiatives form a natural and promising approach, replacing traditional sensor networks, where user communities can contribute sensor information, that can later be exploited by innovative applications. There are already existing deployments and platforms that support a variety of applications like environmental monitoring (OpenSense), transportation (CrowdPark), fitness (BikeTastic), urban sensing (PulsodelaCiudad), and medical research (Apple's ResearchKit).

Unfortunately gathering this information from individually owned devices proved to be not a straightforward task. Some of these platforms have suffered from insufficient participation because users that voluntarily submit their sensing data found no interest in remaining active in the system without being rewarded, or at least, have their cost covered. These undesirable facts have already been observed in [4, 9, 8], which focus on incentive issues arising in the supplier's side. Namely, suppliers may drop out unless there is a positive Return on Investment, which depends on the total cost for collecting data (battery consumption, device resources, privacy, etc). But potential buyers of data may also be reluctant to participate in the market if for instance the prices are prohibitively high, or if the underlying mechanisms do not aim at economic efficiency. How should applications express their demand for sensor information in such an environment and how should prices be determined? At the same time, one also needs mechanisms for matching (elastic) demand with potential sensor providers, exploiting the fact that once a sensor is turned on it can be used simultaneously by multiple applications.

Contribution. To begin with, we develop a framework for operating a large market of sensor data in participatory sensing environments. On one side of the market, we have buyers interested in obtaining data potentially from multiple sources and for different types of sensors. The demand of the buyers can be elastic or inelastic in terms of the number of sensors they require. On the other side of the market, the data suppliers correspond to users or organizations owning sensors, and require a payment to cover their costs. In this work, we are not concerned on how suppliers define their prices. We focus on the buyers' side and the design of mechanisms that match demand with supply for the market environment described above. We assume a centralized platform that is able to execute such pricing schemes, as well as collect and distribute the data and payments. An important aspect of such a marketplace is that we can distribute the same sensor data to multiple interested parties, at no extra cost, i.e., we can view sensors as digital goods. This implies important positive externalities between buyers, since they can profit from each other by sharing sensor costs where possible.

Within this framework, we introduce 2 simple cost-sharing settings regarding the demand of the users, and study various mechanisms. The first scenario involves single-minded buyers, interested in different subsets of sensor types. The second scenario concerns bidders with multi-unit, elastic demand. In both scenarios, the two important and conflicting objectives we care for are *i) budget–balance:* the market management platform does not incur economic loss while operating the system, making it self-sustainable, *ii) economic efficiency, i.e., social welfare maximization:* we would like to satisfy more customer queries if this increases the net surplus in the system.

Given the strong impossibility results of [5, 12], we cannot achieve both at the same time. We show for both scenarios how to achieve each one of these two objectives separately, and with polynomial time complexity. For economic efficiency we prove that the VCG mechanism can be implemented in polynomial time, whereas for budget balance, we utilize and adapt ideas from the Moulin-Shenker mechanisms [10, 11].

For the second scenario, we also propose a natural hybrid mechanism that improves efficiency under budget balance by relaxing strategyproofness. Despite the loss of strategyproofness, our hybrid mechanism has its own merits. It is simple to implement, and is based on a very natural approach for increasing the social welfare. We prove that this mechanism achieves higher welfare than other established cost-sharing mechanisms. Furthermore, our hybrid mechanism satisfies certain fairness properties, in the sense that wealthier players contribute more to the total cost than poorer ones. Finally we also study welfare properties at the equilibria of this mechanism and exhibit cases where socially optimal equilibria exist.

1.1 Related Work

Regarding mechanism design for participatory sensing, in [9], a specialized reverse auction is proposed to incentivize suppliers to increase their participation. Another reverse auction is also proposed in [8]. The work of [4] on the other hand is limited to using a fixed price approach. An issue that is not covered by these works is the modeling of the demand side of the market, which is what we mainly address in this paper.

We have recently performed experimental evaluations for some of the mechanisms we study here, reported briefly in [14]. The main message from these simulations is that certain altruistic versions of budget-balanced mechanisms, where some richer players could contribute a higher payment, may have a practical appeal. The fact that buyers here are able to share the same sensor, implies that in some occasions wealthy buyers may have incentives to help out and contribute a higher cost-share so that the costs are covered and they can still have access to sensors. No theoretical analysis is provided though in [14]. The hybrid mechanism we propose in Section 4 is motivated by such observations (though not implemented or suggested in [14]).

The works from the economics literature that are most relevant to ours are the cost-sharing mechanisms of Moulin and Shenker [10, 11]. These mechanisms work for a setting where each user is either granted the same identical service with all other users or is declined. We also consider the Marginal Cost Pricing mechanism, see [10], which is the adaptation of the VCG mechanism into the cost-sharing setting.

2 Definitions and Notation

In the models we study, we have a set $N = \{1, ..., n\}$ of potential buyers, who have a demand for some sensor data (we use interchangeably the terms buyer or player, to refer to any $i \in N$). Different types of demand (e.g., elastic vs inelastic, or single tuple vs multiple tuples) are examined in Section 3 and Section 4. We also have a set $M = \{1, ..., m\}$, representing the different sensor *basic types*, e.g., accelerometer, temperature, CO_2, etc. Finally, we have a set of suppliers or providers who own sensor data (via their mobile or any other device). Each suppliers may specify a price per sensor type that he needs to be paid for in order to provide access to the value of the sensor. Note also that a value provided by one supplier can be used by many buyers. Suppliers do not all necessarily have the same set of sensor types available.

The main focus in our work will be on the following criteria, and especially on the first two:

- **Budget balance.** A mechanism is *budget-balanced* if for every instance, the payments assigned to the buyers cover exactly the cost of the provider.
- **Social welfare maximization.** Following [10], the *social welfare* or *surplus* in a cost-sharing setting is the sum of the buyers' derived values minus the cost incurred (the payments made by the buyers cancel out with what the providers receive). If x

denotes an outcome of a mechanism, the social welfare is $\sum_i v_i(x) - C(x)$, where v_i is the valuation of buyer i and $C(x)$ is the cost incurred.
– **No Positive Transfers (NPT):** The cost shares are always nonnegative.
– **Voluntary Participation (VP):** The welfare level corresponding to not providing service at no cost is guaranteed to each agent if they report truthfully.

3 Scenario 1: Single-minded buyers

We consider a simple scenario, in which each buyer $i \in N$ is interested in a subset $P_i \subseteq M$ of sensor types. For example P_i could be of the form (speedometer, accelerometer). Furthermore, he requests access to a single tuple with values from these types of sensors, i.e., a tuple (x, y), where x is a value for speed and y a value for the acceleration. These values do not necessarily need to come from the same provider (but buyers can request that all the data come from providers within a certain geographical region, e.g., the city center, in order to collect information about traffic; we omit such implementation aspects from the description of the mechanisms). Hence, the request specified by each buyer $i \in N$, is in the form (v_i, P_i), where v_i is the value derived by i for receiving this tuple, i.e., his willingness to pay. The demand is inelastic in the sense that buyer i is not deriving any utility if he receives only a strict subset of sensors from P_i. We call such buyers *single-minded*, in analogy to single-minded bidders in combinatorial auctions. Clearly such demands can come and go dynamically in the course of time, but we are interested in a static snapshot, i.e., an instance of our problem may correspond to the demands within a given time window during which the centralized platform needs to make a decision on which users to serve.

The cost function $C(S)$ for serving a set of customers $S \subseteq N$ can be easily computed for any S. For any sensor type $j \in M$, let c_j be the cost for the platform of providing a single value for this type. The values of the sensors can be viewed as digital goods, and since each bidder is interested in receiving a single tuple, we can use just one actual sensor for each type requested, to satisfy all customers. Hence, the cost c_j could be taken to be the cheapest price specified by some supplier of type j (it is not though important for the mechanism how c_j is derived). Therefore, for a set $S \subseteq N$ of buyers, the cost $C(S)$ is the sum of the costs of all sensor types required by S:

$$C(S) = \sum_{j \in P(S)} c_j, \text{ where } P(S) = \bigcup_{i \in S} P_i \,. \tag{1}$$

3.1 Social Welfare Maximization

We first look at the objective of maximizing the social welfare. Let $\theta = (\theta_1, ..., \theta_n)$ be the vector of the agents' declared types. Under Scenario 1, $\theta_i = (v_i, P_i)$. If a mechanism chooses $S \subseteq N$ as the set of buyers to be served, then the generated welfare from S is: $SW(S, \boldsymbol{\theta}) = \sum_{i \in S} v_i - C(S)$.

Let us denote by $SW^*(N, \theta)$ the optimal welfare that can be achieved by N, i.e.:

$$SW^*(N, \theta) = \max_{S \subseteq N} \{\sum_{i \in S} v_i - C(S)\}$$

Our main result in this section is the following:

Theorem 1. *The problem of social welfare maximization under Scenario 1 can be solved in polynomial time.*

To prove Theorem 1, we need to avoid the exponential search over all subsets of N. Note also that we do not have any monotonicity properties here (larger sets do not necessarily produce higher welfare). To solve our problem, we resort to a linear programming formulation, which turns out to yield a totally unimodular constraint matrix.

Proof of Theorem 1 : We begin by writing down an ILP for our problem. For this, we use an integer variable x_i for each buyer $i \in N$ and an integer variable y_j for each sensor type $j \in M$. The rationale is clearly that when $x_i = 1$, agent i receives his requested tuple P_i. When $y_j = 1$, this means that the sensor of type j is allocated. Note that determining the set of players who receive service, also determines the set of sensor types that will be set to 1. We claim that the following is an ILP describing our problem.

$$\begin{aligned} \text{maximize: } & \sum_{i \in N} v_i x_i - \sum_{j \in M} c_j y_j \\ \text{subject to: } & x_i \le y_j \ , \qquad \forall i \in N, \forall j \in P_i \\ & x_i \in \{0,1\} \ , \quad \forall i \in N \\ & y_j \in \{0,1\} \ , \quad \forall j \in M \end{aligned}$$

To see why this suffices, let us see the relation between the variables x_i and y_j for each $j \in P_i$. If $x_i = 0$, then the variable y_j could be either 0 or 1, depending on other buyers' demand sets. If $x_i = 1$ however, then we must have that $y_j = 1$. Hence, the only constraint beyond integrality that we need is that $x_i \le y_j$ for $j \in P_i$. It is easy to see now that every solution to our problem corresponds to a feasible solution of the ILP (there are also some feasible solutions in which we can have $y_j = 1$ without allocating the sensor to anybody but these are clearly not optimal solutions).

We relax the ILP to get an LP relaxation, by setting that $x_i, y_j \in [0,1]$. So now we have a linear program, which we can write in the form $\{max\ w^T z \mid Az \le b, z \ge 0\}$. The rest of the proof is devoted to showing that our constraint matrix A is totally unimodular, which implies that the LP always has an integral optimal solution.

Lemma 1. *The constraint matrix A of the LP relaxation is totally unimodular.*

The proof of Lemma 1 is in **Appendix A**. Therefore, we can solve the Social Welfare maximization problem in polynomial time, and the proof of Theorem 1 is complete. □

Theorem 1, implies that we can have strategyproof and efficient mechanisms implemented in polynomial time. For example, we can utilize the VCG mechanism, which we briefly recall for the sake of completeness. The VCG mechanism first computes a set $S^* \subseteq N$, where optimal welfare is attained. Then, if the declared type vector is $\theta = (\theta_1, ..., \theta_n)$, the payment for every player $i \in S^*$, can be written in the form:

$$p_i = b_i - (SW^*(N, \boldsymbol{\theta}) - SW^*(N \setminus \{i\}, \boldsymbol{\theta}_{-i})) \,. \tag{2}$$

Here b_i is the value declared by player i to the mechanism. Agents not picked in the optimal set do not pay anything. This is also known as the pivotal mechanism [3], and also referred to in the cost-sharing context, as the Marginal Cost (MC) mechanism in [11]. Hence, we can conclude with the following:

Corollary 1. *Under scenario 1, the VCG mechanism is strategyproof, satisfies NPT and VP, and can be implemented in polynomial time.*

More generally, we can have a family of strategyproof mechanisms by replacing $SW^*(N \setminus \{i\}, \boldsymbol{\theta}_{-i})$ in (2) with any function of the form $h_i(\boldsymbol{\theta}_{-i})$.

Since VCG is efficient, the impossibility results of [5, 12] imply that it cannot be budget-balanced. In fact, we cannot even hope to be "approximately" budget-balanced, since in the cases where no player is pivotal, the VCG payments are all 0.

3.2 Budget-balanced mechanisms

We now focus on the design of budget-balanced mechanisms. The mechanism we consider is derived directly from the pioneering work of Moulin and Shenker [10, 11]. Their work concerns a setting that differs from ours in 2 respects: first, their model is simpler in terms of the service requested. Namely, they have a binary setup, where there is a single provider, offering the same identical service to everyone, and each agent will be either granted or declined the service. In our case the buyers are interested in different subsets, and hence in a different type of service each. Second, in our model, the cost function is simpler due to the fact that sensors correspond to digital goods and can be shared. This implies that for instances where a set S of buyers requests the same set of sensors, then in our setting $C(S)$ is the same as $C(T)$ for any $T \subseteq S$ with $T \neq \emptyset$. In their work $C(\cdot)$ is an arbitrary submodular set function.

We can easily adapt the approach of Moulin and Shenker for Scenario 1. To do this, we need to define first an underlying *cost-sharing* method. A cost-sharing method is a function $\xi(\cdot, \cdot)$ such that $\xi(i, R)$ determines the cost-share of agent i, when R is the set to be served by the mechanism. We demand that a cost-sharing method satisfies $\sum_{i \in R} \xi(i, R) = C(R)$ for all $R \subseteq N$, i.e., the sum of the payments balance the cost.

We mainly focus on the *egalitarian* cost-sharing method, since this may have more appeal in practice due to its simplicity. To define the share $\xi(i, R)$ for a given set R to be served, we split the cost of each used sensor equally among the people who want it. Let y_j be the number of buyers who have j in their demand set. Egalitarian cost sharing means that each customer i contributes a share c_j / y_j towards the cost of sensor j. Hence for a buyer i, with demand set P_i, his total cost-share is:

$$\xi(i, R) = \sum_{j \in P_i} \frac{c_j}{y_j}. \quad (3)$$

It is obvious that we have: $\sum_{i \in R} \xi(i, R) = C(R)$, for any $R \subseteq N$. Given now any cost-sharing method ξ, one can define parametrically the mechanism below for determining who receives service along with the cost-shares. In the description below, we let $\mathbf{b} = (b_1, ..., b_n)$ be the agents' declared values for their demand sets.

The Mechanism MS(ξ) (Moulin-Shenker mechanism under $\xi(\cdot, \cdot)$):

- Start by trying to serve all agents, with cost-share $\xi(i, N)$. Remove any agent who cannot cover his share, i.e., anyone for which $b_i < \xi(i, N)$. If no one is removed in this step, stop here, otherwise let R^1 be the set of remaining agents.

– Check if we can serve R^1 with a cost-share of $\xi(i, R^1)$ for every $i \in R^1$. Again remove those who cannot afford this price.
– Continue like this and in every round obtain the set $R^{t+1} = \{i \in R^t : b_i \geq \xi(i, R^t)\}$.
– Stop either when we reach the empty set, or when we reach a set in which all agents can afford to pay their cost-share.

This family of mechanisms turns out to have nice properties if the cost function $C(\cdot)$ and the cost-sharing method $\xi(\cdot,\cdot)$ satisfy certain conditions. Regarding $\xi(\cdot,\cdot)$, the following is an important and desirable property, which simply says that the cost-share of an agent should not become higher when more people receive service.

Definition 1. *A cost-sharing method is cross-monotonic if for any $T \subseteq N$,*

$$\xi(i, R) \geq \xi(i, T) \text{ for any } R \subseteq T \text{ and } i \in R\,. \tag{4}$$

Claim 2 *The egalitarian cost-sharing method described by (3) is cross-monotonic.*

We also need submodularity of our cost function, which is easy to establish.

Given Claim 2 , the following theorem is a straightforward extension of the results from [10, 11] to our setting.

Theorem 3. *Given any cross-monotonic cost-sharing method ξ for single-minded bidders, the Mechanism $MS(\xi)$ is budget-balanced, group-strategyproof and satisfies NPT and VP. In particular, if ξ is the egalitarian cost-sharing according to* (3)*, $MS(\xi)$ satisfies these properties and can also be implemented in polynomial time.*

The obvious question is how do these mechanisms perform with respect to social welfare. Unfortunately, they are far from being efficient. Example 1 in Appendix A exhibits instances where the efficiency loss can be made arbitrarily large. The mechanism generates zero welfare, whereas the optimal social welfare is far from zero.

Our discussion in Sections 3.1 and 3.2 highlights the tradeoff between achieving efficiency and budget-balance. In the next section, we will see a way of achieving better trade-offs in a scenario of multi-unit elastic demand (but not applicable to Scenario 1).

4 Scenario 2: Multiple units and Elastic Demand

At this orthogonal scenario all players have the same type of demand, i.e., the set P_i is the same for every player (for example, this could involve buyers who are all interested in the same type of information about the city center, or the same type of environmental sensors in a region). What differentiates the players is that each player i specifies an additional amount d_i, representing the maximum number of tuples that he is interested in acquiring. The demand d_i is elastic, so that player i does not mind receiving less than d_i tuples. Each player also specifies his per-tuple willingness to pay v_i. This encodes an additive[3] valuation up to d_i tuples. We assume that there is a sufficient supply of tuples

[3] Our results also can be adapted for general submodular valuations in the form $v_i = (v_i(1), .., v_i(d_i))$, where $v_i(j)$ is the value for the j-th tuple. We prefer the current exposition, due to its simplicity and more practical appeal for participatory sensing applications.

from the providers, i.e., there are at least d_{max} of them with $d_{max} = \max d_i$. Each tuple has some cost c_k so that we can sort them from the cheapest to the most expensive one, say $c_1 \leq c_2 \cdots \leq c_{d_{max}}$. It is not important for our mechanisms how c_k is derived.

We start with showing that maximizing the social welfare can be solved in polynomial time. The important property is that once we decide for allocating a tuple, we do not lose in welfare by giving the tuple to all customers who have demand for it, since we are only adding more value to the current welfare. Hence, if $\boldsymbol{\theta} = (\theta_1, ..., \theta_n)$ is the type vector, with $\theta_i = (v_i, d_i)$, the optimization problem for the social welfare becomes

$$SW^*(N, \boldsymbol{\theta}) = \max_{1 \leq k \leq d_{max}} \big[\sum_{i \in N} v_i \cdot \min\{k, d_i\} - \sum_{j=1}^{k} c_j \big] \tag{5}$$

We can solve (5) simply by trying all values for k. Hence we have:

Theorem 4. *Under Scenario 2, we can have polynomial time, strategyproof, and efficient cost-sharing mechanisms, that also satisfy NPT and VP.*

4.1 Budget balance: Sequential Moulin-Shenker Mechanisms

The application of Moulin-Shenker mechanisms is not any more straightforward in the case of buyers with multi-unit demand. Each customer i corresponds now to a set of potential service levels, ranging from 0 to d_i tuples. Hence, we cannot just run an analog of $MS(\xi)$ from Section 3. One could consider all combinations of service levels to customers, and run $MS(\xi)$ for each such combination (and then choose the one that is more efficient). But this has prohibitively high complexity to be run in practice.

Instead, one can utilize the Moulin-Shenker approach in a sequential manner.

The Mechanism SMS(ξ) (Sequential Moulin-Shenker):

- Sort the d_{max} cheapest tuples so that $c_1 \leq c_2 \cdots \leq c_{d_{max}}$. Let $A^1 = N$ be the set of active players before the first round (initially all are active).
- At round r (with r ranging from 1 to d_{max}):
  - If A^r is the set of currently active players, run the mechanism MS(ξ) from Section 3 on A^r, to determine who receives the r-th cheapest tuple, along with their cost shares for that round.
  - Remove from A^r all customers who were not selected to be served. Remove also any customer with $d_i = r$.
  - Let A^{r+1} be the set of surviving customers after the previous step. Continue with the next round in the same manner, unless $A^{r+1} = \emptyset$.
- The total cost-share of a player is the sum of the cost-shares from all rounds.

For the remainder of the paper, we fix again ξ to be the egalitarian cost sharing method and denote the mechanism as SMS, rather than SMS(ξ). Since everybody is interested in the same tuple, if there are say k active players in a certain run of SMS at a round r, the cost share is defined as c_r/k. The SMS mechanism is (group) strategy proof by using the same arguments as in Theorem 3 (since the cost shares do not depend on what players declare). Hence:

Theorem 5. *Under Scenario 2, the SMS mechanism with egalitarian cost-shares runs in polynomial time, is budget-balanced, group-strategyproof and satisfies NPT and VP.*

However, as in the previous Section, we can easily construct instances where we have a great loss of efficiency, even with 2 players and 1 round.

4.2 Budget-balance with Better Social Welfare: A Hybrid Mechanism

Motivated by the fact that the Moulin-Shenker mechanisms do not yield high social welfare, we propose in this section a different mechanism, as an attempt to maintain budget-balance but achieve higher welfare. Our mechanism is quite intuitive and uses a very natural approach in order to achieve better welfare. It also proceeds in rounds, but in each round, we start by running the VCG mechanism for sharing the tuple of that round. In order to achieve budget-balance, we also complement the VCG payments with an egalitarian cost-share for the remaining cost. If this results in high costs for some players, we reject them and repeat for the remaining players.

Assume that the input to the mechanism is $\boldsymbol{\theta} = (\theta_1, \cdots, \theta_n)$ with $\theta_i = (b_i, d_i)$. We define first the per-round VCG mechanism, which is quite simple in this setting. If we run VCG only for the tuple at round r, the tuple is allocated if $\sum_j b_j \geq c_r$. A player i is *pivotal* if $\sum_j b_j \geq c_r$ and $\sum_{j \neq i} b_j < c_r$, i.e., player i has an impact on having the tuple allocated. The only players that pay under VCG are the pivotal players, according to (2). Hence, if A^r is the set of active players at round r, and if $\sum_{j \in A^r} b_j \geq c_r$, the VCG payments are:

$$p_i{}^{VCG} = \begin{cases} c_r - \sum_{j \neq i} b_j, & \text{if player } i \text{ is pivotal} \\ 0, & \text{if player } i \text{ is not pivotal} \end{cases} \tag{6}$$

Our mechanism runs as follows:

The Hybrid Mechanism

1. Again sort the tuples so that $c_1 \leq c_2 \cdots \leq c_{d_{max}}$. Let $A^1 = N$.
2. At round r (with r ranging from 1 to d_{max}):
   (a) Check if $\sum_{i \in A^r} b_i \geq c_r$, where A^r is the set of currently active players during round r. If not, the mechanism stops.
   (b) Run the VCG mechanism on A^r, for the tuple of round r, and let $p_i{}^{VCG}$ be the VCG payment for each $i \in A^r$, as defined in (6).
   (c) Let $c_r{}'$ be the reduced cost after the VCG payments: $c_r{}' = c_r - \sum_{i \in A^r} p_i{}^{VCG}$.
   (d) Split the cost c_r' equally among A^r, i.e. let $p^E = c_r'/|A^r|$. Define the candidate cost shares as $p_i{}^H = p_i{}^{VCG} + p^E$.
   (e) If there are players with $b_i < p^H$, then pick the one with the lowest bid, set $A^r = A^r \setminus \{i\}$, and go to step 2a to repeat the process for round r.
   (f) Otherwise, if $b_i \geq p^H$, for each $i \in A^r$, set $A^{r+1} = A^r \setminus \{i : d_i = r\}$, and continue to round $r + 1$, unless $A^{r+1} = \emptyset$.
3. The total payment of each player is the sum of payments over all rounds.

Remark 1. *(i)* Note that at step 2e, we remove only one player, even if there can be more players with $b_i < p_i^H$. This turns out to be crucial regarding the total welfare achieved. See Example 2 in **Appendix B** for more details on this.

(ii) Example 3 in **Appendix B** shows that an analog of the Hybrid mechanism in Scenario 1 does not necessarily produce better social welfare than $MS(\xi)$.

On the positive side, we will prove that the hybrid mechanism can attain much higher social welfare than SMS. On the negative side, this is not a strategyproof mechanism (see Claim 11 and Example 4 in Appendix B). We do not view the lack of strategyproofness as a prohibitive disadvantage for such mechanisms. In the recent literature there have been several studies analyzing simple and non-strategyproof mechanisms that have practical appeal. In the context of auctions for example, see e.g., [1, 2, 7].

Apart from achieving better welfare, the Hybrid mechanism has other merits as well. First, it maintains low complexity like SMS, since the VCG step is very easy to run. Second, we consider it a very natural approach towards increasing the welfare of budget-balanced mechanisms and can be applicable to other settings too. Third, it also satisfies certain fairness properties, in the sense that wealthier players contribute more to the total cost than poorer ones, see Claim 6 below. During the VCG step in each round, the set of players who pay are the richer ones, according to (6). The remaining cost is then an egalitarian cost share for all active players. Hence, the mechanism helps the poorer players to satisfy their demand. But in addition to that, the wealthier players are also rewarded, since as we will see in Lemma 2, a positive payment at the VCG step for player i, ensures that i is never removed during the execution and he will be able to get the desired tuples (as long as the cost of a tuple is covered by the sum of bids).

Claim 6 *In the Hybrid mechanism, if* $b_i \geq b_j$ *then* ${p_i}^H \geq {p_j}^H$.

The main positive result for the Hybrid mechanism is that it dominates the SMS mechanism as follows:

Theorem 7. *For any type vector* $\boldsymbol{\theta}$*, if we run both the Hybrid and the SMS mechanism on input* $\boldsymbol{\theta}$*, then the Hybrid mechanism always achieves at least as good social welfare as the SMS mechanism, w.r.t.* $\boldsymbol{\theta}$.

Proof. The proof is based on two auxiliary lemmas stated below, the proofs of which are in **Appendix C**. The following lemma shows that players who are asked to pay something from the VCG run of a certain round cannot be rejected at that step of the mechanism (in fact this implies that they will not be rejected from any future round where the sum of bids covers the cost).

Lemma 2. *Consider a round* r *in the Hybrid mechanism and let* A^r *be the set of active players just before an execution of step 2b within round* r*. If* $\sum_{j\in A^r} b_j \geq c_r$*, then for every player* $i \in A^r$ *for which* ${p_i}^{VCG} > 0$*, the mechanism cannot remove* i *from* A^r *during that step, i.e.,* $b_i \geq {p_i}^H$ *in the execution of that iteration.*

Using Lemma 2, we can now prove the following fact.

Lemma 3. *Consider a run of the SMS and the Hybrid mechanism on the same instance. At every round* r *of each mechanism, let* N_r^S *and* N_r^H *be the set of players who receive the* r*-th tuple by the SMS mechanism and by the Hybrid mechanism respectively. Then* $N_r^S \subseteq N_r^H$*, for every* r.

Lemma 3 implies that the Hybrid mechanism produces at least as good social welfare as the SMS mechanism in each round. Hence, this completes our proof. □

In the case of 2 players, we can even guarantee optimal welfare by the Hybrid mechanism (see Theorem 12 in Appendix C).

Equilibria under the Hybrid mechanism: An obvious question is whether we can have a price of anarchy analysis in the cost-sharing setting as well. Do all the Nash equilibria of the Hybrid mechanism achieve good social welfare? The answer is negative as there exist many "unreasonable" equilibria. E.g., the profile where everybody declares 0 is an equilibrium and this is inherent in most cost-sharing mechanisms, since no player would be willing to cover the cost of a service on his own. Nevertheless, these are equilibria that will not be attained in practice.

Given the high price of anarchy, the next step is to investigate the existence of equilibria with better guarantees.The Hybrid mechanism is promising in that direction. We briefly mention some results here for the existence of socially optimal pure equilibria.

Theorem 8. *Consider a set of players with the same demand $d_i = d$ for $i \in N$. Then, there is a Nash equilibrium producing optimal social welfare when $d = 1$ or when all the tuples have the same cost, $c_1 = ... = c_d$. In fact, in both cases, if the optimal welfare is positive (i.e., $\sum v_i > c_1$), then every vector $\mathbf{b}$ with $\sum_{j=1}^{n} b_j = c_1$ and $b_i \leq v_i$, is a Nash equilibrium which produces optimal social welfare (w.r.t. the true valuation vector).*

The proof is presented in Appendix C. As we see, there can be a plethora of optimal equilibria in the above cases. Next, we identify some more conditions that enable the existence of socially optimal equilibria. For simplicity, we stick to the case where all players have the same demand d and the optimal welfare is achieved by allocating all d tuples. Note then, that at an equilibirum of the hybrid mechanism, we need to have $\sum b_j = c_d$, i.e., if the bids exceed the cost of the last round, then there are incentives for people to deviate. Second, to enforce an efficient equilibrium, we also need some relation between the values v_i, the parameter d, and possibly the marginal cost increase between rounds. The following conditions that we have identified say that as long as we do not have very poor players (otherwise again some people will have incentives to shade their bids), socially optimal equilibria do exist.

Theorem 9. *Consider an instance with players having the same demand d as before and let $\delta = max_i\{c_i - c_{i-1}\}$. If the following 2 conditions hold, there exists a socially optimal equilibrium.*

1. $v_i > 2(d-1)\delta$,
2. $c_d \in [n(d-1)\delta, \sum_i v_i - n(d-1)\delta]$.

5 Concluding Remarks

To our knowledge, a market-place tailored to the specificities of participatory applications, is missing today. We conjecture that with the wider adoption of devices containing sensors and new types of micro-payments, a market-place for data originating from individually owned devices will be developed.

The implementation details can be very significant and potentially critical in the success of the entire scheme. As an example, on the keyword auctions area, the widely used approach was based on a second-price mechanism, although a VCG mechanism would offer clear theoretical advantages such as strategyproofness. However, the second-price auction was easier to explain and understand which could had played a more important factor towards the success of the model.

There are still many directions and mechanisms that one can explore in the context of sensor-data markets by focusing on the optimization of different criteria. For example, similarly to keyword auctions, is there a very simple mechanism with good properties? We underline that, like in this work, the focus could be on properties with much more practical appeal like budget-balance perhaps at the expense of theoretical guarantees and worst-case scenarios.

References

1. Bhawalkar, K., Roughgarden, T.: Welfare Guarantees for Combinatorial Auctions with Item Bidding. In: Proceedings of the ACM-SIAM Symposium on Disctrete Algorithms (SODA). pp. 700–709 (2011)
2. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian Combinatorial Auctions. In: Proceedings of the International Colloquium on Automata, Languages and Programming (1) (ICALP). pp. 820–832 (2008)
3. Clarke, E.H.: Multipart pricing of public goods. Public Choice 11, 17–33 (1971)
4. Danezis, G., Lewis, S., Anderson, R.J.: How much is location privacy worth? In: WEIS (2005)
5. Green, J., Kohlberg, E., Laffont, J.J.: Partial equilibrium approach to the free rider problem. Journal of Public Economics 6, 375–394 (1976)
6. Heller, I., Tompkins, C.B.: An extension of a theorem of Dantzig's. In: Kuhn, H.W., Tucker, A.W. (eds.) Linear Inequalities and Related Systems, pp. 247–254. Princeton University Press (1956)
7. de Keijzer, B., Markakis, E., Schäfer, G., Telelis, O.: Inefficiency of Standard Multi-unit Auctions. In: Bodlaender, H.L., Italiano, G.F. (eds.) ESA. LNCS, vol. 8125, pp. 385–396. Springer (2013)
8. Koutsopoulos, I.: Optimal incentive-driven design of participatory sensing systems. In: INFOCOM. pp. 1402–1410 (2013)
9. Lee, J.S., Hoh, B.: Dynamic pricing incentive for participatory sensing. Pervasive and Mobile Computing 6(6), 693–708 (2010)
10. Moulin, H.: Incremental cost sharing: Characterization by coalition strategy-proofness. Soc. Choice Welfare 16, 279–320 (1999)
11. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: Budget balance vs efficiency. Econ. Theory 18, 511–533 (2001)
12. Roberts, K.: The characterization of implementable choice rules. In: Laffont, J.J. (ed.) Aggregation and Revelation of Preferences. Amsterdam: North Holland (1979)
13. Schrijver, A.: Theory of Linear and Integer Programming. John Wiley and Sons (1986)
14. Thanos, G.A., Courcoubetis, C., Markakis, E., Stamoulis, G.D.: Design and experimental evaluation of market mechanisms for participatory sensing environments. In: AAMAS 2014,. pp. 1515–1516 (2014)

A Missing material from Section 3

Proof of Lemma 1 : First, we will use the following auxiliary fact.

Fact 10 *If A is a totally unimodular matrix then A^T is also totally unimodular.*

We now apply the above fact to the conditions for total unimodularity as established by [6], see also [13][page 276], which imply the following sufficient condition for total unimodularity:

Proposition 1. *Consider a matrix A, and suppose we are given a partition of its columns into two disjoint sets B and C, such that the following conditions hold.*

(i) Every entry in A is 0, $+1$, or -1, and every row of A contains at most two non-zero entries.

(ii) If there are two non-zero entries in a row of A with the same sign, then the columns corresponding to these entries do not belong to the same class (one is in B, and the other is in C).

(iii) If there are two non-zero entries in a row of A with opposite signs, then the columns corresponding to these entries both are in B, or both in C.

Then A is totally unimodular.

In our case, we have the following LP.

$$\begin{aligned}
\text{maximize: } & \sum_{i \in N} v_i x_i - \sum_{j \in M} c_j y_j \\
\text{subject to: } & x_i \le y_j \ , \quad \forall i \in N, \forall j \in P_i \\
& x_i \le 1 \ , \quad \forall i \in N \\
& y_j \le 1 \ , \quad \forall j \in M \\
& x_i \ge 0, y_j \ge 0 \ , \quad \forall i \in N, j \in M
\end{aligned}$$

If we take the variable vector to be $z = (x_1, \dots, x_n, y_1, \dots, y_m)^T$, then our constraint matrix A will have the form:

$$A = \begin{bmatrix}
1 & 0 & \cdots & 0 & \cdots & -1 & 0 & \cdots & 0 \\
1 & 0 & \cdots & 0 & \cdots & 0 & -1 & \cdots & 0 \\
\vdots & \vdots & & \vdots & & \vdots & \vdots & & \vdots \\
0 & 0 & \cdots & 1 & \cdots & 0 & -1 & \cdots & 0 \\
1 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\
0 & 1 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 0 \\
\vdots & \vdots & & \vdots & & \vdots & \vdots & & \vdots \\
0 & 0 & \cdots & 0 & \cdots & 0 & 0 & \cdots & 1
\end{bmatrix}$$

where the first n columns correspond to the x_i variables and the remaining m columns correspond to the y_j variables. The last $n + m$ rows of A have exactly one $+1$ entry and the rest are 0 (representing the inequalities $x_i \le 1, y_j \le 1$). The first rows, above the last $n + m$ ones, which are $\sum_i |P_i|$ in total, have exactly one $+1$ and one -1 entry, representing the inequalities $x_i - y_j \leq 0$, for $j \in P_i$.

To proceed, we consider a partition of the columns of A into two disjoint sets B, C where we take B to contain all the rows of A and set $C = \emptyset$. We check each of the 3 properties of Proposition 1.

(i) Every entry in A is 0, $+1$, or -1. Also, every row of A contains at most two non-zero entries.

(ii) There is no row in matrix A with elements of the same sign. Hence, we have nothing to check here.

(iii) Rows in A containing entries with opposite signs correspond to constraints of the form $x_i - y_j \leq 0$. Hence, in all such rows there is exactly one $+1$ and one -1 entry. Thus, having all the columns of A in set B, and setting $C = \emptyset$ suffices. Note that the crucial point here is that as we previously saw, there is no row of A having entries of the same sign, and this allows us to leave C empty.

Since all the sufficient conditions hold, the matrix A is totally unimodular. □

Proof of Claim 2 : Let $T \subseteq N$ and $R \subseteq T$. If $\xi(i, R)$ is the cost share of i under R, then adding more buyers to the set R cannot increase the cost share of i. For each $j \in P_i$, the cost share for i towards c_j can only decrease when more buyers are added. Hence, $\xi(i, R) \leq \xi(i, T)$. □

Proof of submodularity of cost function : Consider any two sets $S, T \subseteq N$, with $S \subseteq T$. We claim that adding one more buyer i to each of these sets has a smaller marginal cost in T. This is because the marginal cost $C(T \cup i) - C(T) = \sum_{j \in P_i \setminus P(T)} c_j$. And since, $P(S) \subseteq P(T)$, we have

$$C(T \cup i) - C(T) \leq C(S \cup i) - C(S)$$

This completes the proof. □

We next present an example showing that the loss in social welfare by the Moulin-Shenker mechanisms under Scenario 1, can be arbitrarily large. We show this both for the egalitarian cost-sharing method and also for the Shapley value. But these negative examples do not hold just for these two cost-sharing methods, one can have such bad examples for any other cost-sharing method.

Example 1. Suppose we have two players $N = \{1, 2\}$, and two sensor types, say t_1 and t_2.
Players: Player 1 has a value of $v_1 = 2M$ for some number $M \geq 3$ and he requests the set $P_1 = \{t_1, t_2\}$. Player 2 has a value of $v_2 = 2M - 1$ and he requests the set $P_2 = \{t_1\}$.
Sensors: The cost for the sensor of type t_1 is $c_1 = 2M$ and the cost for the sensor of type t_2 is $c_2 = M + 1$.

We can see now that the costs for serving the 4 different subsets of the 2 players are as follows: $C(\emptyset) = 0$, $C(N) = C(\{1, 2\}) = c_1 + c_2 = 3M + 1$, $C(\{1\}) = \sum_{j \in P_1} c_j = c_1 + c_2 = 3M + 1$, and $C(\{2\}) = \sum_{j \in P_2} c_j = c_1 = 2M$.

Let us consider the egalitarian cost-sharing formula as defined by (3). In the first round of the mechanism $M(\xi)$, the cost shares will be:

$$\xi(1,N) = \frac{1}{2}c_1 + c_2 = 2M+1, \;\; \xi(2,N) = \frac{1}{2}c_1 = M \tag{7}$$

But this implies that player 1 will be kicked out since $v_1 < 2M+1$ and player 2 will continue to the second round since $v_2 \geq M$. In the second round of $M(\xi)$, player 2 has now to cover the cost for the sensor he wants on his own. But since $v_2 < 2M$, he cannot afford to do so and he is rejected as well. This means that $M(\xi)$ produces zero social welfare. In contrast, the optimal social welfare is achieved when we serve both players and it is $SW^* = v_1 + v_2 - (c_1 + c_2) = M - 2 > 0$. Clearly, we can make this arbitrarily large, as M becomes bigger.

Sticking to the same instance, let us consider now a different cost-sharing method, namely the Shapley Value, which is one of the usual alternatives in the literature. One of the reasons that this method has been attractive is that among the family of Moulin-Shenker mechanisms, it achieves the minimal (additive) worst-case efficiency loss, see [11]. A negative aspect of this method however is that computing the cost shares cannot be done in polynomial time. We exhibit that even for this method, the efficiency loss can be arbitrarily large.

For $S \subseteq N$ and $i \in S$, the cost share under Shapley value is defined as follows:

$$\xi(i,S) = \sum_{T \subseteq S-i} \frac{|T|!(|S|-|T|-1)!}{|S|!}[C(T \cup i) - C(T)]$$

If we run now the Moulin-Shenker mechanism, in the first round we need to compute $\xi(1,N)$ and $\xi(2,N)$, which turn out to be:

- $\xi(1,N) = \sum_{T \subseteq \{2\}} \frac{|T|!(|N|-|T|-1)!}{|N|!}[C(T \cup 1) - C(T)] = \frac{1}{2}(C(N) - C(\{2\})) + \frac{1}{2}(C(\{1\}) - C(\emptyset)) = \frac{1}{2}(c_1 + c_2 - c_1) + \frac{1}{2}(c_1 + c_2) = \frac{1}{2}c_1 + c_2$
- $\xi(2,N) = \sum_{T \subseteq \{1\}} \frac{|T|!(|N|-|T|-1)!}{|N|!}[C(T \cup 2) - C(T)] = \frac{1}{2}(C(N) - C(\{1\})) + \frac{1}{2}(C(\{2\}) - C(\emptyset)) = \frac{1}{2}(c_1 + c_2 - c_1 - c_2) + \frac{1}{2}c_1 = \frac{1}{2}c_1$

We see that these are identical to the cost shares defined by the egalitarian cost-sharing method in (7). Hence, the mechanism will run in the same manner as before, and it will result in zero welfare again. □

B Illustrative examples regarding the Hybrid mechanism

In this part of the Appendix we provide some illustrative examples, referenced in Section 4.

The example below shows why we remove only one player at a time at step 2e, and do not remove all together the players for which $b_i < p_i^H$.

Example 2. Suppose that we have five players with

$Players$	1	2	3	4	5
$Bids$	2	2	1.8	1.1	1.1

and the cost of the first round is $c_1 = 6$. The sum of these bids is higher than the total cost and it is easy to see that $p^{VCG} = 0$ for all of them (for every player i we have $\sum_{j \neq i} b_j \geq c_1$). So $c'_1 = c_1$ and for every i, the final payment is $p_i = 6/5 = 1.2$. Players 4 and 5 can not pay their payments, we exclude both of them, thus they stop to participate. However the rest players are now not capable to cover the cost since $b_1 + b_2 + b_3 = 5.8 < 0$, so the procedure stops and no player gets his demanded tuple. Let us run the same example but instead of putting both players 4 and 5 out, we will choose the one with the lowest bid (say arbitarilly player 5 since players 4 and 5 have a tie) and exclude only him. Now we have four players and the procedure continues as follows,

$Players$	1	2	3	4
$Bids$	2	2	1.8	1.1

their new VCG payments are now,

$Players$	1	2	3	4
p_i^{VCG}	1.1	1.1	0.9	0.2

which gives $p_i^E = 2.7/4 = 0.675$ for every i. Thus the final results are

$b_1 = 2 > p_1^H = 1.1 + 0.675 = 1.775$
$b_2 = 2 > p_2^H = 1.1 + 0.675 = 1.775$
$b_3 = 1.8 > p_3^H = 0.9 + 0.675 = 1.575$
$b_4 = 1.1 > p_4^H = 0.2 + 0.675 = 0.875$

So a non zero social welfare is produced and we achieve a great improvement in total (we note here that SMS mechanism also produces a zero social welfare at this example). □

The next example shows that we cannot hope to have an analog of the Hybrid mechanism in Scenario 1 and always beat the Moulin-Shenker mechanisms with respect to social welfare.

Example 3. In order to define an analog of the Hybrid mechanism for Scenario 1, we need first to see how to divide the remaining cost after the VCG step to the players. This is not uniquely defined, because players now have different type of demands. So suppose that a player i demands k out of m sensors and after the VCG step he ends up with a payment ${p_i}^{VCG}$. This amount is divided by k and extracted by the costs of his demanded sensors.

We show now that the mechanism $M(\xi)$ from Section 3 with the Egalitarian cost sharing formula can produce a better social welfare from Hybrid mechanism at Scenario 1. Suppose that we have five players with

$Players$	1	2	3	4	5
$Bids$	25	10	5	9	8

and four sensors with the following costs

$Sensors$	1	2	3	4
$Costs$	20	10	7	5

Now lets take a look at the demand of each player:

$Players$	1	2	3	4	5
$Sensors$	$1, 2, 3, 4$	1	2	3	4

If we run the SMS mechanism with the Egalitarian cost sharing formula, the payment for each player is defined as follows:

$p_1 = 20/2 + 10/2 + 7/2 + 5/2 = 21$
$p_2 = 20/2 = 10$
$p_3 = 10/2 = 5$
$p_4 = 7/2 = 3.5$
$p_5 = 5/2 = 2.5$

Thus for every player i we have that $b_i \geq p_i$ so all players participate in the game. The social welfare which is produced by this procedure is also the optimal one, $SW^{MS} = 25 + 10 + 5 + 9 + 8 - 20 - 10 - 7 - 5 = 57 - 42 = 15$. Let us now proceed with Hybrid mechanism:

$p_1^{VCG} = (9 + 8 - 7 - 5) - (10 + 5 + 9 + 8 - 20 - 10 - 7 - 5) = 5 + 10 = 15$
$p_2^{VCG} = (25+5+9+8-20-10-7-5) - (25+5+9+8-20-10-7-5) = 0$
$p_3^{VCG} = (25+10+9+8-20-10-7-5) - (25+10+9+8-20-10-7-5) = 0$
$p_4^{VCG} = (25+10+5+8-20-10-7-5) - (25+10+5+8-20-10-7-5) = 0$
$p_5^{VCG} = (25+10+5+9-20-10-7-5) - (25+10+5+9-20-10-7-5) = 0$

So the only player who pays is player 1 and his payment is divided by 4, in order to reduce the cost of the sensors that he demands. Thus we have $15/4 = 3.75$ and the new costs are:

$c_1 = 20 - 3.75 = 16.25$
$c_2 = 10 - 3.75 = 6.25$
$c_3 = 7 - 3.75 = 3.25$
$c_4 = 5 - 3.75 = 1.25$

Lets examine now the final payment of player 1:

$p_1^H = p_1^{VCG} + p_1^E = 15 + (16.25/2 + 6.25/2 + 3.25/2 + 1.25/2) = 15 + 13 = 28 > 25 = b_1$

This kicks out player 1 as well as players 2 and 3 who now can not cover the cost of their demanded sensors (since player 1 is out he does not pay anything at the VCG step, in addition $b_2 = 10 < 20 = c_1$ and $b_3 = 5 < 10 = c_2$). Finally player 4 gets his demanded sensor at $p_4 = 7$ and player 5 also gets his demanded sensor for $p_5 = 5$. The new social welfare is $SW^H = 9 + 8 - 7 - 5 = 5 < 15 = SW^{SMS}$. This concludes our proof. $\square$

Claim 11 *The Hybrid mechanism is not strategyproof.*

The proof is by the following example.

Example 4. Suppose that we have 3 players with $v_1 = 4$, $v_2 = 3$, $v_3 = 2$ who demand only one tuple with cost $c = 7$. If these players declare their true values then,
${p_1}^{VCG} = 7 - (3 + 2) = 2$
${p_2}^{VCG} = 7 - (4 + 2) = 1$
${p_3}^{VCG} = 7 - (4 + 3) = 0$
So we compute the new cost, $c' = 7 - (2 + 1 + 0) = 4$ and we proceed with splitting equally the remaining cost. Now every player has to pay an additional amount of $p^E = 4/3 = 1.33$. So the final payments are,
${p_1}^H = 2 + 1.33 = 3.33$
${p_2}^H = 1 + 1.33 = 2.33$
${p_3}^H = 0 + 1.33 = 1.33$

Thus, all the players get the tuple under the Hybrid mechanism. Observe here is that if we run the SMS mechanism, no player gets any service, leading to a social welfare of zero while here we get the maximum possible social welfare of $(4+3+2)-7 = 2$. Now let us take a look at player 3 for example. When player 3 declares his true value, he has a utility of $u_3 = 2 - 1.33 = 0.67$. However, if he declares $b_3 = 0$, then we have bigger VCG payments for the rest of the players and no payment for him. In more detail,

${p_1}^{VCG} = 7 - (3 + 0) = 4$
${p_2}^{VCG} = 7 - (4 + 0) = 3$
${p_3}^{VCG} = 7 - (4 + 3) = 0$

The new remaining cost is now $c' = 7 - (4 + 3 + 0) = 0$ and hence, $p_i^H = p_i^{VCG}$. Thus, all players get again the demanded tuple, but player 3 has now a utility of $u_3 = 2 - 0 = 2 > 0.67$. We conclude that our mechanism is not strategy proof. $\square$

C Missing proofs from Section 4

Proof of Claim 6 : As we stated ${p_i}^H = {p_i}^{VCG} + p^E$. Since p^E is the same for all players, we only have to show that ${p_i}^{VCG} \geq {p_j}^{VCG}$. But this however holds trivially by how the per-round VCG payments are defined in our setting. $\square$

Proof of Lemma 2 : Consider the set A^r of currently active players just before we run the VCG mechanism at step 2b within round r. Let $|A^r| = k$, and let T be the set of players with $p_i^{VCG} > 0$ at that step. Suppose that $|T| = t$ with $t \leq k$. We will prove that $b_i \geq {p_i}^H$. We analyze first p_i^H for players with $p_i^{VCG} > 0$:

$$p_i^H = p_i^{VCG} + p_i^E = c_r - \sum_{j \neq i} b_j + \frac{c_r - \sum_i p_i^{VCG}}{k} \tag{8}$$

$$= c_r - \sum_{j \neq i} b_j + \frac{1}{k} \cdot \left(c_r - \sum_{\ell \in T} \left[c_r - \sum_{j \neq \ell} b_j\right]\right) \quad (9)$$

$$= c_r - \sum_{j \neq i} b_j + \frac{1}{k} \cdot \Big(\sum_{\ell \in T} \sum_{j \neq \ell} b_j - (t-1)c_r\Big) \quad (10)$$

We can now continue with what we want to establish as follows:

$$b_i \geq {p_i}^H \iff kb_i \geq (k-t+1)c_r - k\sum_{j \neq i} b_j + \sum_{\ell \in T} \sum_{j \neq \ell} b_j$$

$$\iff kb_i \geq (k-t+1)c_r - (k-1)\sum_{j \neq i} b_j + \sum_{\ell \in T \setminus i} \sum_{j \neq \ell} b_j$$

$$\iff (k-t+1)b_i \geq (k-t+1)c_r - (k-1)\sum_{j \neq i} b_j + \sum_{\ell \in T \setminus i} \sum_{j \neq \ell, i} b_j$$

$$\iff (k-t+1)b_i \geq (k-t+1)c_r - \sum_{j \in T \setminus i} b_j - (k-t)\sum_{j \neq i} b_j$$

$$\iff (k-t)\sum_{j \in A^r} b_j + \sum_{j \in T} b_j \geq (k-t+1)c_r$$

Note that if $k = t$, i.e., $A^r = T$, the inequality holds trivially, since we have assumed that $\sum_{j \in A^r} \geq c_r$. Suppose $t < k$. Define $b_q = max_{j \in A^r \setminus T}\, b_j$. We can now rewrite the last equivalence above as:

$$(k-t)\sum_{j \neq q} b_j + (k-t)b_q + \sum_{j \in T} b_j \geq (k-t)c_r + c_r$$

The final statement is true because, first:

$$(k-t)\sum_{j \neq q} b_j \geq (k-t)c_r$$

This is because $k > t$ and $\sum_{j \neq q} b_j \geq c_r$ since $p_q^{VCG} = 0$. Finally, we also have that $c_r \leq \sum_{j \in A^r} b_j \leq (k-t)b_q + \sum_{j \in T} b_j$. This concludes our proof. □

Proof of Lemma 3 : The proof is by induction on the number of rounds.
Induction basis: Consider the first round. Suppose that $N_1^S \neq \emptyset$ and let $i \in N_1^S$. We will show that $i \in N_1^H$. We have:

$$i \in N_1^S \Rightarrow b_i \geq c_1 / |N_1^S|$$

Let us examine the Hybrid mechanism now. Recall that each round of the Hybrid mechanism is divided into iterations and at most one player can be removed per iteration. We claim that i cannot be rejected by the Hybrid mechanism during the first $n - |N_1^S|$ iterations. To see this, consider one such iteration. If $p_i^{VCG} > 0$, then i is protected and cannot be removed by Lemma 2. Suppose then that i did not pay during the VCG step. Then, his cost-share will be c_1'/n', where n' is the number of people being active during the iteration. But $c_1' \leq c_1$ and $n' > |N_1^S|$, since we are within the first $n - |N_1^S|$ iterations. Thus, $p_i^H = c_1'/n' \leq c_1/|N_1^S| \leq b_i$, thus i is not removed. Therefore, either the first round of the Hybrid mechanism ends before $n - |N_1^S|$ iterations and $N_1^S \subset N_1^H$, or the mechanism executes exactly $n - |N_1^S|$ iterations and after that precisely the set N_1^S has survived. But by the same arguments again, all these players can afford their cost share in the next iteration and the mechanism will stop by allocating the tuple to them, in which case $N_1^S N_1^H$. This completes the basis.

Induction step: Suppose now that at round $r-1$ we have $N_{r-1}^S \subseteq N_{r-1}^H$. When round r starts, we possibly exclude players with $d_i = r-1$, but these are excluded from both sets. Hence at the beginning of round r, the Hybrid mechanism has a superset of active players with respect to SMS. Let now N_r^S and N_r^H be the players who survived round r in the two mechanisms respectively. Suppose that $N_r^S \neq \emptyset$, and consider $i \in N_r^S$. We can now use very similar arguments as in the induction basis to prove that $i \in N_r^H$. Again the idea is that during each iteration within round r, either $p_i^{VCG} > 0$ and hence i is not rejected, or otherwise, we can show that b_i is at least as big as the cost share in the first $n - |N_r^S|$ iterations. Therefore $i \in N_r^H$ and $N_r^S \subseteq N_r^H$. □

Theorem 12. *The Hybrid mechanism always achieves optimal social welfare when we have only $n = 2$ players.*

Proof of Theorem 12 : We have to prove that if $b_1 + b_2 \geq c_r$ for some round-tuple r, then both players get the tuple. In other words we must show that if $b_1 + b_2 \geq c_r$ then $b_1 \geq {p_1}^H$ and $b_2 \geq {p_2}^H$ for that round. So suppose that $b_1 + b_2 \geq c_r$ for a round r, then we have four possible cases that depend on the possible ${p_1}^{VCG}, {p_2}^{VCG}$ that might occur:

1. ${p_1}^{VCG} = {p_2}^{VCG} = 0$: This means that both $b_1, b_2 \geq c_r \geq c_r/2$. So ${c_r}' = c_r$ and ${p_1}^H = {p_2}^H = c_r/2 \leq b_1, b_2$. Thus both players get the tuple since their bid is bigger than the final price.
2. ${p_1}^{VCG} = 0$ and ${p_2}^{VCG} = c_r - b_1$: This means that $b_2 \geq c_r$. So ${c_r}' = c_r - (c_j - b_1) = b_1$, ${p_1}^H = b_1/2$ and ${p_2}^H = c_r - b_1 + b_1/2 = c_r - b_1/2$. However $b_2 - {p_2}^H = b_2 - c_r + b_1/2 \geq b_2 - c_r \geq 0$ and $b_1 - {p_1}^H = b_1/2 \geq 0$. Thus both players get the tuple since their bid is bigger than the final price.
3. ${p_2}^{VCG} = 0$ and ${p_1}^{VCG} = c_r - b_2$: Similarly, as in Case 2.
4. ${p_1}^{VCG} = c_r - b_2$ and ${p_2}^{VCG} = c_r - b_1$ (remember that $b_1 + b_2 \geq c_r$): So ${c_r}' = c_r - [(c_r - b_1) + (c_r - b_2)] = b_1 + b_2 - c_r$, ${p_1}^H = c_r - b_2 + (b_1 + b_2 - c_r)/2 = (c_r + b_1 - b_2)/2$ and ${p_2}^H = c_r - b_1 + (b_1 + b_2 - c_r)/2 = (c_r + b_2 - b_1)/2$. However $b_1 - {p_1}^H = b_1 - (c_r + b_1 - b_2)/2 = (b_1 + b_2 - c_r)/2 \geq 0$ and $b_2 - {p_2}^H = b_2 - (c_r + b_2 - b_1)/2 = (b_1 + b_2 - c_r)/2 \geq 0$. Thus both players get the tuple since their bid is bigger than the final price.

For $d_1 \neq d_2$ the statement holds trivially beyond the round defined by the lower demand. This concludes our proof. □

□

Proof of Lemma 8 : Suppose that $\sum_{j=1}^{n} b_j = c_1$. Initially notice that for every i we have that $b_i = c_1 - \sum_{j \neq i} b_j$ so it is easy to see that ${p_i}^H = {p_i}^{VCG} + {p_i}^E = b_i + 0 = b_i$, thus all players get their demanded tuple. Now if player i bids a lower amount, then the sum of bids becomes smaller from the total cost and no player (including i) gets the tuple. If player i bids a higher amount, say $b_i + a > b_i$, notice that his payment at the VCG step does not change but the VCG payments of the rest players are decreased. This, according to the structure of Hybrid mechanism, produces a higher egalitarian cost and thus a higher egalitarian payment for all players. So player i now pays ${p_i}^H = {p_i}^{VCG} + {p_i}^E = b_i + {p_i}^E$ where ${p_i}^E > 0$, something that leads to a lower utility. Our proof is complete.

The statement for (ii) is proved by using the same arguments. □

Proof of Theorem 9 : The proof of the Theorem is based on the following useful lemma, telling us what are the properties needed by an equilibrium profile.

Lemma 4. *Consider a set of n players with the same demand d, in addition $\sum_{i=1}^{n} v_i \geq c_r$ where $r \leq d$ is the maximum round that this condition holds and let $\delta = max\{c_i - c_{i-1}\}$ for $i \in [1, r]$. If there exist bids so that for every player i we have $v_i - (r-1)\delta \geq b_i > (r-1)\delta$ and in addition $\sum_{i=1}^{n} b_i = c_r$ then this vector of bids produces a NE that achieves the optimal social welfare.*

Proof. We have to prove that under this hypothesis, no player is excluded up to round r and in addition no player is motivated either to overbid or to underbid.

– No player is excluded up to round r: Basically if $b_i > (r-1)\delta$ and $\sum_{i=1}^{n} b_i = c_r$ then player i is pivotal at all rounds. We will proceed with induction.

  Induction basis: At round 1 an arbitary player i pays at the VCG step, ${p_i}^{VCG} = c_1 - \sum_{j \neq i} b_j = c_r - \sum_{z=2}^{r}(c_z - c_{z-1}) - \sum_{j \neq i} b_j \geq c_r - (r-1)\delta - \sum_{j \neq i} b_j = b_i - (r-1)\delta > 0$. Thus every player i is pivotal in round 1, so none of them is excluded.

  Induction step: Suppose that no player has been excluded up to round $k-1$. As a consequence round k starts with all n players and an arbitary player i pays at the VCG step, ${p_i}^{VCG} = c_k - \sum_{j \neq i} b_j = c_r - \sum_{z=k+1}^{r}(c_z - c_{z-1}) - \sum_{j \neq i} b_j \geq c_r - (r-k)\delta - \sum_{j \neq i} b_j = b_i - (r-k)\delta > 0$. Thus every player i is pivotal in round k, so none of them is excluded.

  The induction is complete and we can conclude that all players are pivotal at all rounds so no player is excluded up to round r.

– No player is motivated to either overbid or underbid: Inially, notice that as with the previous Theorem, when a player i overbids he has no gain since he will pay more at all rounds. His only chance to increase his total utility is to bid something higher from his value in order to gain access to round $r+1$. However in such a case he will increase his payment at all rounds up to r and in addition he will pay something higher from his value at round $r+1$. Thus there is no motivation in overbidding. Now if a player i underbids, he actually has a chance to decrease his per round payments but he also decreases the number of rounds that he participates (at least one round since $\sum_{i=1}^{n} b_i = c_r$). So if player i bids an amount $b_i - a < b_i$ and under his new bid can participate in $x = r - k$ rounds, then since the rest players are pivotal due to our hypothesis (and remain pivotal as the new bid of player i actually increase their VCG payments) the total gain that he has at all the $r-k$ rounds that he participates is $(r-k)[a(n-1)/n] < (r-k)a \le (r-k)\sum_{z=r-k+1}^{r}(c_z - c_{z-1}) \le (r-k)k\delta$. So we have to prove that,

$$u_i^r \ge u_i^x \iff rv_i - \sum_{j=1}^{r-1} p_j - b_i \ge (r-k)v_i - \sum_{j=1}^{r-k} p'_j$$

Notice that $(r-k)v_i - \sum_{j=1}^{r-k} p_j + (r-k)k\delta > (r-k)v_i - \sum_{j=1}^{r-k} p'_j$ so we can proceed instead to,

$$rv_i - \sum_{j=1}^{r-1} p_j - b_i \ge (r-k)v_i - \sum_{j=1}^{r-k} p_j + (r-k)k\delta$$

$$\iff kv_i - \sum_{j=r-k+1}^{r-1} p_j - b_i \ge (r-k)k\delta$$

In addition we have that $kv_i - \sum_{j=r-k+1}^{r-1} p_j - b_i \ge kv_i - \sum_{j=r-k+1}^{r-1} b_i - b_i = k(v_i - b_i)$ thus we can prove instead,

$$k(v_i - b_i) \ge (r-k)k\delta \iff v_i - b_i \ge (r-k)\delta \iff v_i - (r-k)\delta \ge b_i$$

Something that is true by our hypothesis (notice that minimum $k = 1$).

□

We now use the Lemma to prove the theorem. By the first condition of the theorem, we have that $v_i > 2(d-1)\delta \Rightarrow v_i - (d-1)\delta > (d-1)\delta \Rightarrow \exists b_i \in ((d-1)\delta, v_i - (d-1)]$. Adding the respective inequallities of all players, we have that $\sum_{i=1}^{n} b_i \in (n(d-1)\delta, \sum_i v_i - n(d-1)\delta]$. By the second condition, we have that c_d can be realized as a sum of appropriate bids by the players.

□

D Other interesting remarks

The next lemma highlights an important differentiation between the Hybrid and the SMS mechanism, and also yields some insight as to why the Hybrid mechanism can produce higher social welfare.

Lemma 5. *Let A^r be the initial set of active players at the beginning of an arbitrary round r. If $\sum_{j \in A^r} b_j \ge c_r$, then the Hybrid mechanism will always allocate the tuple of this round to some players (i.e., not all players are rejected during this round), unlike the SMS mechanism.*

Proof. The statement of this lemma implies that we can extract nonzero social welfare from such a round, unless it is the case that $\sum_j b_j = c_r$ for the surviving players in that round. In contrast, there are instances where the SMS mechanism can reject all players even when they can collectively cover the cost.

Before stating the formal proof, we first provide some intuition: First, note what we have from Lemma 2. If for a player i we have $p_i^{VCG} > 0$, then we know he is protected during that step and will not be rejected. Hence, we need only worry about players for which $p_i^{VCG} = 0$. But the second important observation now is that if the players can collectively cover the cost, and yet a player i pays zero at the VCG step, this means that the rest of the players can still cover the cost of c_r after excluding i. In fact even more players may become pivotal after excluding i, increasing in that manner the "protected" players.

The formal proof is by induction. At each round r, the Hybrid mechanism may run several iterations, where by an iteration, we refer to an execution of steps 2a till 2e. We prove by induction on the number of iterations, the following statement:
Consider a round r, where initially $\sum_{j \in A^r} b_j \geq c_r$. Then at the end of each iteration within round r,

- either we have $b_i \geq p_i^H$ for every active player i,
- or there is at least one player i, with $b_i < p_i^H$, and after applying step 2e and setting $A^r = A^r \setminus \{i\}$, we still have $A^r \neq \emptyset$ and $\sum_{j \in A^r} b_j \geq c_r$.

Induction basis: Suppose there is a player such that at the end of the first iteration, $b_i < p_i^H$ and that this is the player rejected from the mechanism. By Lemma 2, we know that $p_i^{VCG} = 0$. But this implies that $\sum_{j \in A^r \setminus \{i\}} b_j \geq c_r$. Since, the new set A^r at the end of the first iteration is precisely $A^r \setminus \{i\}$, this means the remaining players can still cover the cost. This completes the basis.
Induction step: Suppose that the induction hypothesis holds up until iteration k and consider the next iteration within round r. The proof is exactly by the same argument as for the induction basis. To see also why A^r can never become the empty set, suppose that as the algorithm keeps removing players, we reach a point where $|A^r| = 2$, consisting say of players 1 and 2, and that we still have $\sum_{j \in A^r} \geq c_r$. If at this point the mechanism removes one more player, say player 2, then since $p_2^{VCG} = 0$, this means that $b_1 \geq c_r$, implying that in the next iteration player 1 will cover the cost and will not be rejected.

Remark 2. Note that in the description of the Hybrid mechanism, at step 2e, we remove only one player, even if there can be more players with $b_i < p_i^H$. This turns out to be crucial for the properties we have established. Namely, in the proof of Lemma 5, we used the fact that *when a player pays zero at the VCG step, this means that the rest of the players can still cover the cost of the round*. If we remove more than one player at a time, the lemma does not hold any more.