

Redistribution in Online Mechanisms

Victor Naroditskiy, Sofia Ceppi*, Valentin Robu and Nicholas R. Jennings
University of Southampton
Southampton, UK

vn@ecs.soton.ac.uk,sc11v@ecs.soton.ac.uk,vr2@ecs.soton.ac.uk,nrj@ecs.soton.ac.uk

ABSTRACT

Following previous work on payment redistribution in static mechanisms, we develop the theory of redistribution in online mechanisms (e.g., [2, 10, 8]). In static mechanisms, redistribution is important as it increases social welfare in scenarios with no residual claimant. Many online scenarios also do not have a natural residual claimant, and redistribution there is equally important. In this work, we adopt a fundamental online mechanism design model where a single expiring item is allocated in each of T periods. Agents with unit demand are present in the market between their arrival and departure periods, which are private information along with the value an agent attributes to the item. For this model, we derive a number of properties characterizing redistribution in online mechanisms (including revenue monotonicity properties, and quantifying the effect an agent can have on the total revenue). We then design two redistribution functions. The first one generalizes the static redistribution proposed by Cavallo [2] making redistribution after the departure of the last agent. For this redistribution function we provide theoretical worst-case guarantees. The second function is truly “online” making redistribution to each agent on her departure. The performance of both functions is evaluated using numerical simulations.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

Keywords

Online mechanism design, payment redistribution

1. INTRODUCTION

Revenue redistribution is a growing area of study within mechanism design. Its importance can be intuitively illustrated by means of a simple example. Consider a situation in which a number of identical items need to be allocated among a group of agents (specific examples might include allocating free tickets for a popular talk, deciding which roommate gets to use the living room for a weekend party, or allocating university parking spots among faculty members). Each agent has a private value for the item, and we

*Double affiliation with Politecnico di Milano, Milan, Italy Email: ceppi@elet.polimi.it

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

would like to distribute the items in a way that maximizes social welfare. In order to incentivize agents to reveal their private values, payments must be introduced. Importantly, there is no revenue-maximizing auctioneer or *residual claimant* to absorb the revenue. Thus, any revenue collected represents the cost of truthfulness, and decreases the social welfare. It has been shown [6] that the cost cannot be zero: i.e., budget balance and allocative efficiency are not compatible. Against this background, the redistribution literature aims to distribute back as much of the revenue as possible.

Much of the work on redistribution mechanisms focuses on finding the best mechanism from the Groves class: i.e., on redistributing VCG payments (e.g., [10, 8, 11]). But some work has addressed non-efficient mechanisms [7, 4]. However, with the exception of [3] discussed in Section 7, all of the redistribution results assume static settings, in which the decisions are made at the same time in the presence of all participating agents. While relevant to some settings, in others like electric vehicle charging and allocation of computational resources in cloud computing this is not a suitable model.

To this end, we provide the first results on redistribution in *online* mechanisms. Specifically, we consider the case in which decisions must be made over time, with a separate decision made each period, and agents who arrive and depart at various times not known by the mechanism, i.e., the private information of each agent also includes her arrival and departure times.

As in the prior work on static models, our goal is to maximize social welfare. Welfare maximization is a natural objective for allocating resources in situations without a revenue-maximizing auctioneer. An example of an online setting where social welfare is the right objective is electric vehicle charging [5, 14], where vehicles, arriving and departing at different times of day, draw electricity from a shared resource such as a community-owned wind turbine or need to divide between them a joint quota made available by the electricity distribution company. Another example is cloud computing, in which computational jobs arriving over time need to be allocated to a number of processors.

In our work we consider the fundamental model where identical items are distributed among agents with unit demand [13]. In particular, we focus on deterministic, individually rational (i.e., each agent should not be worse off after participating in the mechanism), and weakly budget-balanced (i.e., the total payment collected from the agents should be non-negative) mechanisms where truthful reporting is a dominant strategy. We refer to the latter property as *dominant strategy incentive compatibility* (DSIC). A class of online mechanisms satisfying these properties has been described in [13], and we study redistribution within this class. This simple model proves to be a good departure point for the study of redistribution in online mechanisms. The model includes the properties specific

to online settings such as: (i) arrival and departure times are private information of the agents, (ii) in each period no information about agents arriving in the future is available, and (iii) the allocation decisions in each period are interdependent. We start with deriving a number of general properties for this online setting (see Section 3).

Based on the general properties, we design two redistribution functions. One is a generalization of the function proposed for static settings by Cavallo [2]. Under this rule, redistribution occurs only after the last agent departs. We provide analytical guarantees of the performance of the generalized redistribution. The second function, redistributes to each agent on her departure. Performance of both functions is evaluated in numerical simulations.

In summary, the main contributions of this work are as follows:

- We derive general results characterizing properties of redistribution functions in online domains that are required to guarantee dominant strategy truthfulness and weak budget balance.
- Based on these general properties, we design two redistribution functions for online settings: one that redistributes to all agents at the last period, and one that redistributes to each agent at her departure time. Moreover, we provide theoretical guarantees on the performance of the first redistribution function.
- We evaluate the performance of both functions in redistributing collected revenue using numerical simulations.

The remainder of the paper is organized as follows. First, we describe the online mechanism design model we adopt from [13]. We proceed with a series of results pointing out the features that an online redistribution function must satisfy to guarantee weak budget balance and DSIC. Then, we propose two redistribution functions analyzing them in terms of the percentage of revenue redistributed. Finally, we present the results of numerical simulations of the functions described.

2. MODEL OF ONLINE MECHANISMS

Existing literature on online mechanism design discusses several models of online allocation (see [13] for an overview). We focus on a fundamental model proposed in [9, 13]. Specifically, we study the class of *deterministic, model-free* online mechanisms. In such mechanisms, the allocation rule itself is deterministic, and the mechanism does not need a model of future arrivals of the agents in order to compute the allocation. Furthermore, in this work we only consider online mechanisms where truth-telling (i.e., true reporting of types by the agents) is a dominant strategy. Deterministic, model-free, dominant-strategy mechanisms are the most desirable mechanisms to design as they require no prior information on agents' types, and do not make any assumptions about risk-preferences of the agents.

Formally, there are T discrete time periods, and agents may arrive and depart within $[1, T]$. There is an identical item available for allocation in each period.¹ The items are "expiring", and if not allocated within their period, they disappear (this is natural, for example, when items correspond to computational time on a machine). We define the type of agent i as $\theta_i = (a_i, d_i, v_i)$, where a_i is her arrival period, d_i is her departure period ($1 \leq a_i \leq d_i \leq T$), and $v_i \in \mathbb{R}$ is her value for obtaining the item. We sometimes refer to the interval $[a_i, d_i]$ when agent i is present as agent i 's *active window*. We use N to denote the set of agents

¹The case of multi-unit supply per period is not discussed in this paper, but our model can be extended to cover this case.

that are present at some time between 1 and T , with the cardinality denoted by $n = |N|$. We denote² by $\pi^t : \theta \rightarrow \{0, 1\}^n$ the *greedy* allocation function [13] that, at each time t , allocates the good to the previously unallocated agent who is present at time t and has the highest value among all unallocated agents present at t . Allocation to agent i is specified by $\pi_i^t(\theta) \in \{0, 1\}$. Formally, $\pi_i^t(\theta) = 1$ if $i \in N^t$ and $i = \arg \max_{j \in N^t} v_j$ where $N^t = \{j \in N \mid t \in [a_j, d_j] \text{ and } \sum_{t'=a_j}^{t-1} \pi_j^{t'}(\theta) = 0\}$ denotes the set of agents active on day t . If there are multiple active agents with the same highest value, the allocated agent is determined randomly.

The utility of agent $i \in N$ from participating in the market is $u_i(\theta) = v_i \pi_i(\theta) - x_i(\theta)$ where $x_i(\theta)$ denotes the payment of agent i (defined in Equation 1 and generalized in Equation 3), and $\pi_i(\theta)$, defined as $\sum_{t=a_i}^{d_i} \pi_i^t(\theta)$ indicates whether agent i has been allocated ($\pi_i(\theta) = 1$) or not ($\pi_i(\theta) = 0$).

We remark that, for this deterministic, single-unit demand setting, the state of the art characterization was first presented by Hajiaghayi *et al.* [9] (and, in a more extended form, by Parkes [13]). Assuming individual rationality and zero payment from unallocated agents, they show that the allocation function π^t can be truthfully implemented in single-valued domains with *limited mis-reports*³ (i.e., no early arrivals/late departures can be reported) and only if the payment x_i of each agent $i \in N$ takes the form:

$$x_i(\theta) = \hat{x}_i(\theta) = \begin{cases} v^c(\theta_{-i}, a_i, d_i) & \text{if } \pi_i(\theta) = 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $v^c(\theta_{-i}, a_i, d_i)$ denotes the *critical value* of agent i :

$$v^c(\theta_{-i}, a_i, d_i) = \min_{v'_i \in \mathbb{R}} v'_i \mid \pi_i(\theta'_i, \theta_{-i}) = 1, \text{ for } \theta'_i = (v'_i, a_i, d_i) \quad (2)$$

We dispose of the assumption that unallocated agents' payment is zero, and characterize all possible ways to modify the payment function above:

$$x_i(\theta) = \hat{x}_i(\theta) - h(\theta_{-i}, a_i, d_i) \quad (3)$$

where $h(\theta_{-i}, a_i, d_i)$ is the redistribution agent i receives.

In the next section we discuss how redistribution should be defined in order for the allocation mechanism to maintain DSIC and weak budget balance. Note that, individual rationality—the property that each agent has a non-negative utility—is satisfied by the mechanism described above if and only if the redistribution is non-negative. This will be the case throughout the paper.

The last part of the model that needs to be specified is the evaluation metric. As in much of the work on the static case (e.g., [2, 10, 8]), we evaluate mechanisms based on the worst-case performance guarantee: i.e., the welfare that is guaranteed regardless of agents' private information.

In order to better explain our results, we benchmark them against the existing results for the static case. The relevant static case is allocating m identical items among n agents (by definition, there is only one period $T = 1$ in the static case). Note that in the online case, since we consider the scenario with single unit supply (i.e., only one item can be allocated in each time interval), the number of items is the same as the number of periods $m = T$.

²To avoid complicated notation, we use types θ of all agents N as an argument to the allocation function π and of all agents except i , $N \setminus i$, to the payment function x_i . However, allocation at period t is decided based only on the types of the agents that already arrived in the market. The types of agents that have not yet arrived are not used (and, in fact, cannot be known) by these functions.

³We adopt this assumption throughout the paper.

The worst-case ratio for allocating m items among n agents is measured as the percentage of revenue that is guaranteed to be redistributed back to the agents regardless of their types:

$$r(m, n) = \min_{\theta} \frac{H(\theta)}{R(\theta)} \quad (4)$$

where $R(\theta) = \sum_{i \in N} x_i(\theta)$ is the collected revenue, and $H(\theta) = \sum_{i \in N} h(\theta_{-i})$ is the total amount redistributed. Note that the best possible ratio of any mechanism satisfying weak budget balance is one (corresponds to a fully budget-balanced mechanism).

In the static case of allocating identical items to agents with unit demand, the private information of an agent is her value for an item $\theta_i = v_i$. The ratio of various redistribution mechanisms can be described as a function of m and n . For instance, as shown in [8], the Bailey-Cavallo redistribution mechanism [2, 1] achieves the ratio of:⁴

$$r^{\text{BC}}(m, n) = \frac{n - m - 1}{n} \quad (5)$$

In the online case, each agent's private information also includes the arrival and departure dates, $\theta_i = (a_i, d_i, v_i)$. In the next section, we show that it is more difficult to provide positive results for the online case and show how to remedy this in Section 4.

3. CHARACTERIZING REDISTRIBUTION IN ONLINE MECHANISMS

We start with characterizing the class of DSIC redistribution mechanisms.

LEMMA 1. *Under limited misreports, the mechanism specified by the greedy allocation function π^t is DSIC if and only if the payment function is given by Equation 3 where $h(\theta_{-i}, a_i, d_i)$ is an arbitrary function satisfying*

$$h(\theta_{-i}, a_i, d_i) \geq h(\theta_{-i}, a'_i, d'_i) \quad \forall \theta_{-i}, a'_i \geq a_i, d'_i \leq d_i \quad (6)$$

PROOF. First, we show the “if” part: any payment rule (3) with a redistribution function $h(\theta_{-i}, a_i, d_i)$ satisfying (6) is DSIC. Recall that payment rule without redistribution (1) is DSIC [9, 13] and note that the redistribution is independent of agent's reported value v_i . The only manipulation an agent has is with respect to a_i and d_i . The condition in (6) guarantees that an agent never gets a higher redistribution by reporting a later arrival or an earlier departure. By assumption of limited misreports, agents cannot report an earlier arrival or a later departure. Thus, the agents have no incentive to misreport due to the redistribution term, and the mechanism remains DSIC.

Next, we prove the “only if” part: only payment rules stated in the lemma truthfully implement the greedy allocation π^t . For this, we need to argue that in a DSIC mechanism h must be independent of v_i and must satisfy (6). To argue independence of v_i we invoke results for static allocation of m items among n agents, which is a special case of our online mechanisms when arrival and departure dates for all agents are the same. DSIC in the static case (e.g., see [12], Theorem 9.36) states that the difference between the payment of agent i when she is allocated and when she is not allocated (i.e., the price she pays for the item) must be exactly the critical value $v^c(\theta_{-i}, a_i, d_i)$ defined in Equation 2. Thus, redistribution to agent i must be independent of v_i . Finally, we show that DSIC requires (6). We just argued that redistribution to agent i is independent of v_i and consequently of her allocation. Suppose in violation

⁴The allocation rule is set to the efficient allocation function in the Bailey-Cavallo mechanism.

of (6), agent i 's rebate is higher when she reports a later arrival or an earlier departure. Agent i with the true value of zero (i.e., agent who is never allocated, or never pays for being allocated), would want to lie about her arrival/departure violating DSIC. \square

Next we address the issue of weak budget balance in online mechanisms. Denote by $R(\theta)$ the total revenue collected by the mechanism before any redistribution occurs: i.e., when $h(\theta_{-i}, a_i, d_i) = 0, \forall i$. Similarly, denote by $R(\theta_{-i})$ the total revenue the mechanism collects, before redistribution, if agent i were not present.

LEMMA 2. *Given reports of agents other than i , θ_{-i} , let θ'_i denote the report of agent i that minimizes the revenue $R^{\min}(\theta_{-i}) = \min_{\theta'_i \in \Theta} R(\theta'_i, \theta_{-i})$. Redistribution to agent i cannot exceed $R^{\min}(\theta_{-i})$ in a weakly budget-balanced mechanism.⁵*

PROOF. Suppose $h(\theta_{-i}, a_i, d_i) > R^{\min}(\theta_{-i})$. Recall that, in order to maintain truthfulness, the redistribution to agent i must be independent of θ_i . When agent i has the type θ'_i , the revenue collected is exactly $R^{\min}(\theta_{-i})$. Thus, given $h(\theta_{-i}, a_i, d_i) > R^{\min}(\theta_{-i})$, the weak budget balance property is violated as agent i receives more than the collected revenue. \square

Lemma 2 provides an upper bound on the amount that can be redistributed to each agent. We would like to redistribute as much as possible, and are looking for ways to guarantee at least a certain percentage of revenue is redistributed: i.e., we would like to bound redistribution from below. One of the difficulties in doing this is related to non-monotonicity of revenue described next.

LEMMA 3. *The revenue does not monotonically increase with the number of agents.*

PROOF. The proof is by counterexample. Consider case (a) in Figure 1. There are three agents, all present only at t_1 . Given the values reported in the figure, agent 1 is allocated and the total revenue is $R(\theta) = 10$. Now, consider case (b) in Figure 1 in which a new agent, agent 4, whose active window is (t_1, t_2) is added. Given the values reported in the figure, the only agent allocated is agent 4 at t_1 . Since her critical value is zero, the total revenue is $R(\theta) = 0$. Thus, even though the number of agents increased, the total revenue actually decreased. \square

We observe that while revenue does not monotonically increase in n , the critical value of each agent does.

OBSERVATION 1. *Critical value of agent i is non-decreasing in the number of other agents. This holds for all i .*

This observation may seem surprising as, at first, it appears to contradict Lemma 3. However, while the critical values of all agents are monotone in the number of agents, allocation decisions are made based on values of the agents rather than their critical values. It is possible that an agent with the highest value has the lowest critical value because it is patient (i.e., has a late departure date and can wait for a less competitive period). Adding such an agent to an existing set of agents would decrease the collected revenue, but cannot decrease the critical values of other agents. Figure 1 illustrates this.

Our first result on providing a worst-case performance guarantee is negative, showing that for any m and n , there are types of agents such that for each of them $R^{\min}(\theta_{-i}) = 0$. This is in contrast to the static case of allocating identical items where, a non-zero ratio could be achieved for $n > m + 1$ for all agent values (see Equation 5). We show this through the following lemma.

⁵This result is a generalization of Lemma 3 derived in [2] for static mechanisms.

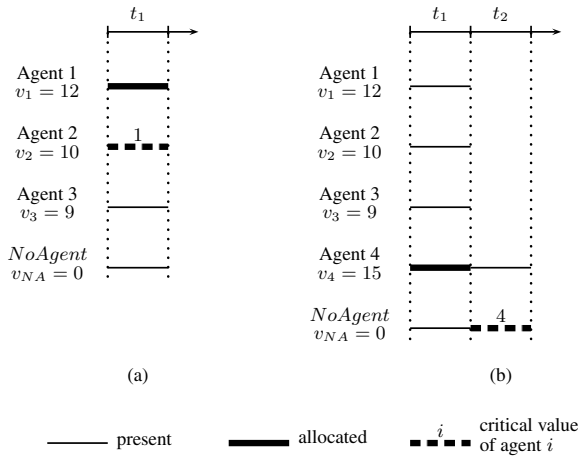


Figure 1: Example of non-monotonicity of revenue in the number of agents.

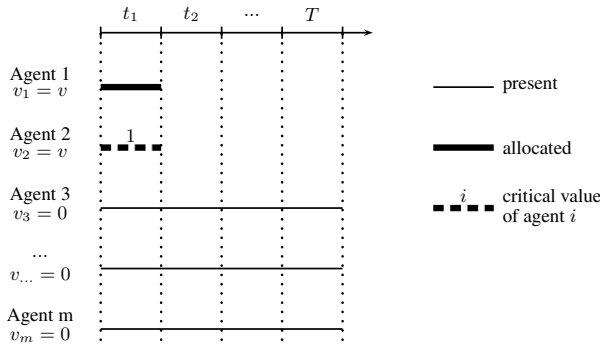


Figure 2: Problem instance in which the worst-case ratio is zero independently of m and n .

LEMMA 4. *The worst-case ratio of an online mechanism is zero for any m and n .*

PROOF. Consider the example depicted in Figure 2 with $T > 1$. Let the first two agents be present only in period t_1 and have the value of v , while the other $n - 2$ agents have the value of 0. The revenue for this problem instance is v : one of the first two agents is allocated and pays the value of the other. The corresponding social welfare before redistribution is zero. We argue that this cannot be improved through redistribution. By Lemma 2, agent 1 can receive no redistribution as the revenue is zero when she reports 0. The same for agent 2. Next, exploiting non-monotonicity of revenue, we show that, for each agent $i > 2$, there is a type θ'_i that results in zero revenue, i.e. $R^{\min}(\theta_{-i}) = R(\theta'_i, \theta_{-i}) = 0$. Suppose the type reported by agent i is $\theta'_i = (t_1, T, v')$, where $v' > v$. By the greedy allocation rule, agent i will be allocated on day one, but will pay 0, resulting in zero revenue. Thus, by Lemma 2, agents $i > 2$ cannot receive any redistribution either. \square

As the lemma above showed, one can construct a type profile θ where $R^{\min}(\theta_{-i})$ is zero for every agent. However, for value profiles θ where revenue comes from multiple agents, not every agent can cancel the entire revenue. We would like to bound exactly how much the report of agent i can decrease the revenue $R(\theta)$. This turns out to be difficult as the next example illustrates.

EXAMPLE 1 (CASCADE). *There are $T + 1$ agents. Agent 1 has the window $[1, 1]$ and the value v , each agent $2 \leq i \leq T$ has the window $[i - 1, i]$ and the value $v - (i - 1)\epsilon$, for some small ϵ , and agent $T + 1$ has the window $[T, T]$ and the value $v - T\epsilon$. Each of the agents $1 \dots T - 1$ is allocated in its departure period and pays the value of the next agent: $v - i\epsilon$. Ignoring the ϵ 's, the total revenue from these agents is $R(\theta) = Tv$. If agent 1 reports the value of 0 (or, equivalently is not present in the market), each of the agents $2 \dots T + 1$ is allocated in his arrival period and pay zero. Generalizing this, if agent i is dropped, the agents $1 \dots i - 2$ each pay v while the agents $i - 1, i + 1 \dots T$ pay nothing. In Figure 3, an example of the cascade is proposed for $T = 4$. In particular, part (a) shows the allocation when agent 1 is in the market and part (b)—when agent 1 is not in the market.*

In the example above, agent $i \geq 2$ cancels a fraction of roughly $1 - \frac{i-2}{T}$ of the revenue. Thus, providing a uniform non-zero bound on $R(\theta) - R^{\min}(\theta_{-i})$ that holds for all agents i is not possible. However, we can bound the difference between $R(\theta_{-i})$ and $R^{\min}(\theta_{-i})$. The difference can be positive due to non-monotonicity of revenue. Intuitively, this difference is the amount by which introducing agent i can reduce the revenue collected had agent i not been present. To bound it, we first characterize the effect a new agent has on the allocation decisions.

LEMMA 5. *Introducing agent i with type θ_i , which is allocated in the market with agents $N \setminus i$, affects the allocation of agents $N \setminus i$ only by forcing out one of the previously allocated agents⁶*

$$W(\theta_i, \theta_{-i}) = W(\theta_{-i}) \cup \{i\} \setminus \{j\}$$

where $W(X)$ refers to the set of allocated agents when the set of agents with types X are in the market, and $j \in W(\theta_{-i})$ is the agent that is forced out by i .

PROOF. Agents $W(\theta_{-i})$ represent the most “competitive” agents in the market θ_{-i} . From the point of view of agent $l \neq i$, arrival of agent i causes the number of available items to decrease by one and the critical values to weakly increase (Observation 1). Suppose for contradiction that agent $j' \notin W(\theta_{-i})$ is allocated. The contradiction is immediate, as j' must have been allocated in the less competitive market θ_{-i} as well. Thus, the only agents allocated after the arrival of i , are the agents from $W(\theta_{-i})$. Since the number of available items decreases by one, all agents $W(\theta_{-i})$ except one (who we call j) remain allocated. \square

Let \bar{v} denote the highest value an agent can have.

LEMMA 6. *The minimum revenue agent i can induce with her report is bounded by $R^{\min}(\theta_{-i}) \geq R(\theta_{-i}) - \bar{v}$.*

PROOF. We saw in Lemma 3 that an agent can reduce the revenue collected had she not been present. In order to influence $R(\theta_{-i})$, agent i must submit a report that affects either the allocation or the critical values or both. If agent i 's report sets a critical value for one of allocated agents but agent i remains unallocated, the revenue increases. Notice that in this case, the allocation does not change for any of the agents, but one of the agents pays a higher price—the critical value provided by agent i . Therefore, agent i may be able to decrease the revenue only if she is allocated.

To prove the bound, we use Lemma 5. When agent i is allocated, the set of other allocated agents remains the same except for one

⁶We implicitly assume that there are no items that are unallocated due to lack of demand. In that case, agent i may claim the unallocated item without forcing anyone out. However, this does not affect any of the results in the paper.

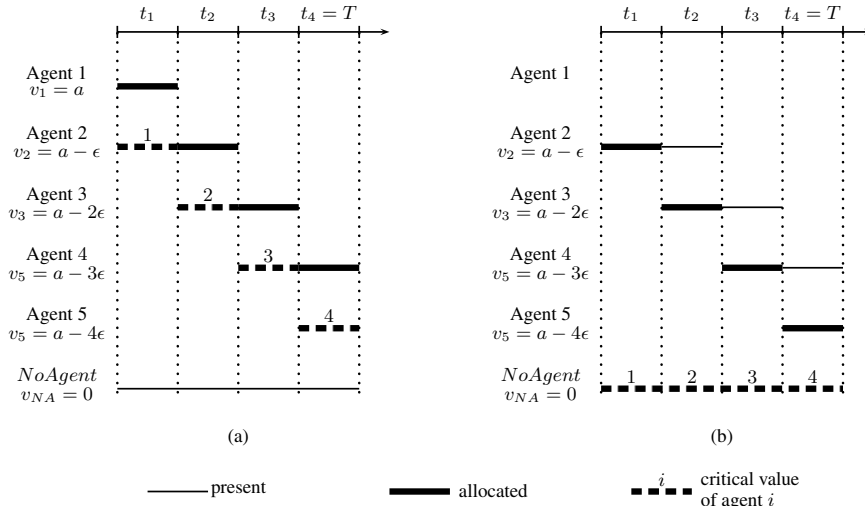


Figure 3: A cascade example. (a) $R(\theta) \approx 4v$ when agent 1 is in the market, (b) $R(\theta_{-1}) = 0$ when agent 1 is not in the market.

agent that is forced out. By Observation 1, the critical values of those agents cannot decrease with the appearance of agent i . Thus, the amount collected from all agents but one is at least as high as it was before agent i showed up. Agent i can reduce the revenue, by forcing out an agent who paid a lot, and paying less herself. The difference is the biggest when the forced out agent has the critical value \bar{v} and agent i has the critical value of zero. The claim in the lemma follows immediately. \square

The discussion in this section focused on general properties of online redistribution mechanisms. In the next two sections, we apply these intuitions to design specific redistribution rules.

4. REDISTRIBUTION AT THE LAST PERIOD

Based on the mechanism of Cavallo [2] and using results from the previous section, we can now design a redistribution mechanism. Let us consider redistribution at period T . By then, we know the entire set of agents N and their types. Our mechanism uses the greedy allocation rule and the following redistribution function:

$$h^T(\theta_{-i}) = \frac{R^{\min}(\theta_{-i})}{n} \quad (7)$$

where $R^{\min}(\theta_{-i})$ is defined as

$$R^{\min}(\theta_{-i}) = \min_{\theta'_i} R(\theta'_i, \theta_{-i}) = \min_{\theta'_i} \sum_{j \in N | \pi_j(\theta'_i, \theta_{-i})=1} v^c((\theta'_i, \theta_{-i})_{-j}, a_j, d_j)$$

Note that the redistribution function is from the class of DSIC functions characterized in Lemma 1. Also, observe that this mechanism is a generalization of Cavallo's redistribution [2], which is recovered by setting $T = 1$.

We remark that $R^{\min}(\theta_{-i})$ is given by one of the following two cases: (1) agent i reporting zero and $R^{\min}(\theta_{-i}) = R(\theta_{-i})$ as in Example 1; or (2) agent i reporting a high value to take an item away from a high-paying agent, and a long enough windows to have a low payment as in the example in Figure 1.

Next, we show that the mechanism is weakly budget-balanced, and prove performance guarantees.

THEOREM 1. *The redistribution function given in Equation 7 is weakly budget-balanced.*

PROOF. The proof is immediate from the definition of redistribution. For each agent j , we have

$$\frac{1}{n} \sum_{i \in N | \pi_i(\theta)=1} v^c(\theta_{-i}, a_i, d_i) \geq \frac{1}{n} R^{\min}(\theta_{-j})$$

Adding up the inequalities for all agents, we obtain the result: total revenue collected $R(\theta)$ (left-hand side in Equation (8)) is above the total redistribution (right-hand side):

$$R(\theta) = \sum_{i \in N | \pi_i(\theta)=1} v^c(\theta_{-i}, a_i, d_i) \geq \frac{1}{n} \sum_j R^{\min}(\theta_{-j}) \quad (8)$$

\square

Recall from Lemma 4 that one can construct a problem instance with a zero ratio for any m and n . However, we can provide positive results when enough agents are present and the total revenue collected is higher than the highest value \bar{v} a single agent may have.

THEOREM 2. *The worst-case ratio guaranteed by the redistribution function in Equation 7 is:*

$$r(m, n, \bar{v}) \geq \frac{n-2T}{n} (R(\theta) - \bar{v})$$

PROOF. Since there are T items, at most T agents can be allocated. Payment of each allocated agent j is $v^c(\theta_{-j}, a_j, d_j) = \min_{i \in N \setminus \{j\} | t_i(\theta_{-j}) \in \{a_j, \dots, d_j\}} v_i$ where $t_i(\theta_{-j})$ denotes the day when agent i is allocated in the market with agents θ_{-j} or 0 if agent i is not allocated. Clearly, each payment is determined by exactly one agent. Thus, we have at most T different agents setting payments for the allocated agents. In total, we have at most $2T$ agents that are either allocated or set payments.

Removing any of the other $n-2T$ agents, does not affect the revenue, and for any such agent j , $R(\theta_{-j}) = R(\theta)$. From Lemma 6, $R^{\min}(\theta_{-j}) \geq R(\theta_{-j}) - \bar{v}$. Thus, for each of the $n-2T$ agents, the amount redistributed is at least $\frac{1}{n}(R(\theta_{-j}) - \bar{v}) = \frac{1}{n}(R(\theta) - \bar{v})$. The total amount redistributed to these agents is $\frac{n-2T}{n}(R(\theta) - \bar{v})$. The amount redistributed to any of the remaining $2K$ is non-negative, thus at least $\frac{n-2T}{n}(R(\theta) - \bar{v})$ is redistributed. \square

The performance guarantee in Theorem 2 is not tight, but it shows that the mechanism is asymptotically optimal.

OBSERVATION 2. *The redistribution function is asymptotically optimal as $n \rightarrow \infty$: for a fixed T , the undistributed revenue goes to zero as the number of agents in the market goes to infinity.*

Redistributing at some last time period may appear to be a significant limitation. However, it fits naturally with many settings encountered in practice, where there is a natural maximal window in which the online redistribution takes place (such as a day). Consider, for example, the problem of allocating electricity charging slots to a set of electric vehicles, discussed in [5, 14]. The electricity supply is scarce only during some interval during each day (e.g. 6am to 11 pm), and we only need to design an online mechanism that allocates in this high-demand interval. Nonetheless, in other applications where online redistribution can be applied there is not such a natural maximal window, and in the next section we present and analyze a function that redistributes continuously, on each agent's departure.

5. REDISTRIBUTION AT EACH AGENT'S DEPARTURE

Previous section described a redistribution mechanism that waits until all agents depart before computing redistribution. It is more desirable (but also more difficult) to provide redistribution sooner. Here we design a function that provides redistribution to each agent on her departure. The main difficulty in designing such a function is ensuring weak budget balance at each period. Specifically, we need to guarantee that the amount of revenue redistributed up to each period, does not exceed the revenue collected until then. However, in determining the part of the revenue that can be redistributed to agent i , DSIC requires us to only include the revenue collected independently of agent i 's report. Furthermore, unlike the static case and the mechanism from the previous section, the number of agents that divide the same revenue is not clear.

We approach this problem by "charging" the rebate of agent i against payments of the agents that contribute to it. Specifically, for agent i we define the set of agents to_i that give part of their payment to agent i . To maintain truthfulness, their payment is computed assuming agent i is not present. Similarly, we specify the set of agents from_j that receive part of agent j 's payment.

We propose a redistribution function, under which agent i receives on d_i the following redistribution:

$$h^d(\theta_{-i}, a_i, d_i) = \left(\sum_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \right) - \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \quad (9)$$

where $\text{to}_i = \{k \in N \mid d_k \in [a_i, d_i]\}$ and $\text{from}_i = \{k \in N \mid d_i \in [a_k, d_k]\}$.

Essentially, this means that agent i gets redistribution from agents who depart when agent i is present. Symmetrically, the payment of agent j is redistributed among the agents who are present when agent j departs. We subtract the max term in order to ensure weak budget-balance in the face of non-monotonicity of revenue. Note that the rule is DSIC as it satisfies the conditions of Lemma 1.

Intuitively, note that the size of set from_j is always known on the departure of agent j from the market, because all agents $i \in \text{from}_j$ (i.e. all agents that must receive a partial redistribution from agent j) must have already arrived in the market by d_j (because $a_i \leq d_j$). Moreover, the number of agents in the set from_j is also known at the time in which agents i get redistribution, i.e., at their departure, because $d_i \geq d_j$.

The following theorem characterizes the function in terms of the crucial property of weak budget balance.

THEOREM 3. *The redistribution function in Equation 9 is weakly budget-balanced.*

PROOF. Recall \hat{x}_i defined in Equation 1.

$$\begin{aligned} \sum_i h^d(\theta_{-i}, a_i, d_i) &= \sum_i \sum_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} - \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \\ &= \sum_j \sum_{i \in \text{from}_j} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} - \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \end{aligned} \quad (10)$$

$$= \sum_j \sum_{\substack{i \in \text{from}_j \\ \pi_j(\theta_{-i})=1}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} - \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \quad (11)$$

$$\begin{aligned} &= \sum_j \sum_{\substack{i \in \text{from}_j \\ \pi_j(\theta_{-i})=1 \\ \pi_j(\theta)=1}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} + \sum_j \sum_{\substack{i \in \text{from}_j \\ \pi_j(\theta_{-i})=1 \\ \pi_j(\theta)=0}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \\ &\quad - \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \\ &\leq \sum_j \hat{x}_j(\theta) + \sum_j \sum_{\substack{i \in \text{from}_j \\ \pi_j(\theta_{-i})=1 \\ \pi_j(\theta)=0}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} - \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \quad (12) \\ &\leq \sum_j \hat{x}_j(\theta) = R(\theta) \end{aligned}$$

Equation 10 follows from the definition of the redistribution function: amount redistributed to each agent comes from the payments of other agents. We obtain Equation 11, by observing that $\hat{x}_j(\theta_{-i})$ is positive only if agent j is allocated in the market with agents θ_{-i} . To prove Equation 12 observe that when agent j is allocated regardless of the presence of i (i.e., $\pi_j(\theta_{-i}) = \pi_j(\theta) = 1$) the payment of j is his critical value, which can only increase with the presence of agent i : $\hat{x}_j(\theta_{-i}) = v^c((\theta_{-i})_{-j}, a_j, d_j) \leq v^c(\theta_{-j}, a_j, d_j) = \hat{x}_j(\theta)$. Finally, we prove the last inequality below. We want to show that

$$\sum_j \sum_{\substack{i \in \text{from}_j \\ \pi_j(\theta_{-i})=1 \\ \pi_j(\theta)=0}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \leq \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \quad (13)$$

We first show that for every i there is at most one agent j such that $\pi_j(\theta_{-i}) = 1$ and $\pi_j(\theta) = 0$. First, note that $\pi_j(\theta_{-i}) = 1$ and $\pi_j(\theta) = 0$ imply $\pi_i(v) = 1$: if agent i had not been allocated, then excluding him, would not have changed the allocation, and agent j would have remained unallocated. By Lemma 5, the allocation when agent i enters the market changes by forcing exactly one agent out. Thus, there is only one agent $j \mid \pi_j(\theta_{-i}) = 1, \pi_j(\theta) = 0$, and we obtain

$$\begin{aligned} \sum_i \max_{j \in \text{to}_i} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} &\geq \sum_i \sum_{\substack{j \mid \pi_j(\theta_{-i})=1, \\ \pi_j(\theta)=0}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \\ &= \sum_j \sum_{\substack{i \mid \pi_j(\theta_{-i})=1, \\ \pi_j(\theta)=0}} \frac{\hat{x}_j(\theta_{-i})}{|\text{from}_j|} \end{aligned}$$

The last equality holds as the summations on either side exhaustively cover all non-zero terms $\hat{x}_j(\theta_{-i})$. \square

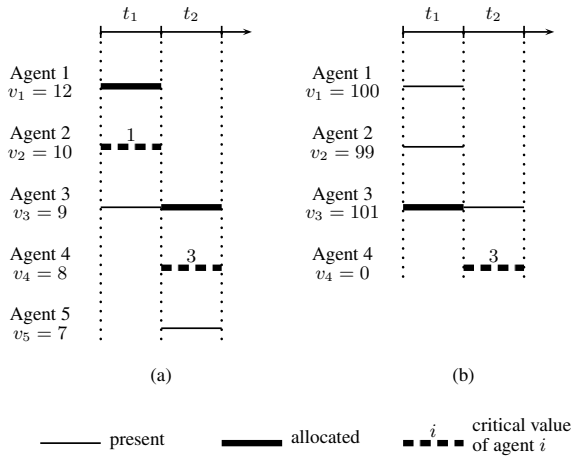


Figure 4: (a) Example of the “distribution-on-departure” rule, (b) Example that is not budget-balanced without subtracting the second term in Equation 9.

We illustrate the redistribution function using the example in part (a) of Figure 4. The only agent that departs when agent 1 is still present (actually, at the same time as agent 1 departs) is agent 2. Thus, $to_1 = \{2\}$. The payment of agent 2 is distributed among agents $from_2 = \{1, 3\}$ present at her departure. Thus, agents 1 is not present, agent 2 is allocated and pays 9). Since part of the payment of agent 2 will also go to agent 3, agent 1 only gets half of agent 2’s payment $h(\theta_{-1}, 1, 1) = \frac{9}{2} - \frac{9}{2}$. Similarly, agent 2 receives redistribution only from agent 1, who pays 9 when agent 2 is not there and $h(\theta_{-2}, 1, 1) = \frac{9}{2} - \frac{9}{2}$. Agent 3 gets redistribution from agents 1, 2, 4, and 5: $h(\theta_{-3}, 1, 2) = \frac{10}{2} + \frac{0}{2} + \frac{7}{2} + \frac{0}{2} - \frac{10}{2}$, etc.

We demonstrate that the second term in Equation 9 is necessary to guarantee budget balance. Consider the instance illustrated in part (b) of Figure 4. The revenue collected is $R(\theta) = 0$. The redistribution agent 3 gets when only the positive part of Equation 9 is considered is $h^d(\theta_{-3}, 1, 2) = \frac{99}{2}$ violating budget balance.

Finally, we apply the redistribution rule to the static case and compare it to the mechanism of Cavallo in the next observation.

OBSERVATION 3. *When applied to the static case, the redistribution function in Equation 9 becomes*

$$\begin{aligned}
 h^d(\theta_{-i}, a_i, d_i) &= \left(\sum_{j \in to_i} \frac{\hat{x}_j(\theta_{-i})}{|from_j|} \right) - \max_{j \in to_i} \frac{\hat{x}_j(\theta_{-i})}{|from_j|} \\
 &= \frac{mv_{-i}^m}{n-1} - \frac{v_{-i}^m}{n-1} = \frac{1}{n-1} (R(v_{-i}) - \frac{1}{m} R(v_{-i})) \\
 &= \frac{m-1}{m(n-1)} R(v_{-i})
 \end{aligned}$$

This redistribution is slightly smaller than the Bailey-Cavallo redistribution of $\frac{R(v_{-i})}{n}$.

6. NUMERICAL SIMULATIONS

In order to get a better insight into the performance of the two redistribution functions we propose, we implemented and tested them in numerical simulations.

Our simulation set-up is as follows. We consider a time horizon of $T = 100$ discrete time periods, in each of which exactly 1

item is sold. The arrival time of each agent is sampled uniformly between $[1, T]$. An important simulation parameter is the agent’s *patience*, defined as the length of the time window she is active in the market before she has to leave, regardless of whether she acquired her desired item or not. We consider two types of agents: *patient agents*, whose active window is sampled from a uniform distribution $U[1, 20]$ and *impatient agents*, whose active window is sampled uniformly from a distribution $U[1, 4]$.

For the agent’s values⁷ we consider two scenarios: one in which they are sampled uniformly between $(0, 1)$, and one in which they are sampled from a normal distribution $N(\mu = 1, \sigma = 0.2)$, truncated at zero. The total number of agents present in the market varied between $n = 20$ and $n = 500$ in different simulation runs, and each simulation set-up was averaged over 200 runs, with the standard error bars computed in each case.

The performance metric we use is the ratio $\frac{H(\theta)}{R(\theta)}$ averaged over multiple simulation runs. Our experiments compare the performance of two redistribution functions: h^T and h^d .

Several effects can be observed from the results shown in Figure 5. First, the h^T function (Equation 7) that distributes payments at the end of the time interval is highly effective, quickly reaching the expected ratio of over 99%. Moreover, it is not affected much by either the type of valuation distribution or the agents’ patience. This can be explained by the fact that end of period redistribution can distribute to any agent present in the market at any time, and does not depend on the amount of overlap between the agents’ active time windows.

For the function h^d (Equation 9), which redistributes immediately on the departure of each agent, several effects can be observed. First, both the number of agents in the simulation and the patience play a crucial role, with a much better redistribution being achieved for markets with patient agents. This can be explained by a greater overlap between agents’ windows (which allows for more redistribution), and by the fact that the largest fractional payment (which we need to leave undistributed, in order to ensure weak budget balance) plays a smaller role in a larger market. For both uniform and normally distributed valuations scenarios, we see that, with a large enough number of agents, the redistribution ratio tends to stabilize around a certain value. For the uniform valuation case, around 85% of total payments collected can be redistributed in a market with patient agents, but only around 55% in a market with impatient agents. For the normal valuation distribution case, this percentages go up to 94%, and 62%, respectively.

Thus, despite the fact that one can construct a worst-case example where the ratio of payments distributed is 0, our numerical simulations show that these cases are quite rare and that the expected ratio, for both the redistribution functions proposed, when there are more than 100 patient agents in the market, is above 80% when the valuations are uniformly distributed and above 90% when the valuations are normally distributed.

7. CONCLUSIONS AND FUTURE WORK

This work has initiated the study of redistribution in online mechanisms. We first characterized properties of online mechanisms relevant to redistribution and then designed two redistribution functions. The first one generalizes the static mechanism of Cavallo, and redistributes at the last period when the types of all agents are available to compute redistribution. This function is asymptotically optimal (as the number of agents increases), but unsatisfying for settings where agents are not available after they depart,

⁷Note that, unlike time which is discretized, the valuations are real numbers.

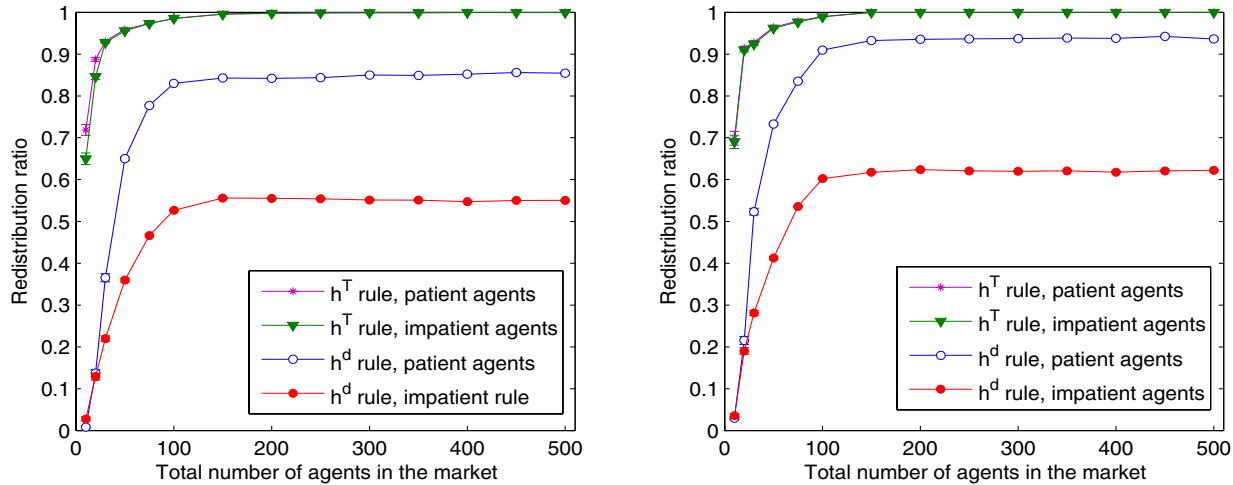


Figure 5: Redistribution amount as a fraction of total revenue for a setting with uniformly distributed valuations (left) and normally distributed valuations (right).

or there is no natural end period. The second function remedies these shortcomings, by making redistribution to each agent on her departure. We do not provide worst-case guarantees for the second function, but evaluate it in numerical simulations. We find this function performs well for various distributions of problem instances redistributing at least one half and in some cases over 90% of the collected revenue.

We know only one other paper studying redistribution in a sequential settings. The work of Cavallo [3] focuses on *dynamic* mechanisms where agents are present for the entire time horizon and it is their types that change from period to period. Importantly, distributional information about future types is available, and an efficient in expectation allocation function is used. In contrast, in the online setting studied here, arrival and departure dates are private information of the agents, and allocation decisions are made at each time period with a greedy rule and without any knowledge about the types of agents arriving in the future.

There are a number of directions for future study. One such direction is to optimize the allocation function along with redistribution function. This has proved fruitful in the static case where inefficient allocations allowed for much higher social welfare when allocating identical items [4, 7]. Another direction involves looking at redistribution in other types of online allocation mechanisms, such as models with items that do not expire, are heterogeneous, or where agents have multi-unit demand. Also, one can study redistribution given a model of future arrivals [13]. Finally, we intend to consider the performance of our online redistribution mechanism in practical applications, such as electric vehicle charging [5, 14].

8. ACKNOWLEDGMENTS

The authors gratefully acknowledge funding from the UK Research Council for project ORCHID (www.orchid.ac.uk), grant EP/I011587/1.

9. REFERENCES

- [1] M. Bailey. The demand revealing process: To distribute the surplus. *Public Choice*, 91(2):107–26, April 1997.
- [2] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *AAMAS’06*, pages 882–889, Hakodate, Japan, 2006.
- [3] R. Cavallo. Efficiency and redistribution in dynamic mechanism design. *EC’08*, pages 220–229, New York, NY, USA, 2008. ACM.
- [4] G. de Clippel, V. Naroditskiy, M. Polukarov, A. Greenwald, and N. R. Jennings. Destroy to save. *Games and Economic Behavior*, 2013.
- [5] E. Gerding, V. Robu, S. Stein, D. Parkes, A. Rogers, and N. Jennings. Online mechanism design for electric vehicle charging. In *AAMAS’11*, pages 811–818, 2011.
- [6] J. R. Green and J.-J. Laffont. *Incentives in public decision-making*. North Holland, New York, 1979.
- [7] M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions with unit demand. In *EC’08*, pages 210–219, 2008.
- [8] M. Guo and V. Conitzer. Worst-case optimal redistribution of vcg payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69 – 98, 2009.
- [9] M. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. C. Parkes. Online auctions with re-usable goods. In *6th ACM Conf. on El. Commerce (EC’05)*, pages 165–174, 2005.
- [10] H. Moulin. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
- [11] V. Naroditskiy, M. Polukarov, and N. Jennings. Optimal payments in dominant-strategy mechanisms for single-parameter domains. *ACM Transactions on Economics and Computation*, 2013.
- [12] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–241. Cambridge University Press, 2007.
- [13] D. C. Parkes. Online mechanisms. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, pages 411–439, 2007.
- [14] V. Robu, S. Stein, E. Gerding, D. Parkes, A. Rogers, and N. Jennings. An online mechanism for multi-speed electric vehicle charging. In *AMMA’11*, LNICST, pages 100–112. Springer, 2012.